Notes on Heptagonal Cusps and the Closing Problem

August 2026

Abstract

These notes concern the arithmetic classification of trace-field directions for the heptagonal triangle group and the geometric height problem behind it.

WARNING: this was created by OpenAI.

Foreword

This volume has two complementary parts. Part I is a systematic account of the heptagonal cusp problem. It states the conjecture, develops the geometric and arithmetic structures that bear on it, records the reductions that have been proved, and isolates the estimate that is still missing. It is intended to be read as a mathematical manuscript: definitions and results are arranged by logical dependence, and obsolete approaches are separated from the active line of argument.

Part II is a research archive. It preserves the calculations, experiments, conditional arguments, counterexamples, and changes of direction from which the organized account emerged. Its value is different. It shows how the obstructions were found, supplies exact computational certificates, and records why several natural ideas do not close the argument. Statements in that part must be read with their local verdicts and subsequent corrections; the later audit of an approach may supersede the optimism with which it was first introduced.

The conjecture itself remains open. The established results give a precise closing criterion and a substantial geometric reduction, but they do not yet show that every rational trace-field direction must close. The distinction between theorem, conditional implication, numerical evidence, and proposed mechanism is maintained throughout the volume.

The opening orientation is independent of the detailed proof ledger. Its purpose is to explain the question, why it is natural, what is already known, and what kind of synthesis appears necessary before the technical accounts begin.

Orientation and Background

1. A geometric classification problem

The hyperbolic plane admits tessellations whose tiles are triangles. A particularly rigid family begins with a triangle having angles $\pi/2$, π/q, and zero. The zero angle means that one vertex lies on the boundary of the hyperbolic plane. Reflections in the sides generate a triangle group, and the orientation-preserving half of that group acts by fractional linear transformations. The quotient has finite hyperbolic area, two cone points, and one cusp.

This volume concerns the heptagonal member of the family. Set

$$\lambda = 2\cos(\pi/7), \qquad K = \mathbf{Q}(\lambda),$$

and let

$$S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \qquad T = \begin{pmatrix} 1 & \lambda \\ 0 & 1 \end{pmatrix}.$$

The group

$$G_7 = \langle S, T \rangle \subset \mathrm{PSL}_2(\mathbf{R})$$

is the orientation-preserving triangle group of signature $(2, 7, \infty)$. The transformation T fixes infinity and translates along a horocycle there. Its fixed point is the distinguished cusp of the quotient. Every translate of infinity by an element of G_7 is another representative of the same cusp on the boundary.

The field K is a totally real cubic number field. Since the matrices of the group have entries in its ring of integers, the group acts on the projective line $\mathbf{P}^1(K)$. This produces a concrete classification question: which points of $\mathbf{P}^1(K)$ occur in the orbit of the cusp, and what accounts for those that do not?

There are two evident mechanisms by which a point of $\mathbf{P}^1(K)$ can be distinguished by the group. A parabolic element has one boundary fixed point; its fixed points are precisely the cusps. A hyperbolic element has two boundary fixed points. Occasionally those eigenlines are defined over K, and then they too give arithmetic points on the boundary. The first kind is analogous to a rational number with a finite continued fraction. The second is analogous to a quadratic irrational with a periodic continued fraction.

The problem is to show that these familiar alternatives exhaust the field points. In other words, arithmetic origin is expected to force dynamical recurrence. A point defined over the trace field should not be able to wander forever without either reaching the cusp or returning to the same complete state.

2. The conjectural dichotomy

The central conjecture can be stated without any height theory.

Conjecture 0.1 (Heptagonal closing conjecture)**.** *Every point of $\mathbf{P}^1(K)$ is either a cusp of G_7, or belongs to the G_7-orbit of a K-rational eigenline of a hyperbolic element of G_7.*

The phrase "belongs to the orbit" is important. A continued-fraction expansion may have an initial nonrepeating segment before it enters a periodic cycle. Acting by the matrices represented by that initial segment moves the original point onto the eigenline responsible for the cycle. Thus the geometric statement naturally includes eventual periodicity, not only periodicity from the first step.

The two alternatives are genuinely different. A cusp is fixed by a parabolic element, while the special points in the second class are fixed by hyperbolic elements. Discreteness prevents a nontrivial parabolic and a hyperbolic element from sharing a boundary fixed point. Nor is the hyperbolic alternative an artificial exception inserted to save the statement. Explicit examples occur already among very small elements of the group, and computations reveal numerous distinct special orbits.

There are several equivalent ways to hear the conjecture.

- In hyperbolic geometry, it says that every trace-field boundary point is stabilized, after translation by the group, by a nontrivial element.
- In continued-fraction language, it says that every trace-field input has a finite or eventually periodic Rosen expansion.
- In arithmetic-group language, it says that the thin group meets every relevant rational Borel subgroup nontrivially.
- In the geometry of translation surfaces, it says that every trace-field direction is either parabolic or governed by a special affine automorphism.

These formulations emphasize different structures, but the desired conclusion is the same exact event: a determinant vanishes, a complete state repeats, or a stabilizer appears. Approximate recurrence is not enough. Any proof must eventually convert an inequality into an algebraic equality.

Historical context

The question belongs to a line of work on the arithmetic of Hecke triangle groups and their continued fractions. Rosen gave the foundational arithmetic characterization of parabolic points for the pentagonal group [18], and Rosen and Schmidt developed the continued-fraction coding for Hecke groups more generally [26]. These results established the model in which cusp membership is tested by termination of a group-adapted reduction algorithm.

Hanson, Merberg, Towse, and Yudovina studied generalized continued fractions and projective orbits for Hecke triangle groups directly [62]. Their work includes arithmetic obstructions to cusp membership and orbit predictions that helped make the trace-field classification problem explicit. The computations now available for the heptagonal group show a richer special hyperbolic spectrum than a two-orbit picture, but the organizing question is the same one: which trace-field points are detected by the Hecke-group action?

The translation-surface literature supplies a related, but importantly different, notion of periodicity. Arnoux and Schmidt proved the existence of trace-field directions that are not periodic cylinder directions on Veech surfaces [7]. A hyperbolic eigenline can already be nonperiodic in that

cylinder sense while being periodic as a Rosen or geodesic state. Their theorem therefore motivates the distinction between cusps and special hyperbolic directions; it does not contradict the closing conjecture stated here.

McMullen introduced the Hodge-height method that provides the principal geometric model for the present approach [17]. His argument connects billiards, trace-field arithmetic, and contraction of Hodge norms in a setting where the unit geometry is effectively one-dimensional. The heptagonal problem asks whether that architecture can be enlarged to retain two independent unit directions without losing integrality.

Most recently, Winsor has investigated special hyperbolic elements in Hecke triangle groups and carried out extensive computations for the heptagonal case [64]. Those computations exhibit many special orbits and find termination or eventual periodicity throughout the tested heptagonal sample. They provide strong evidence for the dichotomy while also showing that the hyperbolic alternative is a substantial part of the answer.

3. Why one should expect a continued-fraction theorem

For the modular group $\mathrm{PSL}_2(\mathbf{Z})$, ordinary continued fractions provide a perfect dictionary between arithmetic and dynamics. Rational numbers have finite expansions, real quadratic irrationals have eventually periodic expansions, and the associated matrices record motion through the modular tessellation. The statement for G_7 asks for an analogue in a setting where the coefficient field is cubic and the group is far smaller than the full arithmetic group over that field.

The appropriate algorithm is the Rosen continued fraction. It repeatedly uses a translation by an integral multiple of λ, followed by inversion, to return a real number to a fixed interval. Each digit records a cusp excursion in the triangle tessellation. A finite expansion reaches a boundary state representing infinity. Repetition of the complete projective state produces a group element fixing the corresponding point; see [26, 62] for the Hecke-group continued-fraction setting.

The elementary direction of the conjecture is therefore clear. If the algorithm terminates, the original point is a cusp. If a complete state repeats, the intervening matrix fixes the point; except in the terminating case, this matrix is hyperbolic. The difficulty is the converse: why should membership in a cubic number field prevent an infinite aperiodic sequence of digits?

One cannot answer this by counting field elements. Both K and the group are countable. One cannot answer it from the abstract presentation of the triangle group either. The conclusion depends simultaneously on the exact integral matrices, the three real embeddings of their coefficients, and the geometry that couples those embeddings. The continued fraction supplies the states on which these structures must be compared; it does not by itself provide the comparison.

The correct state is richer than the current real number. Homogeneous coordinates remember a pair of algebraic integers, and consecutive pairs remember a determinant. This affine information distinguishes genuine return from projective resemblance. It also records finite divisibility, which is invisible if the orbit is observed only at the distinguished real embedding. Much of the later machinery can be understood as an effort to choose the smallest complete state that retains all of this information.

4. The classical analogy and its limits

Ordinary continued fractions begin with a remarkably simple operation. From a real number one subtracts its integer part and then takes a reciprocal. The result is again placed in a fixed interval, and iteration produces a sequence of positive integers. Behind this elementary process is the action of $\mathrm{PSL}_2(\mathbf{Z})$ on the upper half-plane. Subtraction and inversion are its standard generators, and the continued-fraction digits record the sequence of tiles crossed by a geodesic.

Two classical theorems make this coding arithmetic. A real number has a finite continued fraction exactly when it is rational. It has an eventually periodic continued fraction exactly when it is a real quadratic irrational. The proofs use different features of the algorithm, but together they show a striking principle: low algebraic complexity is recognized by recurrence in a geometric reduction process.

For a rational number, Euclid's algorithm decreases a positive integer denominator until it reaches zero. For a quadratic irrational, there need not be monotone descent. Instead, the conjugate bounds the complete quotients, leaving only finitely many integral states. One of those states must repeat, and the matrix accumulated around the cycle fixes the original quadratic irrational. Thus even classically the periodic case is proved by finiteness rather than by strict descent at every step.

The heptagonal conjecture asks for a cubic analogue with a different balance between field and group. The Rosen algorithm still has subtraction and inversion, with multiples of λ replacing ordinary integers. Its geodesic meaning is equally concrete. But the field K has degree three, while a fixed point of an arbitrary matrix over K would usually have degree two over K. A K-rational fixed point is therefore a special degeneracy of the matrix, not the generic source of a cubic number.

The group also differs sharply from the modular group. The modular group is the full group of integral determinant-one matrices, whereas G_7 is thin in the corresponding arithmetic group over K. A Euclidean reduction in the ambient ring may use matrices that do not belong to G_7. It is not enough to show that a pair of algebraic integers can be simplified by some integral operation; the simplifying operations must follow the exact heptagonal word.

These differences prevent a direct imitation of Lagrange's theorem, but the classical logic remains instructive. One seeks a complete integral state, a global bound that leaves only finitely many such states, and an exact interpretation of repetition. The present program follows this logic. Its new ingredients are the three real embeddings, the Hodge metrics supplied by the translation surface, and a rank-two arithmetic state capable of remembering both unit directions.

The analogy also warns against demanding too much. A periodic ordinary continued fraction can have very large partial quotients before its state repeats. Likewise, a special heptagonal orbit can make a large cusp excursion or acquire a badly unbalanced unit representative before it closes. A successful height need not be small at every instant. It must control the number of genuinely new complete states, or charge a large excursion only once.

5. What closing means exactly

It is useful to separate three levels of recurrence. A real orbit may return close to a previous point in the reduction interval. A projective arithmetic state may return exactly after multiplication of

its homogeneous coordinates by an element of $K^\times$. A complete affine state may return with its integral frame and determinant data. Only the latter two can produce the algebraic conclusion, and the distinction between them is controlled by units.

Suppose a Rosen prefix is represented by a matrix $g \in G_7$. If the current projective state equals the initial state, then g preserves the corresponding line in K^2. Thus g lies in the K-rational Borel subgroup stabilizing that point. If the orbit has not terminated, the geometry of the triangle group rules out a nontrivial elliptic or parabolic return, and one obtains a hyperbolic element with the given eigenline.

If a remainder vanishes instead, the corresponding matrix sends the original point to infinity. Reversing the prefix shows that the point lies in the cusp orbit. These deductions are exact and elementary. They explain why the conjecture is often called a closing problem: once an exact return or zero has been produced, no further analytic argument is needed.

Approximate return is subtler. Hyperbolic dynamics naturally supplies shadowing and recurrence statements, but a line can pass arbitrarily close to itself without being defined by a stabilizer. Diophantine approximation can show that certain coincidences are rare, yet rarity is not the same as impossibility along one prescribed algebraic orbit. A proof must insert integrality at the final step so that sufficiently strong approximation forces an actual zero determinant.

This exactness requirement is why the affine frame matters. Two projective vectors may define the same point while differing by a large unit. In a single real embedding they may look close even when their conjugates and integral lattices are far apart. The determinant of two consecutive vectors records both incidence and normalization. It is the arithmetic bridge between geometric closeness and equality.

There is also a group-theoretic restatement. Every point $x \in \mathbf{P}^1(K)$ determines a Borel subgroup B_x of SL_2, consisting of matrices preserving the line x. The conjecture says that $G_7 \cap B_x(K)$ is nontrivial for every such x. A unipotent element in the intersection gives the cusp alternative; a semisimple element gives the hyperbolic alternative. This formulation removes the continued-fraction coordinates, but it does not remove the difficulty: proving that a thin subgroup meets every one of these rational stabilizers is a strong rigidity assertion.

The different formulations therefore serve different stages. Rosen dynamics chooses a canonical sequence approaching the stabilizer. Hodge geometry measures that sequence at the real places. Arakelov geometry keeps the determinant integral. The Borel formulation identifies the exact event that ends the proof.

6. The arithmetic landscape

The number λ satisfies

$$\lambda^3 - \lambda^2 - 2\lambda + 1 = 0.$$

Its three conjugates are real. Consequently every element of K has three real sizes, one for each embedding $\sigma_i : K \hookrightarrow \mathbf{R}$. The first embedding is the one in which $\lambda = 2\cos(\pi/7)$ and G_7 acts discretely on the hyperbolic plane. The other two embeddings send the same matrix entries to different real numbers. They do not produce two further copies of the same triangle tessellation,

but they carry indispensable information about arithmetic size.

The distinction between reduction in the full ring of integers and reduction inside the triangle group is important here. Euclidean questions for septic cyclic fields have been studied by Lezowski and McGown [13]; those ambient results illuminate the arithmetic of K, but they do not force the reducing matrix to lie in the thin subgroup G_7.

If $[p : q] \in \mathbf{P}^1(K)$, its height combines the sizes of p and q at all three embeddings and at all finite primes. A central principle of Diophantine geometry is that these contributions must be combined globally. Multiplying (p, q) by a nonzero element changes every local size, but the product formula makes the total height unchanged.

Units make this point vivid. A unit has no finite prime divisors, yet its three real absolute values can be extremely uneven. Their product is one, so expansion at one embedding is paid for by contraction at the others. This means that a proposed descent function can look excellent at the geometric embedding and fail badly after multiplication by units, even though the projective point has not changed. A valid height must either be invariant under this redistribution or keep an exact account of it.

The field has unit rank two. That simple fact is one source of the difference between quadratic and cubic trace fields. In a real quadratic field there is essentially one direction of unit imbalance. In a totally real cubic field there are two. A scalar comparison that controls one direction may leave another completely free. This is why vector-valued or rank-two objects appear naturally in the heptagonal problem.

Finite primes create the complementary difficulty. A remainder in the continued-fraction recurrence is an algebraic integer, and its norm records all of its prime divisors at once. It is tempting to follow those prime divisors individually. But new primes can enter at successive stages, so a prime-by-prime loss can migrate indefinitely. The product formula suggests the opposite viewpoint: regard the entire remainder ideal as one global divisor and compare it in a single height calculation with all three real metrics.

Thus the arithmetic setting points toward Arakelov geometry even before any technical construction is made. Arakelov language packages an integral lattice at every finite place together with a norm at every real place. Its degree is independent of the choice of basis because of the product formula. For this problem, that is exactly the bookkeeping one wants: one finite divisor, three archimedean behaviors, and no possibility of paying for the same denominator more than once.

7. Heights as global memory

A height is often introduced as a measure of the size of rational or algebraic coordinates. In this problem it is better understood as global memory. The Rosen map repeatedly changes coordinates. A quantity observed at only one real embedding forgets how those changes redistribute size among the conjugates and finite primes. A height remembers the total redistribution because the product formula ties all places together.

Take a projective point $[p : q]$ and choose integral homogeneous coordinates. At a real embedding σ_i, one can measure $\max(|\sigma_i(p)|, |\sigma_i(q)|)$. At a finite prime, one measures the common divisibility of the coordinates. Multiplying both coordinates by $a \in K^\times$ shifts every one of these local

measurements. The sum of the shifts is zero. The resulting global height therefore depends on the point, not the representative.

Northcott's theorem says that points of bounded degree and bounded height form a finite set. This is precisely the sort of compactness needed for a closing argument. If all complete Rosen states could be represented by points of uniformly bounded height, repetition would follow. The obstacle is that the ordinary projective point is not a complete state, while the natural affine state can acquire large unit imbalance even on an orbit that eventually repeats.

One may balance the three real sizes of the initial coordinates by multiplying them by a suitable unit. Dirichlet's unit theorem makes this possible because the logarithms of units form a lattice in the plane where the three coordinates sum to zero. But this balancing can be done only once. Choosing a new unit after every Rosen step would alter the affine recurrence and erase the very data needed to recognize repetition.

For that reason, the initial height should be viewed as a budget rather than a uniform bound on every later coordinate. The orbit may borrow heavily in one real direction, provided the global accounting detects when the debt has already been represented elsewhere. A stopped height implements this idea by measuring only at first-hitting times for an arithmetic record. Between two records, large archimedean motion may be holonomy that eventually cancels; it does not create a new finite obligation.

This viewpoint is close in spirit to descent arguments in Diophantine geometry. One does not insist that every local expression decrease. Instead, one constructs a global degree whose positivity has an arithmetic meaning. If the degree were to increase too often, a fixed initial point would produce infinitely many independent integral conditions. Geometry is expected to rule this out.

The relevant global degree cannot be merely the norm of one remainder. Field norms are blind to units, and units are exactly where large intermediate excursions hide. Nor can it be merely a Hodge norm, because a Hodge norm has no finite integrality by itself. The desired object must carry the remainder ideal as its finite lattice and the Hodge comparisons as its real metrics. That combination is the essence of the Arakelov formulation.

8. Thinness and the local–global tension

The group G_7 sits inside an arithmetic matrix group over the ring of integers of K, but it has infinite index there. At the same time it is Zariski dense in the relevant algebraic group. Such a subgroup is called thin. Thinness is not merely a label for small index; it explains why the problem lies between geometry and arithmetic.

The ambient arithmetic group has many tools for moving one primitive vector to another. Strong approximation describes its reductions modulo ideals, and reduction theory controls its real fundamental domains. A thin subgroup can have very large images modulo almost every prime while still missing most integral matrices. Congruence information therefore gives necessary local conditions but does not automatically produce the required global element.

The conjecture asks for an intersection with a stabilizer. Given a point $x \in \mathbf{P}^1(K)$, all matrices preserving its line form a Borel subgroup. The desired conclusion is that the thin group meets this Borel in something other than the identity. Zariski density does not imply such an intersection: a

dense subgroup can avoid a prescribed lower-dimensional subvariety. Nor does surjectivity modulo many primes force an exact characteristic-zero intersection.

On the other hand, the group is not arbitrary. It is the monodromy group of a geometric family and the Veech group of a translation surface. Its words carry Hodge-theoretic contraction, and its parabolic subgroup is tied to a genuine degeneration of that family. The plausible route is therefore not a generic theorem about thin groups. It is a rigidity theorem for this particular thin orbit, using the geometric metrics to turn approximate incidence into exact incidence.

There is a useful division of labor. Symbolic dynamics selects canonical prefixes and stopping times. Arithmetic records their integral content. Hyperbolic and Hodge geometry provide contraction at the real places. Diophantine finiteness turns a bounded global state into repetition. None of these ingredients alone appears sufficient, but their interfaces are rigid enough that the missing assertion can be stated quite precisely.

9. Translation surfaces and the source of extra geometry

The triangle group is also realized as the symmetry group of a translation surface built from two regular heptagons. Parallel sides are identified by translations, producing a compact surface of genus three equipped with a holomorphic one-form. Applying a real linear map to the polygons deforms the surface. The resulting orbit in the moduli space of Riemann surfaces is a Teichmuller disk, and quotienting by the stabilizer produces the same hyperbolic orbifold that arose from the triangle.

This realization adds a family of Jacobians to the problem. The Jacobian of each genus-three surface has real multiplication by K. Its first cohomology consequently splits into three rank-two pieces, one for each real embedding of K. Each piece contains a line of holomorphic forms, and that line carries a natural Hodge norm.

As the surface moves along a geodesic, these three norms change at different rates. The distinguished piece is the tautological one: it directly records the stretching and contracting directions of the polygons. The two conjugate pieces vary more slowly. Holomorphic period maps from the triangle disk to the upper half-plane express this comparison. By the Schwarz–Pick theorem, the conjugate period maps cannot expand hyperbolic distance, and in the interior they contract it strictly.

For the heptagonal family the normalized rates are

$$1, \qquad \frac{3}{5}, \qquad \frac{1}{5}.$$

They are geometric invariants, not constants fitted to numerical data. Their inequality says that motion visible in the triangle tessellation dominates the corresponding conjugate motion. Combining all three pieces gives a positive margin. This is the most compelling source of rigidity presently available. The Teichmuller-curve and modular-embedding calculations behind these rates are developed by Bouw–Moller and Moller–Zagier [7, 40].

The challenge is that Hodge contraction is naturally stated for a flat cohomology class transported along the family, whereas the continued-fraction algorithm continually changes the remainder or

covector whose finite divisor is relevant. The object with the correct real contraction and the object with the correct integral content are close, but they are not automatically the same. Their discrepancy is an incidence or angle term. Controlling that term at the right stopping times is the central geometric difficulty.

10. The geometry of a cusp excursion

A finite-area hyperbolic surface is not compact. A geodesic may spend a long time high in a cusp before returning to the thick part. In the Rosen coding, such an excursion appears as a large digit or a long run of parabolic letters. This is not an exceptional defect of the coding; it is the correct geometric description of motion near infinity.

Parabolic motion behaves differently from compact motion. A hyperbolic element has positive translation length and distinct expanding and contracting eigenlines. A parabolic element has translation length zero and only one fixed line. Under the three real embeddings, a large parabolic power has roughly the same logarithmic size in every component. Consequently the Hodge advantage is not accumulated uniformly during the deepest part of a cusp excursion.

The advantage appears when the orbit passes through the thick part, where the conjugate period maps contract strictly. A long word can therefore be decomposed conceptually into neutral cusp motion and charged compact returns. This decomposition is visible both in the hyperbolic geodesic and in the finite Rosen alphabet. It explains why estimates for individual generators are too crude: the amount of useful contraction depends on how the entire word enters and leaves the cusp.

The cusp also controls integrality. Near a degeneration, Hodge norms have logarithmic singularities whose coefficients can be computed. Those singularities match the parabolic factors in the recurrence. An adapted metric can therefore subtract the predictable cusp growth and leave a bounded transition term. This is the geometric reason to hope that very large digits are not intrinsically dangerous.

The dangerous event is a junction between excursions. The moving affine frame may change orientation, and a unit accumulated in one corridor may no longer cancel in the next. Explicit examples show that such junctions can create large intermediate split-character height even on an orbit that later terminates or becomes periodic. Any global inequality that places a fixed ceiling on every junction is therefore stronger than the classification.

A stopped argument treats the junction differently. It asks whether the junction creates a genuinely new finite record. If it does not, its large real size can be absorbed into previously recorded arithmetic data. If it does, the Hodge surplus accumulated since the previous record should pay for the new determinant divisor. The required inequality is thus naturally attached to an interval of the orbit, not to one generator or one current coordinate.

This is also why periodic points are compatible with contraction. Along a periodic orbit, a chosen representative may become enormous before a unit returns the projective state to its starting value. The total Hodge degree around the closed cycle is constrained, but no single coordinate need remain small. Equality in the global argument should recognize this unit holonomy as the hyperbolic alternative rather than attempt to eliminate it.

11. McMullen's height argument as the model

McMullen's work on nonarithmetic groups supplies the guiding precedent. In a real quadratic trace field, the cohomology splits into two conjugate rank-two systems. An integral class has one projection in each system, and the product of their Hodge norms is insensitive to multiplication by a unit: the two real absolute values of a unit multiply to one. Once the finite lattices are included, this product behaves like a genuine arithmetic height [17].

Point a geodesic toward a trace-field boundary point. In the distinguished component, the covector defining that point contracts strongly because the geodesic is aimed at its kernel. The conjugate component can grow, but the period map contracts hyperbolic distance. The conjugate growth is therefore too small to cancel the distinguished decay. The global Hodge height tends downward.

An arithmetic height cannot decrease indefinitely through infinitely many distinct integral states of bounded complexity. Northcott finiteness then forces recurrence, and recurrence gives a stabilizing group element. This is the ideal architecture for the present problem: geometric contraction creates a global height descent, arithmetic discreteness turns descent into a finite set, and dynamics turns repetition into a cusp or hyperbolic fixed point.

Several features make the argument especially clean in the quadratic case. There is only one conjugate component to compare with the distinguished one. The unit group has only one logarithmic direction. The product of the two Hodge norms is already the natural Galois-invariant scalar. Finally, the flat class transported by the local system can be kept aligned with the arithmetic state used in the height.

The heptagonal problem retains the overall architecture but breaks each of these simplifications. There are two conjugate components and two unit directions. A product that is Galois invariant can average away the direction one needs to control. More importantly, the exact Rosen recurrence selects a changing affine frame. The flat endpoint line has clean Hodge transport, but the moving remainder line carries the finite divisor. Passing between them has a cost, and an orbit can accumulate large intermediate archimedean size even when it later closes.

The lesson is not that height methods are irrelevant. It is that the height must be attached to a richer state and evaluated at selected times. A scalar that decreases at every step is too rigid to accommodate periodic unit orbits. A scalar ceiling on every intermediate state is likewise stronger than the conjecture requires. The natural replacement is a stopped global height: one waits until an arithmetic quantity first reaches a new record and then compares all finite and infinite contributions in a single degree.

12. Why the most natural shortcuts stop short

The shape of a likely proof becomes clearer when compared with several standard approaches. Each sees a genuine part of the problem, but each forgets one piece of the complete state.

The first idea is ordinary descent. Find a norm on homogeneous coordinates that decreases under every nonterminal Rosen step, and conclude by well-ordering. This works for Euclidean algorithms but cannot accommodate the special hyperbolic alternative. A periodic state would

have to return to itself after a sequence of strict decreases. Any successful descent must therefore permit equality along a nontrivial unit holonomy or operate only at selected stopping times.

The second idea is to use a scalar height at the distinguished embedding. Hyperbolic geometry gives strong contraction there, especially for a covector aimed at the endpoint. But a scalar at one embedding is not projectively well defined over K. Multiplication by a unit can change it arbitrarily while leaving the point fixed. Adding the other embeddings is necessary, and doing so introduces two independent conjugate directions.

The third idea is a direct application of Roth's theorem or the Subspace Theorem. Consecutive Rosen convergents do approximate the endpoint, and Diophantine approximation is designed to limit unusually good rational approximations to algebraic numbers. The target here, however, moves with the prefix, and its height grows at the same critical scale as the approximation. Standard fixed-target finiteness does not retain the thin-word relation or the shared denominator at the required strength.

The fourth idea is congruence detection. Reduce the orbit modulo a prime or an ideal, use expansion of the finite action, and seek a return to the stabilizer. The finite images of a thin group can be very large, and this is valuable structural information. Yet equality modulo an ideal is only divisibility of a determinant. To deduce an actual zero, the modulus must exceed an archimedean bound for that determinant. Producing such a lift with the right Hodge shape is already a quantitative form of the closing problem.

The fifth idea is to track every prime divisor of the remainders. Local recurrences give exact information about how a fixed prime can persist or return. The difficulty is migration: a different prime may be responsible at each new stage. Summing local estimates separately can charge indefinitely many primes without contradiction. A global divisor degree is needed to couple all finite places at once.

The sixth idea is to multiply all Hodge components into one scalar. This has the advantage of Galois invariance and honest integrality. It also averages the two unit directions. The best real direction may change with the state, and the scalar product can cancel the information needed to select it. Two pairwise products retain the favorable contractions, but treating them as two independent heights duplicates their common finite denominator.

The seventh idea is to use a spectral gap for the group action. Uniform expansion in congruence quotients or contraction of transfer operators can control averages over many words. The conjecture, by contrast, concerns the one canonical word selected by each algebraic point and asks for an exact stabilizer. An average theorem must be supplemented by an inverse theorem showing that every persistent exceptional trajectory is a rational Borel orbit. That inverse statement is close to the desired rigidity itself.

These limitations do not make the corresponding theories irrelevant. Roth methods suggest how approximation and height should balance. Congruence methods describe the finite lattice and may support a bilinear amplification. Spectral and transfer methods quantify the thick-part Hodge gain. Unit equations provide the final finiteness. The point is that they must act on one common framed object rather than be applied as independent black boxes.

13. From scalar descent to a rank-two state

The three Hodge components suggest several scalar products, and each captures part of the desired contraction. Comparing the distinguished component with either conjugate component gives a favorable pairwise slope. But choosing the better pair separately at every stage risks charging the common denominator twice. Taking the full product avoids that duplication but can lose the ability to select the effective direction of unit imbalance.

This tension points toward rank two. Instead of forcing one scalar formula to make every choice in advance, one keeps two channels in a common integral lattice. At the real places, an analytic argument shows that at least one channel retains a definite portion of the Hodge contraction. At the finite places, both channels share a determinant ideal, which must enter the degree only once. A maximal-slope principle can then select the larger real gain without duplicating the finite cost.

The word "slope" here has a concrete meaning. An adelic vector bundle is a finite-dimensional vector space with integral lattices at finite primes and norms at real places. A line inside it has a global degree. The maximal slope is the largest normalized degree among its subspaces. In rank two it is the coordinate-free version of choosing the more favorable of two directions while keeping their common lattice fixed.

The relevant two directions arise from consecutive Rosen states. Their determinant measures the failure of the two affine frames to coincide. At the real places their ratio is an explicit cross ratio, or equivalently a Green-function displacement between endpoint lines. At the finite places the same determinant controls the common denominator. This agreement is why the rank-two formulation is more than a convenient repackaging: it is the first object in which the analytic choice and the arithmetic cost visibly refer to the same data.

A successful proof would establish a stopped inequality for this bundle. At the first time a remainder norm reaches a new record, the Hodge gain in one of the two channels would dominate the degree of the determinant lattice, up to a constant depending only on the initial point. Repeated records would then be impossible unless the affine frame entered an equality corridor. In that corridor the exact recurrence would force either a vanishing remainder or a repeated frame.

This formulation also explains why several weaker ideas do not suffice. Controlling only operator norms ignores the evaluated rational line. Tracking only the projective point forgets the determinant and the unit coordinate. Bounding prime divisors separately allows losses to migrate. Applying a scalar maximum after the fact risks paying the shared denominator twice. The rank-two stopped bundle is designed precisely around these failures.

14. Record times and the equality corridor

Stopping at record times is a mathematical choice, not merely a way to make plots easier to read. Let a remainder norm be called a record when it exceeds all preceding remainder norms. If these records are bounded, then every remainder norm is bounded. If they are unbounded, the first transition creating each record isolates the exact point at which new finite arithmetic complexity enters the orbit.

The transitions between records may contain long parabolic runs and large unit excursions.

Their internal coordinates can fluctuate without creating new finite content. By evaluating a global degree only at the first-hitting transition, one asks the Hodge geometry to pay precisely for the increase that has arithmetic significance. This is analogous to optional stopping in probability, but the desired estimate is deterministic and arithmetic.

At a record transition there are two possibilities. The Hodge surplus is strict and controls the new divisor, or the orbit lies in an equality corridor. The equality corridor should not be treated as a negligible error set. It contains the periodic alternative that the conjecture predicts. The proof must classify it exactly.

Geometrically, equality means that a rational line remains aligned with a distinguished boundary or Levi direction under transport. Arithmetically, it means that the unit shape absorbs the expected contraction without producing a new divisor. Dynamically, it should force recurrence of the complete Borel frame. Once that frame repeats, the intervening group element fixes the original line.

This trichotomy is the intended endgame. Strict positivity bounds a new record. A vanishing incidence gives a cusp. Neutrality with a repeated affine frame gives a hyperbolic fixed point. These are not three unrelated outcomes: they are the positive, degenerate, and equality cases of one global degree comparison.

The rank-two state is particularly well suited to this endgame because its maximal slope can select a positive channel while its determinant detects degeneracy. When the two directions collapse, the determinant vanishes. When they remain distinct but the degree is neutral, their relative unit coordinate records the possible hyperbolic holonomy. Thus the same object contains both the inequality and its equality classification.

15. The role of Diophantine finiteness

Height contraction is not expected to prove closing by itself. Its role is to place the exact continued-fraction states in a finite or finitely generated arithmetic configuration. A second ingredient must then turn that bounded configuration into repetition.

The most useful finiteness theorem here concerns linear equations in units. The heptagonal field has a special splitting identity associated with two adjacent parabolic directions. Applied to consecutive homogeneous remainders, it produces four multiplicative quantities satisfying a fixed linear relation. If the norms, or more generally the finite support, of these quantities remain bounded, they lie in a fixed finitely generated multiplicative group. The projective S-unit theorem says that a linear equation has only finitely many nondegenerate solutions in such a group.

This yields a rigorous conditional closing result. Once the appropriate split norms are bounded along a complete orbit, only finitely many projective states can occur; the Rosen recurrence must terminate or repeat. The remaining issue is therefore not the final finiteness step. It is the global estimate needed to obtain the bounded arithmetic state from Hodge geometry.

This division is conceptually useful. The Hodge side need not identify the period or the stabilizing element. It need only prevent indefinite growth of the complete arithmetic state. Conversely, the S-unit theorem need not understand hyperbolic geometry. It receives bounded support and returns finiteness. The exact recurrence supplies the glue between the two.

The equality cases matter. A vanishing term in the unit equation corresponds to termination at

a cusp. A repeated nonvanishing projective state corresponds to a hyperbolic eigenline. Thus the Diophantine theorem naturally returns exactly the two alternatives in the conjecture rather than a larger collection of exceptional behaviors.

16. What computation contributes

Exact computation has played three distinct roles. First, it verifies algebraic identities and finite certificates that would be tedious to check by hand. Matrix products, field norms, branch decisions, and projective returns can all be evaluated exactly in $\mathbf{Z}[\lambda]$. Such calculations can legitimately support lemmas whose remaining content is finite.

Second, computation is an adversarial test for proposed heights. Random samples are useful, but structured unit families are far more revealing. They can keep finite support small while moving far in the archimedean unit lattice. A formula that appears monotone on generic inputs may fail sharply on these rays. Several attractive scalar bounds were abandoned for exactly this reason.

Third, computation maps the dynamical landscape. Large collections of trace-field inputs have been observed to terminate or become periodic, and the periodic alternatives occur in many distinct orbits. These experiments support the formulation of the conjecture and prevent an oversimplified classification. They do not prove that an orbit that has run for many steps will eventually close.

The distinction between a finite certificate and finite evidence is crucial. An exhaustive check of a finite alphabet inequality can be a proof when a separate argument reduces every length to that finite check. Observing no counterexample among a million initial states is evidence only. Likewise, an exact table showing the failure of a proposed inequality disproves that finite claim but may not establish an asymptotic theorem about the family.

The research archive records these roles in detail. It includes successful certificates, diagnostic searches, and routes that were withdrawn after a targeted counterexample. The organized account retains only those computations that support a current statement or clarify the precise shape of the obstruction.

17. The expected synthesis

The evidence points toward a proof with four layers.

First, the Rosen algorithm supplies a canonical sequence of integral affine states and selects the moments immediately preceding a new norm record. It also gives the exact algebraic identities that interpret termination and repetition.

Second, the geometry of the heptagonal Teichmuller curve supplies unequal Hodge growth at the three real embeddings. Strict contraction of the conjugate period maps creates a definite surplus whenever an orbit spends time away from the cusp. Cusp excursions are neutral rather than harmful; their contribution must be normalized by the exact parabolic recurrence.

Third, an Arakelov construction must place the two surviving real channels and the common determinant ideal in one rank-two object. A stopped maximal-slope inequality should retain the larger Hodge gain while counting the finite divisor once. This is the point where a new argument is needed. It may draw on arithmetic holonomy, relative height inequalities, Bott–Chern transgression,

or a coefficient-sensitive form of a slopes method, but it must respect the exact thin-orbit framing.

Fourth, once the resulting global state is bounded, the split-character identity and the projective S-unit theorem force finiteness. The exact Rosen dynamics then turn finiteness into termination or periodicity, which is the desired classification.

There are two conceptual tests for any proposed replacement. It must be invariant under changing homogeneous coordinates by a unit, or explicitly account for both unit directions. And it must attach finite and infinite data to the same moving state. A real inequality for one line and a finite ideal attached to another do not form a global height merely because their formulas look parallel.

The proposed rank-two formulation passes these tests at the level of definitions and exact identities. The analytic two-channel dichotomy is known, and the determinant lattice can be written explicitly. What remains is the global stopped comparison that fuses them. Proving it would not just improve a numerical exponent; it would supply the mechanism by which the Hodge surplus survives the changing affine frame.

18. Mathematical tools likely to be relevant

The remaining problem has been reduced far enough that the relevant areas of mathematics can be described by the jobs they would perform.

Teichmuller and Hodge theory provide the archimedean source of positivity. The period maps compare the distinguished and conjugate representations, while the degeneration at the cusp determines the exact singular metrics. Further progress on this side would most likely take the form of a relative or stopped comparison for a moving pair of lines, rather than a new calculation of the Hodge exponents.

Arakelov geometry and slopes methods provide the language for fusing real and finite data. The desired bundle is only rank two, but its metrics depend on consecutive points of a thin orbit. A useful theorem would be coefficient-sensitive: it would exploit the fact that the two sections are successive Rosen states, not arbitrary vectors in the same lattice. A Bott–Chern or height-jump formula may be relevant because the crucial real quantity is a difference between two metric splittings of one extension.

Diophantine approximation provides models for turning small incidence into algebraic dependence. Moving-target versions of the Subspace Theorem, arithmetic holonomy bounds, and auxiliary-function methods all resemble parts of the needed comparison. Their standard forms lose either the critical height margin or the word-dependent coefficient data. A successful adaptation would have to build the Rosen frame into the auxiliary space from the outset.

Thin-group dynamics provide quantitative control of the allowed words. Strong approximation, expansion, symbolic transfer operators, and local–global dispersion can constrain how a rational flag moves through finite and real quotients simultaneously. Their likely role is to prove an inverse statement for the exceptional set: a trajectory that repeatedly evades the positive Hodge channel must carry a genuine stabilizer.

Homogeneous dynamics offers a broader conceptual setting. The three real embeddings place the diagonal group in a product of three rank-one groups, while units form a rank-two torus. Rigidity

of torus or Borel orbits may help classify the equality corridor. The required statement is not plain equidistribution, since the orbit is thin and the stopping rule is arithmetic, but measure rigidity suggests what persistent failure of transversality should look like.

Unit equations and transcendence methods provide the closing step and possibly an effective version. Once the finite support is controlled, the four-term split relation gives finiteness. Effective bounds for unit equations or linear forms in logarithms could in principle bound the period or reduce the remaining equality analysis to a finite computation. They do not supply the missing height bound on their own.

Exact symbolic computation remains relevant wherever the infinite argument reduces to a finite cone, automaton, or polynomial-positivity certificate. Such certificates are especially valuable for the cusp normalization and word combinatorics. They should be designed so that a small checker can reproduce the result independently and so that the mathematical reduction clearly states why the finite check covers every word.

The likely proof will not consist of a single black-box theorem from one of these subjects. The problem sits at an interface: Hodge theory knows the correct real contraction, Arakelov geometry knows how to count one divisor, thin dynamics knows which coefficients occur, and Diophantine finiteness knows how boundedness becomes repetition. The missing result must preserve enough structure for all four statements to apply to the same state.

One plausible architecture is to regard consecutive Rosen quotients as a point on an arithmetic correspondence over the triangle curve. The two candidate splittings form a rank-two extension with a canonical determinant. Equip it with the cusp-adapted Hodge metrics and its natural integral lattice. A relative slopes inequality along the correspondence would compare the maximal real direction with the determinant degree. The equality case would be a horizontal algebraic leaf, hence a rational stabilizer. This is a program rather than a theorem, but every term in it now corresponds to an explicit object in Part I.

Another plausible architecture is dynamical. Accelerate the Rosen map past deep cusp runs, and attach to each return the two Hodge cocycles and the determinant cocycle. Prove an inverse theorem saying that failure of the stopped maximal-slope inequality forces concentration on a periodic Borel orbit. The strict Hodge gap handles the generic returns; arithmetic integrality rules out diffuse exceptional recurrence. This route and the Arakelov route may be two descriptions of the same mechanism.

The decisive feature in either architecture is equality. A theorem that merely gives a positive average for generic trajectories would leave open the individual algebraic points in the exceptional set. The desired theorem must say that the only persistent zero-slope trajectories are the cusp and special hyperbolic trajectories already present in the conjecture.

19. Present mathematical status

Several parts of the program are established.

- Termination of the exact Rosen process gives a cusp, and repetition of a complete nonterminal state gives a K-rational hyperbolic fixed point.

- The geometric family, its real-multiplication splitting, and the Hodge growth rates $1, 3/5, 1/5$ provide a strict global contraction margin.
- The arithmetic fixed-point and split-character identities are exact. Bounded split norms imply the full termination-or-periodicity conclusion through a projective S-unit theorem.
- The discrepancy between the flat Hodge line and the moving remainder line is an explicit incidence term, not an unspecified analytic error.
- Two real channels can be constructed so that at least one retains a fixed portion of the Hodge gain, and their common determinant lattice is explicit.

The missing theorem is a global stopped fusion statement for those two channels. It must compare their real Hodge metrics with their one common finite divisor at consecutive record states. Equality must lead to a vanishing incidence or recurrence of the complete affine frame. Either event would close the orbit.

This is narrower than the original classification but not routine. It is a coefficient-sensitive arithmetic statement along a thin orbit. Standard height inequalities tend to forget the consecutive Rosen relation; standard thin-group estimates tend to give distribution rather than exact incidence; standard Hodge contraction applies to a fixed transported line rather than the moving record line. The desired result lies exactly at their intersection.

The conjecture is supported by extensive exact computation and by the fit between its two alternatives and the dynamics. Those computations also show why a proof must allow large intermediate excursions. A periodic or terminating orbit can pass through states of great archimedean height before closing. The required invariant should charge a complete excursion or a new arithmetic record, not impose a uniform ceiling at every step.

The pages that follow give the full account. Part I develops the organized argument, including the proved reductions and the active formulation of the missing estimate. Part II preserves the broader investigation: alternative routes, exact audits, finite certificates, counterexamples, and the sequence of corrections by which the present formulation was reached.

Contents

Foreword 2

Orientation and Background 3
A geometric classification problem . . . 3
The conjectural dichotomy . . . 3
Historical context . . . 4
Why one should expect a continued-fraction theorem . . . 5
The classical analogy and its limits . . . 6
What closing means exactly . . . 6
The arithmetic landscape . . . 7
Heights as global memory . . . 8
Thinness and the local–global tension . . . 9
Translation surfaces and the source of extra geometry . . . 10
The geometry of a cusp excursion . . . 11
McMullen's height argument as the model . . . 12
Why the most natural shortcuts stop short . . . 12
From scalar descent to a rank-two state . . . 14
Record times and the equality corridor . . . 14
The role of Diophantine finiteness . . . 15
What computation contributes . . . 16
The expected synthesis . . . 16
Mathematical tools likely to be relevant . . . 17
Present mathematical status . . . 18

I The Organized Account 24

1 The question 24
1.1 The present status and what the $q = 11$ comparison decides . . . 25
1.2 Why $q = 11$ is the decisive control case . . . 26
1.3 Vocabulary and conventions . . . 28

2 What the triangle group contributes 29

3 Heights over the cubic field 30
3.1 Why units are both harmless and dangerous . . . 31
3.2 Metrized lines in concrete terms . . . 31

4 McMullen's height argument 32
4.1 The geometric family . . . 32
4.2 Why the quadratic argument works . . . 35

5 Why degree three is different 35

6 The $p = 11$ comparison and the rank predicted by Hodge theory 37
6.1 The exact quintic field . . . 37
6.2 The five Hodge weights . . . 48

6.3 The threshold is an average, not a uniform contraction . . . 49
6.4 What degree five reveals: the natural rank is the unit rank . . . 51
6.5 Why this does not contradict Galois averaging . . . 53
6.6 Exact Rosen experiments . . . 54
6.7 The $q = 9$ control . . . 56

7 The same question as a thin-group stabilizer problem 58
7.1 The proof spine . . . 60
7.2 Which structure of the thin group is relevant . . . 62
7.3 There is no larger discrete group to interpolate through . . . 63

8 A group-wide integral Asai obstruction 82
8.1 The equality corridor as a Borel-frame problem . . . 84

9 The Rosen reduction 109
9.1 The exact algorithm . . . 109
9.2 Why bounded norms force closing . . . 117

10 The stopped scalar 118

11 Replacement ideas and what they actually establish 119
11.1 The ordinary three-factor height . . . 119
11.2 The two pairwise Hodge products . . . 119
11.3 The negative Hodge line . . . 120
11.4 Real word contraction and finite-state fusion . . . 120
11.5 Roth-type approximation . . . 120
11.6 Vojta and global height conjectures . . . 121
11.7 Finite-divisor bookkeeping (not a proof route) . . . 121
11.8 Operator norms versus evaluated tensors . . . 121

12 The pairwise Arakelov energy 122

13 The Hodge tensor behind the distinguished metric 125
13.1 What "global" means here . . . 127

14 The flat-line/record-line obstruction 127

15 What the pairwise tensor fixes, and what it does not 129

16 How the Hodge metric crosses a cusp 130

17 The remaining degree inequality 130
17.1 The rank-two pairwise bundle . . . 133
17.2 Why the affine framing cannot be discarded . . . 167
17.3 The relative integral-vector target . . . 170

18 The fixed-point form and the nearest-wall method 191
18.1 A thin-group stammering theorem . . . 195

19 Numerical audit 202

20 What has and has not been achieved 203

A Archived finite-quotient experiments 208
A.1 Discarded finite-quotient route: a presentation-free Fourier problem 208
A.2 Earlier congruence formulations and their limitation 217
A.3 Reduction by the entire parabolic orbit 222
A.4 The joint compression functor of the full cusp orbit 223
A.5 A whole-orbit pigeonhole which does not give rigidity 227
A.6 Optimal congruence lifting of a cusp 229
A.7 The relevant thin-orbit method: a pointed local–global theorem 235
A.8 Why a black-box spectral gap still does not use enough of the group 242
A.9 Comparison: a congruence shrinking-target criterion 248
A.10 What a full-group transfer theorem would have to say 250

II Research Archive and Proof Audits 255

II.1 Precise target and notation 255

II.2 Proved algebraic statements and exact computational certificates 255

II.3 Route I: ordinary moving-target Roth 282

II.4 Route II: Vojta's quantitative theorem with positive target height 283

II.5 Route III: optimal and conjectural symmetric Roth inequalities 284

II.6 The critical-scale obstruction and the Legendre near miss 285

II.7 Exact group-theoretic reformulation 286
II.7.1 What can be obtained from the whole thin group 287

II.8 Discarded route: the already-known 2-adic image 301
II.8.1 A precise 2-adic height criterion for the global theorem 304

II.9 Route V: a cubic Panti block criterion 305

II.10 What a standard number-theory conjecture actually gives 335

II.11 Exact ordinary target height on $A_2A_1^N$ 336

II.12 Constructing the alleged rational approximant 337

II.13 Route IV: twisting by the Hodge transfer section 339

II.14 Comparison with `curt11.tex`: nearest-wall Hodge charge and the S-unit boundary 340

II.15 Experiments and proof attempts still in progress 375

II.16 The global incidence reformulation 377

II.17 Reproducibility index **394**

II.18 Pairwise Arakelov energy: the August 5 advance **402**

II.19 Theoretical correction: fixed flat line versus moving record line **404**
II.19.1 The exact local loss . . . 404
II.19.2 The exact global identity and why Roth points the wrong way . . . 405
II.19.3 Exact energy cocycle . . . 405
II.19.4 What a correct Hodge–Arakelov route must construct . . . 406
II.19.5 A relative-cohomology construction of the moving cusp . . . 406
II.19.6 The exact negative line and the common source of the failed arguments . . . 408
II.19.7 The formal line and its higher-rank integral realization . . . 409
II.19.8 Why the expected alternatives are negative and zero Hodge weights . . . 410
II.19.9 Audit of the arithmetic-holonomy route . . . 411
II.19.10 A relative-dominated representation for the outer gap . . . 412
II.19.11 What the computations did and did not establish . . . 412

II.20 The rank-two correction: what the thin group actually supplies **413**
II.20.1 Why the determinant cannot control the units . . . 413
II.20.2 Why ordinary principal-bundle stability is circular here . . . 414
II.20.3 The order-seven monodromy splits the defect . . . 414
II.20.4 The full Asai orbit retains the same two coordinates . . . 422
II.20.5 A record-free route through the integral Asai exponent . . . 423
II.20.6 Audit of the integral-vector theorem for log-Anosov systems . . . 424
II.20.7 Full-group parabolic coding versus stammering theorems . . . 426
II.20.8 Advance: the shrinking real target is unnecessary . . . 428
II.20.9 A two-prime amplification of the fixed-target problem . . . 430
II.20.10 Why even a Ramanujan bound for the thin quotients is not enough . . . 433
II.20.11 A precise horizontal Hodge local limit statement . . . 433
II.20.12 A second full-orbit target: rigidity of a large incidence fibre . . . 466
II.20.13 Why Hilbert-modular torus rigidity does not directly close the argument . . . 476
II.20.14 The earlier, less efficient, congruence shrinking-target route . . . 477
II.20.15 Exact transfer-operator target and its entropy obstruction . . . 478
II.20.16 Arithmetic holonomy bounds: an exact applicability audit . . . 480
II.20.17 Algebraic leaves: why Bost's theorem proves the wrong kind of stabilizer . . . 480
II.20.18 The norm-Euclidean comparison: a successful reformulation . . . 510
II.20.18.1 The Eisenstein correction lattice . . . 512

II.21 Exact audits of the rank-two Arakelov and holonomy routes **515**

II.22 The $q = 11$ comparison: scalar sign versus unit-rank fusion **519**
II.22.1 Reassessment: the natural rank is $[K : \mathbb{Q}] - 1$ . . . 524
II.22.2 The codifferent H_4 reduction: a sharper scalar target . . . 526
II.22.3 Fubini–Study/Bott–Chern audit of the one-divisor pencil . . . 531
II.22.4 Two-time period determinant and the exact holonomy audit . . . 531
II.22.5 Relative-Anosov/Morse audit after the $p = 11$ test . . . 551

II.23 Primary references used for the Roth comparison **552**

Part I
The Organized Account

1 The question

Put

$$\lambda = 2\cos(\pi/7), \qquad K = \mathbf{Q}(\lambda),$$

and let

$$G_7 = \left\langle S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \quad T = \begin{pmatrix} 1 & \lambda \\ 0 & 1 \end{pmatrix} \right\rangle \subset \mathrm{PSL}_2(\mathbf{R}). \tag{1}$$

This is the orientation-preserving triangle group of signature $(2, 7, \infty)$. It acts on the upper half-plane by fractional linear transformations. The quotient has one cone point of order 2, one cone point of order 7, and one end of finite area. That end is a cusp. On the boundary line, a *cusp* means a point fixed by a parabolic element of G_7; equivalently it belongs to the orbit $G_7 \cdot \infty$.

The problem is to describe all points of $\mathbf{P}^1(K)$. There is one obvious second kind besides cusps. If $g \in G_7$ is hyperbolic and one of its eigenlines is defined over K, that eigenline gives a point of $\mathbf{P}^1(K)$ which need not be a cusp.

Conjecture 1.1 (Heptagonal cusp classification)**.** *Every point of $\mathbf{P}^1(K)$ is either a cusp of G_7, or belongs to the G_7-orbit of an eigenline of a hyperbolic element whose eigenvalues lie in K.*

The second alternative is not an error term. It is the periodic alternative in the Rosen continued-fraction algorithm. The theorem sought in this paper is therefore a termination-or-periodicity theorem, not a statement that every field point is a cusp.

There are two different uses of the word "periodic" in the surrounding literature. A cusp direction gives a periodic cylinder decomposition of the translation surface. A special hyperbolic direction is generally not periodic in that sense; it is the stable or unstable direction of a pseudo-Anosov map. It is periodic only as a geodesic/Rosen state. In this paper, "periodic alternative" always has the latter meaning [7].

The assertion is surprising because G_7 is much smaller than the full arithmetic group $\mathrm{PSL}_2(\mathcal{O}_K)$. There is no elementary congruence reason for its orbit of infinity to account for nearly every point of $\mathbf{P}^1(K)$. The missing points are predicted to have a rigid dynamical explanation: they are precisely the points on special periodic orbits.

For reference, λ has minimal polynomial

$$f(X) = X^3 - X^2 - 2X + 1. \tag{2}$$

The three real values of λ are

$$\lambda_0 = 1.8019377358\ldots, \qquad \lambda_1 = 0.4450418679\ldots,$$
$$\lambda_2 = -1.2469796037\ldots.$$

We write $\sigma_i : K \to \mathbf{R}$ for the embedding carrying λ to λ_i. The ring of integers is $\mathcal{O}_K = \mathbf{Z}[\lambda]$, the field discriminant is 49, and $\mathrm{N}(\lambda) = -1$. In particular, λ is a unit. These facts are elementary calculations from (2); the genuinely geometric input begins only with the variation of Hodge structure described below.

The two alternatives in Conjecture 1.1 are disjoint. Indeed, if a hyperbolic element and a nontrivial parabolic element fixed the same boundary point, conjugate that point to infinity and

the other hyperbolic fixed point to zero. The two transformations then have the forms $z \mapsto az$, $a > 1$, and $z \mapsto z + b$, $b \neq 0$. Their conjugates $z \mapsto z + a^{-n}b$ approach the identity, contradicting the discreteness of the group.

There really are points of the second kind. Put $\alpha = \lambda^2 - 1$. The identity

$$\alpha^2 - \lambda\alpha - 1 = 0$$

shows that α is fixed by

$$\begin{pmatrix} 1+\lambda^2 & \lambda \\ \lambda & 1 \end{pmatrix} \in G_7.$$

This matrix has determinant one and trace greater than two at σ_0, so it is hyperbolic. Any proposed height must therefore allow periodic orbits; strict descent at every nonterminal step would prove a false statement.

The example α is not expected to represent the only special orbit. Recent computations of Winsor find at least twenty-four G_7-orbits in $\mathbf{P}^1(K)$, including the cusp orbit, and find twenty-seven integers below 10^6 which are not cusps but are fixed by special hyperbolic elements [64]. Every tested heptagonal Rosen expansion in that search was either finite or eventually periodic. This is evidence for Conjecture 1.1, but it also corrects an important possible misunderstanding: the conjecture predicts a dichotomy between termination and periodicity, not a list consisting of the cusp orbit and one exceptional periodic orbit. Earlier two-orbit predictions [62] do not describe the observed special spectrum.

1.1 The present status and what the $q = 11$ comparison decides

The situation can be stated before introducing the later machinery.

1. The group-theoretic reformulation is proved: a point of $\mathbf{P}^1(K)$ has the required form exactly when its K-rational Borel stabilizer meets G_7 nontrivially.

2. The Hodge calculation is proved: the three normalized weights are $1, 3/5, 1/5$. Hence the integral Asai endpoint tensor has the strict formal weight

$$-1 + \frac{3}{5} + \frac{1}{5} = -\frac{1}{5}.$$

   This is a global advantage of the heptagonal problem, not a numerical feature of selected generators.

3. A whole-group arithmetic reduction is also proved. Trace dominance along every conjugate parabolic ray bounds every outer coefficient of the fixed-point form Q_M by its central coefficient height. The product formula then shows that little-oh of $H_0(Q_M)^{-2}$ forces exact closing at $q = 7$; in degree five the corresponding exponent is 4.

4. The adjacent-parabolic quadratic splits only at $p = 7$. Its four linear factors are denoted A, B, C, D below. Theorem 6.6 proves, using the ordinary projective S-unit theorem and the exact consecutive Rosen recurrence, that the classification would follow from a bounded split norm along each complete orbit. The previously proposed uniform estimate

$$\log \max\{\min(|\,\mathrm{N}\,A_j|, |\,\mathrm{N}\,B_j|), \min(|\,\mathrm{N}\,C_j|, |\,\mathrm{N}\,D_j|)\} \leq 6h(x) + B_7$$

   has new counterevidence. For the Pisot-unit family $x_N = (\lambda(\lambda^2 - 1))^{-N}$, its excess grows from 1.28 at $N = 10$ to 18.21 at $N = 60$ in the exact audit. No proof that this excess is

unbounded has been obtained, so the estimate is withdrawn as a route rather than declared disproved.

5. What is not proved is a replacement which distinguishes those large split norms on terminating or periodic unit orbits from a nonclosing excursion. The Hodge calculation supplies the numerical margin $1/5$, but no argument in this manuscript transports that margin through the changing remainder without losing it at the junctions.

1.2 Why $q = 11$ is the decisive control case

The degree-five case is not merely a more cumbersome version of the degree-three case. It separates two genuinely different conjectures which happen to give the same prediction at $q = 7$.

The first is a *scalar-defect conjecture*: tensor all Galois components together, use the resulting integral line only once, and let the sign of its Hodge degree decide whether an aperiodic trace-field ray is possible. The second is a stronger *simultaneous-slope conjecture*: compare the central component separately with every outer component, but prove that a single finite denominator can support all of those comparisons at once. The following table is the entire logical point of the comparison.

	$q = 7$	$q = 11$
full integral tensor	$-1 + \frac{3}{5} + \frac{1}{5} = -\frac{1}{5}$	$-1 + \frac{7}{9} + \frac{5}{9} + \frac{3}{9} + \frac{1}{9} = \frac{7}{9}$
central–outer pairs	$-\frac{2}{5}, -\frac{4}{5}$	$-\frac{2}{9}, -\frac{4}{9}, -\frac{6}{9}, -\frac{8}{9}$
adjacent quadratic wall	split over K_7	nonsplit over K_{11}

Thus the scalar integral object contracts at seven and expands at eleven, whereas every pairwise object contracts in both cases. A proof which uses one global integral height should therefore stop working at eleven. A proof which still works there must contain the new theorem that the four negative pairwise slopes can be fused without charging the same denominator four times. Merely listing four negative degrees does not provide that theorem.

This is also a Galois-descent obstruction, not just an unfortunate sum of numbers. Write the five conjugate endpoint lines at eleven as $\ell_0, \dots, \ell_4$. A favorable pair $\ell_0 \otimes \ell_m$ is not defined over $\mathbf{Q}$: Galois conjugation cycles it through the five pairs $\ell_j \otimes \ell_{j+m}$. By contrast, $\ell_0 \otimes \cdots \otimes \ell_4$ is Galois invariant and carries the canonical integral Asai lattice. Consequently the scalar exponent $7/9$ is the cost of descending one honest line to $\mathbf{Q}$. The simultaneous-slope proposal asks Harder–Narasimhan theory to choose among several conjugate directions after the Hodge metric has distinguished the uniformizing place. That is a nonlinear selection theorem, not an alternative tensor identity.

For odd q, let χ_m be the normalized Hodge weights of the Galois-conjugate rank-two systems and put

$$\Delta_q = 1 - \sum_{m \neq 1} \chi_m.$$

Section 6 proves that

$$\Delta_q > 0 \quad \Longleftrightarrow \quad q \in \{3, 5, 7, 9\}.$$

Thus $q = 11$ is the first odd parameter at which the full integral endpoint tensor changes from

contraction to expansion. Around the threshold one has

q	$[K_q : \mathbf{Q}]$	Hodge weights	Δ_q
5	2	$1, 1/3$	$2/3$
7	3	$1, 3/5, 1/5$	$1/5$
9	3	$1, 5/7, 1/7$	$1/7$
11	5	$1, 7/9, 5/9, 3/9, 1/9$	$-7/9$

The scalar interpretation gives the following family statement.

Conjecture 1.2 (Working odd-q phase-transition conjecture)**.** *If $\Delta_q > 0$, every point of $\mathbf{P}^1(K_q)$ has a terminating or eventually periodic Rosen expansion. If $\Delta_q < 0$, at least one point of $\mathbf{P}^1(K_q)$ has an infinite aperiodic Rosen expansion.*

The exact and computed cases line up with this statement. For $q = 5$, Rosen–Leutbecher proved the stronger assertion that every field point is a cusp. The cases $q = 7$ and $q = 9$ are the two nonarithmetic degree-three cases on the contracting side; the classification studied here is the $q = 7$ instance, and every tested $q = 9$ integer closes. At $q = 11$, only 1 among the integers $1, \dots, 100$ terminates within 500 steps, and none of $2, \dots, 100$ repeats there. The exact orbit of 2 has no detected return through 3000 steps. Its four Galois shape slopes are

$$1.787, \quad 1.565, \quad 1.327, \quad 1.106,$$

close to the Hodge predictions $(16/9, 14/9, 12/9, 10/9)$. Moreover, all 2989 blocks of length 12 in its first 3000 digits are distinct. These are finite diagnostics, not an infinite-orbit proof, but they make 2 the natural first candidate for the second half of Conjecture 1.2. Winsor's independent height experiment for the same pair $(q, x) = (11, 2)$ shows the same escaping behavior [64].

There is also exact earlier evidence at precisely this threshold. Hanson–Merberg–Towse–Yudovina proved by reduction modulo 2 that λ_{11}^2 is not a cusp; equivalently, its Rosen expansion does not terminate. They then conjectured that, for every $q > 9$, some point of $\mathbf{Q}(\lambda_q)$ has a nonperiodic Rosen expansion [62, Example 3 and Remarks 7, 9]. These are two different assertions. The congruence calculation does not exclude the possibility that λ_{11}^2 is fixed by a hyperbolic element after a long preperiod. Nevertheless, the fact that their conjectural transition begins at 11, exactly where Δ_q changes sign, is independent support for the scalar-defect conjecture rather than for uniform pairwise fusion. Our exact calculation detects no projective repetition for either λ_{11}^2 or λ_{11}^4 through 500 steps.

One terminology point is essential here. Arnoux–Schmidt prove that regular double q-gons of trace-field degree greater than two have trace-field directions which are not periodic cylinder directions [7]. A special hyperbolic eigenline already has that property, while its Rosen expansion is eventually periodic. Their theorem therefore proves the existence of noncuspidal directions, not the aperiodic alternative in Conjecture 1.2. It is consistent both with the heptagonal classification and with the proposed change at eleven.

Finally, two adjacent parabolics give the quadratic

$$F_q(X, Y) = X^2 + \lambda_q XY - Y^2.$$

It splits over K_7, but not over K_{11}; indeed its discriminant has norm 5741 at 11. This splitting is useful extra arithmetic for proving the heptagonal theorem. It is not the family-level criterion, because the same quadratic is nonsplit at $q = 5$, where every field point is nevertheless a cusp. The conclusion used below is therefore:

Verdict. The most economical uniform conjecture is the Hodge-defect phase transition: closure at $q = 5, 7, 9$, followed by an aperiodic field orbit at $q = 11$. The competing conjecture is not another scalar estimate; it is the stronger one-divisor, simultaneous-slope theorem, and it would predict closure also at eleven. The split torus is additional arithmetic available only for the $q = 7$ proof. Accordingly, the heptagonal argument should use the $1/5$ scalar margin unless it really constructs the simultaneous maximal-slope object. The $q = 11$ case is the first exact falsification test for that distinction.

1.3 Vocabulary and conventions

A *Rosen fraction* is the continued fraction adapted to the Hecke triangle group (1). Its digits record successive crossings of a geodesic through translates of the standard triangle. A *remainder* is one of the algebraic integers appearing in the corresponding second-order recurrence.

A *real embedding* of K is a field map $K \to \mathbf{R}$. There are three. We label them $\sigma_0, \sigma_1, \sigma_2$, with $\sigma_0(\lambda) = 2\cos(\pi/7)$. The first is distinguished because it is the embedding in which G_7 acts discretely on $\mathbb{H}$.

An *Arakelov degree* is a single global size obtained by adding the finite valuations of an algebraic line to its sizes at all real places. The product formula makes that degree independent of a change of algebraic basis. This is the feature needed here: the norm of a remainder must be handled as one divisor, not as a list of its prime factors.

A *Hodge line* is the one-dimensional space of holomorphic forms in one real-multiplication eigenspace of a varying Jacobian. Its metric is the usual L^2-norm of a holomorphic form. The variation of these metrics along the heptagonal Teichmüller curve is the source of contraction.

A *Fuchsian group* is a discrete subgroup of $\mathrm{PSL}_2(\mathbf{R})$. It is *cofinite* if a fundamental domain has finite hyperbolic area. A hyperbolic element has two boundary fixed points; a parabolic element has one. These elementary distinctions are responsible for the periodic and terminating alternatives above.

A subgroup of an arithmetic matrix group is called *thin* when it has infinite index but is Zariski dense: it satisfies no extra polynomial matrix identities, even though it occupies only a small part of the ambient integral group. The diagonal Galois representation of G_7 has this property in $\mathrm{Res}_{K/\mathbf{Q}}\,\mathrm{SL}_2$. Zariski density and strong approximation make its finite quotients large, but neither fact alone forces the group to meet a prescribed rational stabilizer.

A *variation of Hodge structure* is a family of cohomology spaces of algebraic curves, together with holomorphic subspaces varying over the base. The double-heptagon family has genus three. Its real multiplication splits the six-dimensional real cohomology into three two-dimensional pieces, one for each σ_i, and produces one Hodge line in each piece.

The *Asai vector* attached to $P \in K^2$ is simply the tensor

$$\sigma_0(P) \otimes \sigma_1(P) \otimes \sigma_2(P).$$

The name refers to tensor induction, but no general Asai formalism is needed here. The point of the tensor is that it remembers all three real embeddings while still belonging to one rational integral lattice. The product formula can therefore compare its Hodge norm with arithmetic integrality.

A *norm record* is a remainder whose absolute field norm is larger than the norms of all preceding remainders. A *stopped inequality* is an inequality asserted only immediately before the transition that creates such a record. This qualification is essential: unit holonomy can make archimedean quantities very large on a periodic orbit, but a periodic orbit creates no indefinitely growing norm records.

Finally, the *maximal slope* of an adelic vector bundle is the largest Arakelov degree, divided by

rank, among its subbundles. In the argument below it is simply a coordinate-free way to select the largest norm seen so far. No knowledge of the general Harder–Narasimhan formalism is needed to read the reduction.

2 What the triangle group contributes

The phrase "triangle group" can obscure which facts do mathematical work. At the distinguished real embedding one has

$$S^2 = 1, \qquad (ST)^7 = 1,$$

and T is parabolic. A fundamental triangle has angles $\pi/2, \pi/7, 0$. The zero angle places its third vertex on the boundary of $\mathbb{H}$, giving the cusp. The resulting tessellation supplies a finite symbolic coding away from arbitrarily long excursions around the cusp.

All matrices have entries in $\mathcal{O}_K$, so applying the three embeddings gives representations

$$\sigma_i : G_7 \longrightarrow \mathrm{SL}_2(\mathbf{R}), \qquad i = 0, 1, 2.$$

Only $\sigma_0(G_7)$ is the discrete tessellation just described. The other two representations should not be pictured as two additional hyperbolic triangle tessellations. They occur as monodromy on the two conjugate Hodge summands and are related to the first action through holomorphic period maps.

For some combinatorial calculations it is useful to set $R = ST$ and $A_j = -SR^j$, $1 \le j \le 6$. If $\alpha = \lambda^2 - 1$, then

$$A_1 = \begin{pmatrix} 1 & \lambda \\ 0 & 1 \end{pmatrix}, \qquad A_2 = \begin{pmatrix} \lambda & \alpha \\ 1 & \lambda \end{pmatrix}, \qquad A_3 = \begin{pmatrix} \alpha & \alpha \\ \lambda & \alpha \end{pmatrix},$$
$$A_4 = \begin{pmatrix} \alpha & \lambda \\ \alpha & \alpha \end{pmatrix}, \qquad A_5 = \begin{pmatrix} \lambda & 1 \\ \alpha & \lambda \end{pmatrix}, \qquad A_6 = \begin{pmatrix} 1 & 0 \\ \lambda & 1 \end{pmatrix}.$$

The entries are nonnegative at σ_0, but not necessarily at the other embeddings. The end letters A_1, A_6 are the two orientations of motion around the cusp. The four middle letters force a return through a compact part of the quotient. This cusp/compact decomposition is the real dynamical structure used later.

It is also useful to say what is not used. Residual images at a fixed prime do not distinguish a cusp from a special periodic point and play no role here. The abstract presentation of G_7 is not enough either: the argument uses its integral matrices, all three real embeddings, the finite Markov coding, and the Hodge modular embedding. Real spectral estimates for words are valuable, but they become a closing argument only after they are incorporated into a global arithmetic height.

> **Geometric background.** The discreteness and cofinite area of the Hecke triangle group, and the existence of the Teichmüller curve with this Veech group, are global geometric inputs. The displayed matrices and their elementary parabolic/compact interpretation are direct calculations.

The six matrices above are coordinates for the geodesic flow; they are not the source of the proposed rigidity. In fact powers of a parabolic letter have essentially the same logarithmic size in all three real embeddings, so no worst-letter estimate can produce the Hodge exponents needed below. The positive uses of the group are instead global. They are proved in later sections: the Galois-diagonal image is Zariski dense in SL_2^3; a fixed integral Asai vector forces every noncuspidal rational ray into an exceptional Lyapunov set; the triangle curve is the full stabilizer of its lift in

the Hilbert modular threefold; and its pointed congruence covers are connected ambient pullbacks. These statements concern all long elements and all cusp conjugates at once. The remaining estimate is a quantitative rigidity theorem for that whole thin orbit, not another inequality for the displayed matrices.

3 Heights over the cubic field

Before bringing in Hodge theory, it is useful to write the arithmetic height in coordinates. Let $[p : q] \in \mathbf{P}^1(K)$, and choose a primitive integral pair (p, q), meaning that the ideal generated by p, q is $\mathcal{O}_K$. Put

$$H_i(p, q) = \log \max\{|\sigma_i(p)|, |\sigma_i(q)|\}.$$

With the normalization used in this paper, the absolute projective height satisfies

$$3h([p : q]) = H_0(p, q) + H_1(p, q) + H_2(p, q). \tag{3}$$

The finite terms vanish for a primitive pair. If the pair is not primitive, the common ideal contributes the usual finite correction.

Formula (3) is independent of the chosen homogeneous coordinates. Indeed, multiplying both coordinates by $a \in K^\times$ adds $\log |\sigma_i(a)|$ at the three real places and adds the opposite valuation terms at the finite places. The product formula says that the total change is zero. This cancellation is the elementary core of an Arakelov height.

The height used in the recent computations is

$$H_\Sigma([p : q]) = \prod_{i=0}^{2} \big(|\sigma_i(p)| + |\sigma_i(q)|\big). \tag{4}$$

It is useful for plotting an orbit, but it does not introduce a new coercive quantity.

Lemma 3.1 (The two projective heights differ by a bounded factor)**.** *For a primitive pair (p, q), put*

$$H_{\max}([p : q]) = \prod_{i=0}^{2} \max\{|\sigma_i(p)|, |\sigma_i(q)|\}.$$

Then

$$H_{\max}([p : q]) \le H_\Sigma([p : q]) \le 8 H_{\max}([p : q]).$$

Both quantities are unchanged by multiplying (p, q) by a unit.

Proof. At each real embedding,

$$\max\{|a|, |b|\} \le |a| + |b| \le 2 \max\{|a|, |b|\}.$$

Multiply the three inequalities. Unit invariance follows from $\prod_i |\sigma_i(\varepsilon)| = 1$. □

Thus $\log H_\Sigma$ and $3h$ differ by at most $\log 8$. Replacing the max norm by the sum norm can improve the appearance of a finite plot, but it cannot repair an unbounded failure of descent. The new information in Winsor's experiment is the observed termination-or-periodicity, not a different asymptotic height.

3.1 Why units are both harmless and dangerous

If ε is a unit, the finite ideal of $(\varepsilon p, \varepsilon q)$ is unchanged and

$$\sum_i \log|\sigma_i(\varepsilon)| = 0.$$

The total height is unchanged, but its three real summands can move by an arbitrarily large vector in the plane $x_0 + x_1 + x_2 = 0$. Consequently a claim involving only one embedding, or unequal fixed weights at the three embeddings, must be checked for unit invariance.

Dirichlet's unit theorem also gives a useful normalization.

Lemma 3.2 (Balanced representative)**.** *There is a constant C_K such that every point of $\mathbf{P}^1(K)$ has a primitive representative (p, q) satisfying*

$$|H_i(p, q) - h([p : q])| \le C_K, \qquad i = 0, 1, 2.$$

Proof. The logarithmic images of units form a lattice of full rank in the plane $x_0 + x_1 + x_2 = 0$. Translate

$$(H_0 - h, H_1 - h, H_2 - h)$$

by this lattice into a fixed bounded fundamental parallelogram. The corresponding unit multiple of (p, q) has the required property. □

Only one unit is chosen here, at the beginning. Rebalancing after every Rosen step would change the recurrence matrices and destroy the deterministic orbit. A valid proof must allow the three archimedean components to become unbalanced later while retaining control by the initial height.

3.2 Metrized lines in concrete terms

For the purposes of this paper, a metrized K-line consists of a one-dimensional K-vector space, a lattice at every finite prime, and a norm at each of the three real embeddings. If e is a nonzero basis, its Arakelov degree is the sum of the finite lattice indices and the three logarithmic real norms. Replacing e by ae changes these summands, but their total is unchanged by the product formula.

For example, if $r \in \mathcal{O}_K \setminus \{0\}$, let $\overline{\mathcal{N}}(r)$ have finite unit ball $r^{-1}\mathcal{O}_{K,\mathfrak{p}}$ at $\mathfrak{p}$ and standard real norms. Then

$$\widehat{\deg}\,\overline{\mathcal{N}}(r) = \log|\,\mathrm{N}\, r|. \tag{5}$$

This one degree already sums all prime ideals dividing r, with their multiplicities. Factoring r may be useful for diagnosis, but it is not needed to define or compare the degree.

An Arakelov argument in this problem must therefore do two things at once: identify a finite lattice whose degree is the norm of the selected remainder, and identify real metrics on that same line whose variation is controlled by Hodge theory. A real contraction without the lattice is not a height; a finite norm without the real contraction is not bounded.

> **Background, not a new theorem.** The height formula, the balanced-representative lemma, and (5) are standard consequences of the product formula and Dirichlet's unit theorem. They are included because they explain why the final estimate must be global and why repeated renormalization by units is illegitimate.

4 McMullen's height argument

4.1 The geometric family

Take two regular heptagons and identify parallel sides by translation. The resulting translation surface has genus three. Applying matrices in $\mathrm{SL}_2(\mathbf{R})$ to its polygons produces a Teichmüller disk; after quotienting by the stabilizer, this disk becomes an algebraic curve. Its projective stabilizer is G_7. Thus the hyperbolic triangle-group dynamics and the variation of genus-three Jacobians are two descriptions of the same one-dimensional family.

Real multiplication by K splits the real first cohomology into three rank-two systems

$$H^1 = \mathcal{V}_0 \oplus \mathcal{V}_1 \oplus \mathcal{V}_2. \tag{6}$$

Here $\mathcal{V}_i$ is the summand on which $a \in K$ acts as the real number $\sigma_i(a)$. The line of holomorphic forms in $\mathcal{V}_i$ is denoted by $\mathcal{L}_i$.

There is a necessary sign normalization at the third embedding. Put

$$\varepsilon_i = \mathrm{sign}(\lambda_i), \qquad R_i = \begin{pmatrix} \varepsilon_i & 0 \\ 0 & 1 \end{pmatrix}, \qquad \rho_i(g) = R_i \sigma_i(g) R_i^{-1}.$$

Thus $R_0 = R_1 = 1$, while R_2 has determinant minus one, and

$$\rho_i(T) = \begin{pmatrix} 1 & |\lambda_i| \\ 0 & 1 \end{pmatrix}.$$

The conjugation by R_2 is the change from the raw Galois basis to a positively oriented symplectic basis on $\mathcal{V}_2$. It is essential: a holomorphic map $\mathbb{H} \to \mathbb{H}$ equivariant at a cusp can intertwine only parabolics with the same orientation. Conjugation by R_i preserves traces, translation lengths, Zariski closure, and all algebraic stabilizer statements.

Taking periods in these symplectic coordinates gives an equivariant map

$$I : \mathbb{H} \longrightarrow \mathbb{H}^3, \qquad I(z) = (z, F_1(z), F_2(z)). \tag{7}$$

The first coordinate uniformizes the triangle quotient. The other two coordinates are holomorphic and satisfy

$$F_i(\sigma_0(g)z) = \rho_i(g) F_i(z), \qquad g \in G_7.$$

The Schwarz–Pick lemma says that they do not expand the hyperbolic metric. Because neither is a hyperbolic automorphism, the inequality is strict at every interior point.

This pointwise statement already gives a theorem about every element of the thin group, without choosing a word or a generating set.

Proposition 4.1 (Global length domination)**.** *If $g \in G_7$ is hyperbolic in the distinguished representation and $\ell_i(g)$ denotes the translation length of $\sigma_i(g)$ on $\mathbb{H}$, then*

$$\ell_i(g) < \ell_0(g), \qquad i = 1, 2.$$

In particular

$$2\ell_0(g) - \ell_1(g) - \ell_2(g) > 0.$$

Proof. Choose z on the axis of $\sigma_0(g)$. Equivariance and strict Schwarz–Pick contraction give

$$\ell_i(g) \le d_{\mathbb{H}}\big(F_i(z), \rho_i(g)F_i(z)\big) = d_{\mathbb{H}}\big(F_i(z), F_i(gz)\big) < d_{\mathbb{H}}(z, gz) = \ell_0(g).$$

□

The same proof gives a quantitative relative gap without inspecting generators. Fix a standard cusp neighborhood in the triangle quotient and let $L_{\mathrm{thick}}(g)$ be the length of the part of the closed geodesic of g outside that neighborhood.

Proposition 4.2 (Thick-part spectral charge)**.** *There is $\kappa_7 > 0$, depending only on the chosen cusp neighborhood, such that every hyperbolic $g \in G_7$ satisfies*

$$2\ell_0(g) - \ell_1(g) - \ell_2(g) \geq \kappa_7 L_{\mathrm{thick}}(g).$$

Proof. The hyperbolic derivative norms of F_1, F_2 descend to continuous functions on the triangle quotient. On its compact thick part there are numbers $k_i < 1$ with $\|DF_i\|_{\mathrm{hyp}} \leq k_i$. Everywhere else Schwarz–Pick gives the weaker bound one. Integrate along one period of the axis of g. The length of the image path under F_i is at most

$$\ell_0(g) - (1 - k_i)L_{\mathrm{thick}}(g),$$

and the translation length $\ell_i(g)$ is no larger than the length of that path. Add the two estimates and take $\kappa_7 = 2 - k_1 - k_2$. □

The same argument is not restricted to periodic axes. Fix $z_* \in \mathbb{H}$, write $F_0(z) = z$, and set

$$d_i(g) = d_{\mathbb{H}}\big(F_i(z_*), F_i(gz_*)\big) = d_{\mathbb{H}}\big(F_i(z_*), \rho_i(g)F_i(z_*)\big).$$

Let $L_{\mathrm{thick}}(z_*, gz_*)$ be the length outside the chosen cusp neighborhood of the geodesic segment from z_* to gz_*, after projection to the triangle surface.

Proposition 4.3 (Open-segment relative domination)**.** *With the same $\kappa_7 > 0$, every $g \in G_7$ satisfies*

$$2d_0(g) - d_1(g) - d_2(g) \geq \kappa_7 L_{\mathrm{thick}}(z_*, gz_*). \tag{8}$$

If $a_i(g)$ denotes the top Cartan exponent in the i-th factor for any fixed Euclidean norms, then

$$2a_0(g) - a_1(g) - a_2(g) \geq \frac{\kappa_7}{2} L_{\mathrm{thick}}(z_*, gz_*) - O_7(1). \tag{9}$$

Thus the block representation

$$\rho_\star = \mathrm{Sym}^2 \rho_0 \oplus (\rho_1 \otimes \rho_2) \tag{10}$$

has a singular-value gap which grows with thick-part displacement and is neutral on pure parabolic motion.

Proof. Integrate the hyperbolic derivative of F_i along the geodesic segment from z_* to gz_*. On its thick part the derivative norm is at most $k_i < 1$, and on the remaining part it is at most one. Hence

$$d_i(g) \leq d_0(g) - (1 - k_i)L_{\mathrm{thick}}(z_*, gz_*).$$

Add the two inequalities. This gives (8). In $\mathrm{SL}_2(\mathbf{R})$, half the displacement of a fixed basepoint differs by a bounded amount from the top Cartan exponent computed in any other fixed norm. Finally, the top exponents of the two blocks in (10) are $2a_0$ and $a_1 + a_2$, respectively. □

> **What this adds.** The relative gap is now proved for every prefix of the whole group, not only for generators or closed words. It supplies the analytic relative-domination input suggested by cusped-Anosov theory. It still does not prove the arithmetic classification: the decreasing block $\mathrm{Hom}(\rho_1 \otimes \rho_2, \mathrm{Sym}^2 \rho_0)$ is not a rational one-dimensional subrepresentation after Galois descent. The record-framed Asai tensor is the integral replacement, and controlling its moving incidence angle is exactly the stopped-height problem below.

The proposition is the group-wide source of the two pairwise spectral gaps used later. Parabolics are the neutral boundary because their translation length is zero in every factor. Passing from this eigenvalue statement to a stopped estimate for an evaluated rational flag still requires control of endpoint angles; that is exactly the incidence problem isolated below. The finite-state word estimates sharpen the size of the gap on induced returns, but do not create it.

The Hodge degrees give a global summary of these contractions. The logarithmic cotangent bundle of the orbifold $(2, 7, \infty)$ has degree

$$\deg \Omega^1_{\log} = 1 - \frac{1}{2} - \frac{1}{7} = \frac{5}{14}.$$

The Lyapunov exponents are $1, 3/5, 1/5$. With the embeddings numbered as in (2), the value $\lambda_1 = 2\cos(3\pi/7)$ belongs to the exponent $1/5$, whereas the projective class of $\lambda_2 = -2\cos(2\pi/7)$ belongs to the exponent $3/5$. Thus

$$\deg \mathcal{L}_0 = \frac{5}{28}, \qquad \deg \mathcal{L}_1 = \frac{1}{28}, \qquad \deg \mathcal{L}_2 = \frac{3}{28}. \tag{11}$$

These fractions are not fitted parameters; they are invariants of the variation of Hodge structure [7, 40].

They can also be read directly from the full triangle-group representations. Write $R = ST$, so projectively

$$G_7 = \langle S, R \mid S^2 = R^7 = 1 \rangle.$$

Normalize relative Euler number using the positive parabolic at the cusp. In the three signed representations, the rotation numbers of S are $1/2$, while those of R are respectively $1/7, 3/7, 2/7$. The relative Euler numbers are therefore

$$e(\rho_0) = 1 - \frac{1}{2} - \frac{1}{7} = \frac{5}{14}, \quad e(\rho_1) = 1 - \frac{1}{2} - \frac{3}{7} = \frac{1}{14}, \quad e(\rho_2) = 1 - \frac{1}{2} - \frac{2}{7} = \frac{3}{14}. \tag{12}$$

For a polarized rank-two variation, the relative Euler number is twice the parabolic degree of its holomorphic line. Thus (12) recovers (11) without referring to a Rosen alphabet.

More importantly, the defect used later is the full-group class

$$2e(\rho_0) - e(\rho_1) - e(\rho_2) = \frac{3}{7}. \tag{13}$$

This is exactly $\deg \mathcal{N}^{-2}$, and the Higgs-ratio section realizes it as the effective orbifold divisor $3P_7$. The "charges" seen in a symbolic return are consequently not properties of selected generators. They are local records of one bounded-Euler-class defect of the entire triangle-group representation. What this class controls is the total closed-orbit or degree defect. It does not by itself evaluate a moving rational Borel, which is why it supplies the Hodge surplus but not the missing stopped incidence transport.

> **Deep geometric input.** The real-multiplication splitting, modular embedding, and Hodge degrees come from the geometry of the Teichmüller curve. They are not consequences of the Rosen recurrence. The later arithmetic argument uses them through strict Schwarz–Pick contraction and the behavior of the Hodge metric at the cusp.

4.2 Why the quadratic argument works

McMullen's theorem for a real quadratic trace field is the model [17]. In that setting the Jacobian has two rank-two real eigenspaces. Let C be a nonzero integral homology class, and let C_0, C_1 be its two projections. The product

$$\|C_0\|_{\mathrm{Hdg}}\,\|C_1\|_{\mathrm{Hdg}} \tag{14}$$

is unchanged when C is multiplied by a unit: the two real absolute values of a unit have product one. Including the integral lattice at the finite places makes (14) a height on a projective line over the trace field.

Now point a geodesic toward a field point x. Choose a covector (a, b) with $a + bx = 0$, and let $\gamma(s)$ be the unit-speed geodesic toward $\sigma_0(x)$. In a standard Hodge frame the relevant logarithmic norms are

$$\eta_0(s) = \log \frac{|a_0 + b_0\gamma(s)|^2}{\Im\gamma(s)}, \tag{15}$$

$$\eta_i(s) = \log \frac{|a_i + b_i F_i(\gamma(s))|^2}{\Im F_i(\gamma(s))}. \tag{16}$$

Here $a_i = \sigma_i(a)$ and $b_i = \sigma_i(b)$. A direct Busemann calculation in the uniformizing coordinate, followed by Schwarz–Pick in a conjugate coordinate, gives

$$\eta_0'(s) = -1, \qquad \eta_i'(s) \le \|DF_i(\gamma(s))\|_{\mathrm{hyp}} < 1. \tag{17}$$

The covector (a, b) is fixed throughout this calculation. More precisely, it is a flat integral homology class, viewed in changing Hodge metrics. The identity $a + bx = 0$ is what makes the first derivative exactly -1. A sequence of different cusp approximants does not satisfy the same differential equation merely because its pointwise norms have the same formal weights. With only one conjugate factor, the derivative of $\eta_0 + \eta_1$ is negative. On a fixed compact portion of the triangle quotient it is at most $-\kappa$ for some $\kappa > 0$. Thus each return through the compact part consumes a definite amount of a height which is bounded below.

The cusp requires care because the derivative norm of the conjugate period map approaches one there. McMullen compares the Hodge metric with the metric coming from the nilpotent orbit of the degenerating abelian variety. The parabolic contribution is then an endpoint term of bounded size. An orbit cannot make infinitely many compact returns, and hence must eventually remain in a cusp. This is the geometric origin of the continued-fraction termination statement.

The important lesson is not merely that one has found a decreasing real function. The function is the archimedean part of a projective arithmetic height; the product formula has already included every finite place. This is what makes the descent a classification argument rather than a spectral estimate for words.

> **Elementary calculation versus deep theorem.** Once the modular embedding is given, the derivatives in (17) are a short hyperbolic calculation. The existence of the modular embedding, strictness of its conjugate factors, and uniform comparison with a nilpotent-orbit metric are the substantive inputs in McMullen's work.

5 Why degree three is different

The genus-three heptagonal family has three real-multiplication factors. The period map has the form

$$t \longmapsto (t, F_1(t), F_2(t)) \in \mathbb{H}^3.$$

Equation (17) gives two genuine inequalities:

$$(\eta_0+\eta_i)' \le -1+\|DF_i\|_{\mathrm{hyp}} < 0, \qquad i=1,2. \tag{18}$$

Adding them gives

$$(2\eta_0+\eta_1+\eta_2)' \le -2+\|DF_1\|_{\mathrm{hyp}}+\|DF_2\|_{\mathrm{hyp}} < 0. \tag{19}$$

Thus there is enough analytic contraction. The obstruction is arithmetic: the expression on the left is not the ordinary three-factor projective height.

This obstruction can be seen without Hodge terminology. Suppose that a weighted product

$$\|C_0\|^{w_0}\|C_1\|^{w_1}\|C_2\|^{w_2}$$

is unchanged when C is multiplied by every unit of $\mathcal{O}_K$. Its logarithm changes by $\sum_i w_i \log|\sigma_i(\varepsilon)|$. Dirichlet's unit theorem says that the logarithms of units span a lattice in the plane $x_0+x_1+x_2=0$. Hence a weight annihilating every unit must satisfy

$$w_0=w_1=w_2. \tag{20}$$

The analytically contracting weight $(2,1,1)$ does not pass this test if it is applied as a bare product of three independent Hodge norms.

The smooth three-factor quantity has derivative bounded by

$$-1+\|DF_1\|_{\mathrm{hyp}}+\|DF_2\|_{\mathrm{hyp}},$$

which can be positive. In fact the two outer derivative norms tend to one at the cusp, so this bound tends to 1. The failure is one missing copy of the central contraction.

The Hodge degrees display the same fact. Let $\mathcal{L}_i$ be the three Hodge lines on the compactified Teichmüller curve. Their normalized degrees are

$$\deg\mathcal{L}_0=\frac{5}{28}, \qquad \deg\mathcal{L}_1=\frac{1}{28}, \qquad \deg\mathcal{L}_2=\frac{3}{28}.$$

Consequently

$$\mathcal{M}=\mathcal{L}_1\otimes\mathcal{L}_2\otimes\mathcal{L}_0^{-1}, \qquad \deg\mathcal{M}=-\frac{1}{28}. \tag{21}$$

The negative line $\mathcal{M}$ measures the missing central copy. It does not, by itself, give a projective arithmetic height; an integral finite-place model is also required.

It is worth being precise about the conclusion. There is no shortage of real contraction for a *fixed flat covector*: both pairwise sums in (18) decrease. A Rosen prefix, however, selects a new covector at every return. There are consequently two missing operations. One must first join the repeated central factor to the finite divisor of an algebraic remainder in a way that is homogeneous under $K^\times$. The pairwise Arakelov line in Section 12 does this. One must then transport that changing line through the Hodge variation. The second operation is the genuine unproved step inside the argument; a stopped estimate cannot be inferred merely by applying the fixed-covector calculation repeatedly.

> **Substantive diagnosis.** The deep fact is the existence of two strict Hodge contractions. The unit-weight test is elementary. Together they locate the problem exactly: the desired real weight exists, but it must be completed by a finite Arakelov term rather than treated as an ordinary weighted Hodge height.

6 The $p = 11$ comparison and the rank predicted by Hodge theory

This section compares the heptagonal problem with

$$G_{11} = \langle S, T_{11} \rangle, \qquad T_{11} = \begin{pmatrix} 1 & \lambda_{11} \\ 0 & 1 \end{pmatrix}, \qquad \lambda_{11} = 2\cos(\pi/11).$$

The comparison is useful only if the same arithmetic, dynamics, and Hodge normalizations are used on both sides. We therefore make each of them explicit.

6.1 The exact quintic field

Put $K_{11} = \mathbf{Q}(\lambda_{11})$. Since

$$\lambda_{11}^2 - 2 = 2\cos(2\pi/11),$$

this is the field $\mathbf{Q}(\zeta_{11})^+$ in the question. Direct elimination from $\zeta_{22} + \zeta_{22}^{-1}$ gives

$$\lambda_{11}^5 - \lambda_{11}^4 - 4\lambda_{11}^3 + 3\lambda_{11}^2 + 3\lambda_{11} - 1 = 0. \tag{22}$$

The five embeddings send λ_{11} to

$$2\cos(k\pi/11), \qquad k = 1, 3, 5, 7, 9.$$

An exact PARI calculation gives

$$\mathcal{O}_{K_{11}} = \mathbf{Z}[\lambda_{11}], \qquad \operatorname{disc}(K_{11}) = 11^4, \qquad h(K_{11}) = 1, \qquad M(K_{11}) = \frac{1}{11}.$$

Only the polynomial and the embeddings are needed for the comparison; the class-number statement merely permits the numerical height to be written using primitive pairs without ideal-class notation. The last identity is the exact Euclidean-minimum calculation of Cerri, as recorded by Lezowski–McGown [13]. It will be used below to separate the norm-Euclidean input from the Hodge sign.

Proposition 6.1 (The adjacent word is split only at seven in the prime family)**.** *Let p be an odd prime, put $K_p = \mathbf{Q}(\lambda_p)$, and*

$$F_p(X, Y) = X^2 + \lambda_p XY - Y^2, \qquad E_p = K_p[t]/(t^2 - \lambda_p t - 1).$$

Then

$$F_p(X, Y) = \operatorname{Nm}_{E_p/K_p}(X + tY). \tag{23}$$

If P_n denotes the Pell sequence

$$P_0 = 0, \qquad P_1 = 1, \qquad P_{n+2} = 2P_{n+1} + P_n,$$

then

$$\operatorname{Nm}_{K_p/\mathbf{Q}}(\lambda_p^2 + 4) = P_p. \tag{24}$$

Consequently E_p is split over K_p if and only if $p = 7$. More explicitly:

1. *In K_7, if*

$$\alpha = \lambda_7^2 - 1, \qquad s = 2\lambda_7^2 - \lambda_7 - 2,$$

then

$$\alpha^{-1} = \alpha - \lambda_7, \qquad s^2 = \lambda_7^2 + 4, \qquad \operatorname{Nm}_{K_7/\mathbf{Q}}(s) = -13,$$

and

$$F_7(X, Y) = (X + \alpha Y)(X - \alpha^{-1} Y). \tag{25}$$

Thus $E_7 \simeq K_7 \times K_7$.

2. *In K_{11},*

$$\mathrm{Nm}_{K_{11}/\mathbf{Q}}(\lambda_{11}^2+4)=5741. \tag{26}$$

The integer 5741 *is prime. Hence λ_{11}^2+4 is not a square in K_{11}, F_{11} is irreducible, and E_{11}/K_{11} is a quadratic field extension.*

The matrix

$$W_p=\begin{pmatrix}1+\lambda_p^2 & \lambda_p\\ \lambda_p & 1\end{pmatrix}$$

is the product of two adjacent primitive cusp parabolics, up to the orientation convention. Its eigenvalues are t^2, t^{-2}. In particular W_7 has K_7-rational eigenlines and is special, while W_p has no K_p-rational eigenline for every odd prime $p\neq 7$.

Proof. The two roots of $t^2-\lambda_p t-1$ are t and $\lambda_p-t=-t^{-1}$. Multiplying the two conjugates of $X+tY$ proves (23). The identities in the first part follow by reduction modulo $\lambda_7^3-\lambda_7^2-2\lambda_7+1$; they also give (25).

We next prove the prime-family norm formula. Let $\zeta=e^{2\pi i/p}$ and $v=1+\sqrt{2}$. The conjugates of λ_p^2+4 are

$$6+\zeta^j+\zeta^{-j},\qquad 1\le j\le (p-1)/2.$$

Writing $u=v^2=3+2\sqrt{2}$, taking the product over all $1\le j\le p-1$, and using $(Z+u)(Z+u^{-1})=Z^2+6Z+1$, we obtain

$$\mathrm{Nm}_{K_p/\mathbf{Q}}(\lambda_p^2+4)^2=\Phi_p(-u)\Phi_p(-u^{-1})=\left\{\frac{u^p+1}{u^{(p-1)/2}(u+1)}\right\}^2.$$

Every real conjugate of λ_p^2+4 is positive, so the positive square root is the norm. Since p is odd, the expression in braces is

$$\frac{v^p+v^{-p}}{2\sqrt{2}}=P_p,$$

which proves (24).

Ljunggren's theorem says that the only square Pell numbers are $P_0=0, P_1=1, P_7=169$ [1, 2]. If λ_p^2+4 were a square in K_p, its absolute norm would be an integer square. Hence this is possible for an odd prime only when $p=7$, where the explicit factorization above proves that it actually occurs. For $p=11$, (24) gives $P_{11}=5741$; trial division by the primes at most $\sqrt{5741}<76$ also verifies directly that it is prime.

Finally, $\mathrm{tr}(W_p)=2+\lambda_p^2$ and $\det(W_p)=1$. Since

$$t^2+t^{-2}=2+\lambda_p^2,$$

its characteristic polynomial has roots t^2, t^{-2}. Its eigenlines are defined over K_p exactly when λ_p^2+4 is a square. The last assertion follows. □

What an alternating block says at $p=11$. The comparison is sharper than a change from three to five real places. Over E_p, put

$$u(P)=\frac{tX+Y}{\iota(t)X+Y},\qquad P=(X,Y)^t,$$

where $\iota(t)=-t^{-1}$. The two rows are eigenvectors of W_p, so

$$u(W_pP)=t^4u(P),\qquad \iota(u)=u^{-1},\qquad \iota(t^4)=t^{-4}.$$

Thus an alternating digit block is a translation on the norm-one torus for every p, but the height of that translation behaves differently at seven and eleven.

At $p = 7$, the torus splits over the trace field. One can select a single K_7-valued character, and the signed logarithms of its three conjugates add to zero by the product formula. The Rosen interval fixes the sign of the distinguished logarithm. Proposition 6.3 below then shows that the remaining positive parts cancel it: an arbitrarily long alternating block does not raise projective height by more than a fixed constant.

At $p = 11$, E_{11} is a field. For each embedding σ of K_{11}, its two extensions to E_{11} send u to reciprocal real numbers. If $a_\sigma = \log|\widetilde{\sigma}(u)|$ and $b_\sigma = \log|\widetilde{\sigma}(t^4)|$, their combined contribution to height after m pairs is

$$\max(0, a_\sigma + mb_\sigma) + \max(0, -a_\sigma - mb_\sigma) = |a_\sigma + mb_\sigma|.$$

The Rosen interval controls the sign in the distinguished component, but it cannot cancel the four outer absolute values. They may grow linearly. This explains, without appealing to the degree alone, why the neutral-corridor reduction available at seven has no literal K_{11} analogue.

The resulting family conjecture has two layers. Alternating blocks are always norm-one-torus translations. Their reduced Arakelov drift should be governed by the Hodge defect Δ_q: contraction is expected on the side $q \le 9$, while escape is permitted once $q \ge 11$. The heptagonal case has the extra accident that the torus splits, turning its neutral part into the elementary signed-height calculation proved below. The $p = 11$ experiment therefore clarifies the conjecture; it does not merely present a harder version of the same calculation.

Proposition 6.2 (Exact transport of the heptagonal split characters)**.** *For a Rosen remainder recurrence*

$$r_{j+1} = -a_j\lambda_7 r_j - r_{j-1},$$

put

$$\begin{aligned} A_j &= r_j + \alpha r_{j-1}, \quad B_j = r_j - \alpha^{-1} r_{j-1}, \\ C_j &= r_j - \alpha r_{j-1}, \quad D_j = r_j + \alpha^{-1} r_{j-1}, \end{aligned}$$

where $\alpha = \lambda_7^2 - 1$. *Thus*

$$F_7(r_j, r_{j-1}) = A_j B_j, \qquad F_7(r_j, -r_{j-1}) = C_j D_j.$$

For the digit $a_j = 1$,

$$A_{j+1} = \alpha^{-1} C_j, \qquad B_{j+1} = -\alpha D_j, \tag{27}$$

whereas for $a_j = -1$,

$$C_{j+1} = -\alpha^{-1} A_j, \qquad D_{j+1} = \alpha B_j. \tag{28}$$

For an arbitrary digit the errors in these four identities are, respectively,

$$-(a_j - 1)\lambda_7 r_j, \quad -(a_j - 1)\lambda_7 r_j, \quad -(a_j + 1)\lambda_7 r_j, \quad -(a_j + 1)\lambda_7 r_j.$$

In particular an alternating $1, -1$ *block is exactly the unit-scaling channel of the split torus. Its two-step matrix is*

$$-W_7 = -\begin{pmatrix} 1 + \lambda_7^2 & \lambda_7 \\ \lambda_7 & 1 \end{pmatrix}.$$

Proof. Use $\alpha^{-1} = \alpha - \lambda_7$ and substitute the recurrence. For example, when $a_j = 1$,

$$A_{j+1} = (-\lambda_7 + \alpha) r_j - r_{j-1} = \alpha^{-1}(r_j - \alpha r_{j-1}) = \alpha^{-1} C_j.$$

The other three identities are identical one-line calculations. Before specializing a_j, the difference between the two sides leaves the displayed multiple of $a_j - 1$ or $a_j + 1$. Multiplying the two transition matrices for the digits $1, -1$ gives $-W_7$. □

Proposition 6.3 (Alternating blocks cannot amplify projective height)**.** *Let*

$$I_7 = [-\lambda_7/2, \lambda_7/2)$$

be the interval on which the Rosen map is defined, and put $x_j = r_j/r_{j-1} \in I_7$. Suppose that a consecutive block of digits is made only of 1 and -1, with the signs alternating. There is a constant C_{alt}, depending only on the heptagonal field, such that every state in the block satisfies

$$h(x_{j+k}) \le h(x_j) + C_{\mathrm{alt}}. \tag{29}$$

The constant is independent of the length of the block.

More precisely, at every second state in a $1, -1$ block the coordinate

$$z_j = \frac{C_j}{D_j} = \frac{x_j - \alpha}{x_j + \alpha^{-1}}$$

satisfies

$$z_{j+2m} = \alpha^{-4m} z_j, \qquad h(z_{j+2m}) \le h(z_j). \tag{30}$$

For a $-1, 1$ block the same statements hold with

$$w_j = \frac{A_j}{B_j} = \frac{x_j + \alpha}{x_j - \alpha^{-1}}$$

in place of z_j.

Proof. Apply Proposition 6.2 twice. Along a $1, -1$ pair it gives

$$C_{j+2} = -\alpha^{-2} C_j, \qquad D_{j+2} = -\alpha^2 D_j,$$

and hence $z_{j+2} = \alpha^{-4} z_j$. Along a $-1, 1$ pair it gives the identical formula for $w = A/B$.

We prove the height assertion for z; the proof for w is the same. Put $\kappa = \alpha^{-4}$ and $b_i = \log|\sigma_i(\kappa)|$. Directly from the three real conjugates of λ_7,

$$b_0 < 0, \qquad b_1 > 0, \qquad b_2 > 0, \qquad b_0 + b_1 + b_2 = 0.$$

The last equality is the product formula: κ is a unit. The Rosen interval supplies exactly the sign information which an arbitrary unit orbit would lack. Indeed $\alpha - \alpha^{-1} = \lambda_7$, and for every real $x \in I_7$,

$$\left| \frac{x - \alpha}{x + \alpha^{-1}} \right| \ge 1. \tag{31}$$

If $x + \alpha^{-1} \ge 0$, this is the inequality $2x \le \lambda_7$; if $x + \alpha^{-1} < 0$, it is immediate.

Let $u_i = \log|\sigma_i(z_j)|$. Multiplication by κ^m does not change any finite-place contribution to height. At the distinguished real place, both endpoints of the pair block lie in I_7, so (31) says that $u_0 \ge 0$ and $u_0 + mb_0 \ge 0$. Therefore its contribution to the change of height is exactly mb_0. At either outer real place,

$$\max(0, u_i + mb_i) - \max(0, u_i) \le mb_i.$$

Adding the three inequalities and using (6.1) proves $h(\kappa^m z_j) \le h(z_j)$. If $D_j = 0$, the projective coordinate $[C_j : D_j] = [1 : 0]$ is fixed by the two-step diagonal action, so the same conclusion holds without using an affine ratio.

Finally, $x \mapsto (x - \alpha)/(x + \alpha^{-1})$ and its inverse are fixed fractional-linear maps over K_7. Weil height changes by at most a field constant under either map. The intervening odd state is obtained from the preceding even state by one of the two fixed maps $x \mapsto \mp\lambda_7 - x^{-1}$. This proves (29) for every state in the block. For $-1, 1$, use w; the central inequality $|x + \alpha| \ge |x - \alpha^{-1}|$ follows from $x \ge -\lambda_7/2$. □

This proposition is an elementary product-formula calculation, not the missing global estimate. It shows that a long string of alternating unit digits may be arbitrarily long, but its length does not consume any new height budget. Cut a Rosen word immediately before and after each digit of absolute value greater than one, and between two equal adjacent unit digits. Call the resulting cut vertices *junction states*; the pieces between them are alternating blocks. Then Proposition 6.3 gives

$$\max_{0\le k\le n} h(x_k) \le \max_{\substack{b\le n\\ b \text{ a junction}}} h(x_b) + C_{\mathrm{alt}}. \tag{32}$$

Consequently the junction-only estimate

$$h(x_b) \le 2h(x_0) + B_7 \quad \text{for every junction state } b \tag{33}$$

is sufficient for the all-time split target (40), by Corollary 6.8, and hence for the classification. The Pisot-unit family recorded below produces a growing tested excess over its proposed right side, so it is withdrawn as a route. The corridor calculation therefore removes only the unbounded duration of an alternating block; it does not supply a usable bound at the junction where the next block begins.

Why the absolute factor norm is not the missing height. The exact program `rosen_split_torus_double_coset_audit.py` enumerates the free-product normal forms in $G_7 = \langle S, R \mid S^2 = R^7 = 1\rangle$. For $e = (\alpha, 1)$ and $W_7 = SRSR^{-1}$, it quotients by both actions of $\langle W_7\rangle$. Through word lengths $0, 2, 4, 6, 8, 10$, the numbers of double cosets with $|\operatorname{Nm}\det(e, ge)| \le 1$ are

$$1, \quad 5, \quad 11, \quad 17, \quad 33, \quad 87.$$

For bounds 13 and 100, the final counts are 510 and 1528. This finite calculation is not a proof that a bounded-norm set is infinite. It is a sharp negative control on the proposed norm-only argument: proliferation persists after the cyclic normalizer has been removed. An algebraic integer can have norm one while its three real sizes diverge. Thus Proposition 6.2 identifies the correct special channel, but an all-place metric is still required to force a general stopped word into that channel.

What this changes in the conjectural picture. Formula (23) is uniform, but only at $p = 7$ does it split into two K_p-valued characters. To state the uniform all-place object, let ι be the involution of E_p and put

$$u_p(X, Y) = \frac{X + tY}{X + \iota(t)Y} \in \mathcal{T}_p(K_p), \qquad \mathcal{T}_p = \operatorname{Res}^1_{E_p/K_p} \mathbf{G}_m.$$

The projective Weil height

$$\mathfrak{h}_p(X, Y) = h_{E_p}\big([X + tY : X + \iota(t)Y]\big)$$

includes the finite divisor and every real size. At $p = 7$ it is the metrized ratio of the two factors in (25); at $p = 11$ it remains a K_{11}-defined point of the nonsplit norm-one torus although the individual characters are not defined over K_{11}.

The heptagonal target is therefore a *metrized split-character stopping theorem*, modulo the cyclic scaling in Proposition 6.2. A mere bound for $|\operatorname{Nm}(X + \alpha Y)|$ is insufficient. The uniform prime-family conjecture is not that $\mathfrak{h}_p$ stays bounded for all p. Rather, the Hodge-defect conjecture below predicts that the reduced torus height cannot escape along a trace-field orbit when

$$\Delta_p = \frac{-g^2 + 4g - 2}{2g - 1} > 0, \qquad p = 2g + 1,$$

but that some trace-field orbit does escape when $\Delta_p < 0$. The last prime with positive defect is $p = 7$, and the first prime with negative defect is $p = 11$. Independently, Proposition 6.1 says that $p = 7$ is the unique prime with a split adjacent torus. These are not the same criterion: $p = 5$ has positive Hodge defect but a nonsplit adjacent torus. What the consecutive-prime comparison shows is that both heptagonal advantages are lost upon passing from 7 to 11. The Hodge sign is the plausible family-level criterion; splitting is extra structure available only in the $p = 7$ proof.

There is no formal way to recover the two heptagonal norm minima merely by passing to the splitting field at $p = 11$. When E_p is a field, the two conjugate factors have the identical absolute norm

$$\left|\mathrm{N}_{E_p/\mathbf{Q}}(X + tY)\right| = \left|\mathrm{N}_{K_p/\mathbf{Q}}\, F_p(X, Y)\right| = \left|\mathrm{N}_{E_p/\mathbf{Q}}(X + \iota(t)Y)\right|. \tag{34}$$

Thus the choice of the smaller of two K_p-character norms, which is the input to Theorem 6.6, exists only in the split algebra $E_7 = K_7 \times K_7$. In the nonsplit cases one must retain the archimedean distribution of the conjugate factors; that distribution is exactly the metric on the norm-one torus. This is the precise sense in which the uniform conjecture is Hodge–Arakelov, whereas the final $p = 7$ closing argument is an S-unit argument.

Proposition 6.4 (The split estimate leaves a sparse unit-return problem)**.** *Keep the four incidences*

$$A = r + \alpha s, \quad B = r - \alpha^{-1} s, \quad C = r - \alpha s, \quad D = r + \alpha^{-1} s.$$

For every $B_0 \geq 1$, the projective states satisfying

$$\min\{|\,\mathrm{N}\, A|, |\,\mathrm{N}\, B|\} \leq B_0, \qquad \min\{|\,\mathrm{N}\, C|, |\,\mathrm{N}\, D|\} \leq B_0$$

belong to a finite union of sets

$$\Phi(c\,\mathcal{O}_K^\times), \qquad c \in K^\times,$$

where Φ is one of the four explicit Möbius transformations

$$\alpha\frac{z+1}{z-1}, \qquad \frac{\alpha - \alpha^{-1} z}{z-1}, \qquad \frac{\alpha z - \alpha^{-1}}{z-1}, \qquad \alpha^{-1}\frac{z+1}{1-z}.$$

For each member of this finite union,

$$h([r : s]) = h(\epsilon) + O_{B_0}(1), \qquad \epsilon \in \mathcal{O}_K^\times.$$

Proof. Choose a bounded factor from $\{A, B\}$ and one from $\{C, D\}$. Their coefficient determinants are

$$-2\alpha, \qquad -\lambda, \qquad -\lambda, \qquad 2\alpha^{-1},$$

respectively. There are only finitely many integral principal ideals of absolute norm at most B_0. After choosing one generator for each, write the two selected factors as

$$L = \beta u, \qquad M = \gamma v, \qquad u, v \in \mathcal{O}_K^\times.$$

Then $L/M = (\beta/\gamma)(u/v)$ lies in one of finitely many unit cosets. Solving the four two-by-two systems for r/s gives the displayed transformations. A fixed Möbius transformation changes absolute logarithmic height by $O(1)$. □

Proposition 6.5 (One-divisor quadratic reformulation)**.** *Define the four fixed quadratic forms*

$$Q_{AC}(r,s) = r^2 - \alpha^2 s^2, \qquad Q_{AD}(r,s) = r^2 + (\alpha + \alpha^{-1})rs + s^2,$$
$$Q_{BC}(r,s) = r^2 - (\alpha + \alpha^{-1})rs + s^2, \qquad Q_{BD}(r,s) = r^2 - \alpha^{-2}s^2.$$

If none of (A,B,C,D) vanishes, then

$$\min_{\mu \in \{AC,AD,BC,BD\}} |\,\mathrm{N}\, Q_\mu(r,s)| = \min\{|\,\mathrm{N}\, A|, |\,\mathrm{N}\, B|\}\, \min\{|\,\mathrm{N}\, C|, |\,\mathrm{N}\, D|\}. \tag{35}$$

In particular, since every nonzero algebraic-integer norm has absolute value at least one,

$$\max\{m_+(r,s), m_-(r,s)\} \le \min_\mu |\,\mathrm{N}\, Q_\mu(r,s)| \tag{36}$$
$$\le \max\{m_+(r,s), m_-(r,s)\}^2, \tag{37}$$

where

$$m_+ = \min\{|\,\mathrm{N}\, A|, |\,\mathrm{N}\, B|\}, \qquad m_- = \min\{|\,\mathrm{N}\, C|, |\,\mathrm{N}\, D|\}.$$

Thus the bounded all-time split estimate is equivalent to boundedness of the minimum norm of these four quadratic forms. Quantitatively, the bound

$$\log \min_\mu |\,\mathrm{N}\, Q_\mu(r_j, r_{j-1})| \le 6h(x) + B_7$$

already implies the target (40)*, while that target implies the same quadratic-form inequality with right side* $12h(x) + 2B_7$*.*

Proof. Direct multiplication gives

$$Q_{AC} = AC, \qquad Q_{AD} = AD, \qquad Q_{BC} = BC, \qquad Q_{BD} = BD.$$

For four nonnegative real numbers a, b, c, d,

$$\min\{ac, ad, bc, bd\} = \min\{a, b\}\min\{c, d\}.$$

Apply this identity to the four absolute field norms. If $u, v \ge 1$, then $\max\{u, v\} \le uv \le \max\{u, v\}^2$, proving (37) and its stated consequences. If a factor vanishes, the state is already one of the special eigenlines and no inequality is needed. □

The four forms in this proposition are not artificial products chosen after the fact. Proposition 7.9 identifies them, up to fixed scalars, as the fixed-point forms of four explicit hyperbolic elements in the rank-two intermediate group $H_2 = \langle G_7, T(2)\rangle$. Their four axes form a square on the four split-character lines. This is the group-theoretic structure behind the one-divisor rewrite.

Theorem 6.6 (The all-time split-norm inequality is sufficient)**.** *Let* (r_j, r_{j-1}) *be the primitive exact remainder states of* $x \in K$*. If one constant* B_0 *satisfies*

$$\max\{\min(|\,\mathrm{N}\, A_j|, |\,\mathrm{N}\, B_j|), \min(|\,\mathrm{N}\, C_j|, |\,\mathrm{N}\, D_j|)\} \le B_0 \tag{38}$$

at every nonterminal state, then the orbit terminates or repeats. Therefore x *satisfies the heptagonal cusp/special-fixed-point classification.*

Proof. If a selected factor is zero, the state is an eigenline of one of the adjacent special words. Otherwise choose one bounded factor from each orientation and write $X_j = (L_j, M_j)$. The coordinate type is one of AC, AD, BC, BD. Bounded absolute norm leaves finitely many principal ideals; after choosing their generators, all L_j, M_j lie in one finitely generated subgroup $\Gamma \subset K^\times$.

Let P_μ be the matrix whose two rows are the linear factors of type μ, so $X_j = P_{\mu_j}(r_j, r_{j-1})^t$. Write R_μ for the first row of P_μ^{-1} and S_μ for its second row. The shared remainder in two consecutive states gives

$$S_{\mu_{j+1}} X_{j+1} - R_{\mu_j} X_j = 0. \tag{39}$$

For fixed types this is a homogeneous four-term equation in Γ. The projective S-unit theorem leaves finitely many nondegenerate projective solutions, hence finitely many corresponding states.

There are only two nonterminal degeneracy pairings: each pairs one term from time j with one from time $j+1$. Substituting

$$X_{j+1} = P_{\mu_{j+1}} \begin{pmatrix} -n_j\lambda & -1 \\ 1 & 0 \end{pmatrix} P_{\mu_j}^{-1} X_j$$

shows that $z_j = L_j/M_j$ is a rational function of the integer digit n_j. The exact 32-case calculation in `rosen_split_sunit_degeneracy_certificate.py` proves the following dichotomy: either $\mathrm{N}(z_j)$ is a nonconstant quotient of cubic polynomials in n_j, or z_j itself is one of

$$-1, \qquad \lambda^2 + \lambda, \qquad 2 - \lambda.$$

The script works in the exact field $\mathbf{Q}(n)[\lambda]/(\lambda^3 - \lambda^2 - 2\lambda + 1)$, and the three constant values are checked directly. But $\mathrm{N}(z_j) = \mathrm{N}(L_j)/\mathrm{N}(M_j)$ belongs to a finite set determined by B_0. A nonconstant rational function can take one of those values at only finitely many integer digits, while a constant z_j already fixes the projective state. Thus degenerate states are finite as well.

Every infinite deterministic orbit in this finite state set repeats. The closing criterion gives the special hyperbolic alternative. □

The proposed uniform $6h$ estimate is withdrawn. Theorem 6.6 remains a valid conditional statement, but the quantitative estimate proposed for its hypothesis,

$$\log \max\{\min(|\,\mathrm{N}\,A_j|, |\,\mathrm{N}\,B_j|), \min(|\,\mathrm{N}\,C_j|, |\,\mathrm{N}\,D_j|)\} \le 6h(x) + B_7 \tag{40}$$

is contradicted by the numerical claim previously made here. Put

$$U = \lambda\alpha, \qquad x_N = U^{-N}.$$

The element U is a Pisot unit at the distinguished embedding: $|\sigma_0 U| > 1$ and $|\sigma_1 U|, |\sigma_2 U| < 1$. The exact recurrence in `rosen_split_factor_all_states_audit.py` gives the following maximal excesses in the first one hundred states:

N	10	20	40	60
$\max_j\{E_j - 6h(x_N)\}$	1.2821496	3.5023095	12.4091459	18.2149036,

where E_j denotes the logarithm on the left of (40). All norm and recurrence calculations are exact; only the displayed logarithms are numerical. The first large digits arise because traces of powers of the Pisot unit approximate its distinguished real power exponentially closely.

It also gives counterevidence to the stronger junction-height proposal (33): for $N = 20, 40, 60$, a state among the first five has height respectively

$$2.1237\,h(x_N), \qquad 2.1935\,h(x_N), \qquad 2.2168\,h(x_N).$$

No asymptotic lower bound for these excesses is proved here. These particular orbits nevertheless close in the tested cases; for example $N = 20$ becomes period two and $N = 40$ terminates. Hence large split norm or height at a junction is compatible with the desired classification. A successful inequality must recognize the Pisot trace cancellation or the subsequent closing word, rather than impose a uniform $6h$ ceiling. Accordingly neither quantitative estimate is used as a proposed route in the status summary.

Proposition 6.7 (Equivalent two-axis proximity formulation)**.** *For a primitive state $X = (r, s)^t$, put*

$$e_A = [-\alpha : 1], \quad e_B = [\alpha^{-1} : 1], \quad e_C = [\alpha : 1], \quad e_D = [-\alpha^{-1} : 1] \in \mathbf{P}^1(K),$$

and, for $e = [u : v]$, define its total archimedean projective proximity by

$$\Lambda_e(X) = \sum_{i=0}^{2} \log \frac{\max(|\sigma_i r|, |\sigma_i s|) \max(|\sigma_i u|, |\sigma_i v|)}{|\det(\sigma_i X, \sigma_i e)|}.$$

Then

$$\begin{aligned} &\log \max\{\min(|\operatorname{N} A|, |\operatorname{N} B|), \min(|\operatorname{N} C|, |\operatorname{N} D|)\} \\ &\quad = 3h(X) + 3h(\alpha) - \min\{\max(\Lambda_{e_A}(X), \Lambda_{e_B}(X)), \max(\Lambda_{e_C}(X), \Lambda_{e_D}(X))\}. \end{aligned} \tag{41}$$

Consequently (40) is equivalent, after changing the field constant, to

$$\min\{\max(\Lambda_{e_A}(X_j), \Lambda_{e_B}(X_j)), \max(\Lambda_{e_C}(X_j), \Lambda_{e_D}(X_j))\} \geq 3h(X_j) - 6h(x) - B_7'. \tag{42}$$

Proof. The four determinants $\det(X, e_A), \ldots, \det(X, e_D)$ are respectively A, B, C, D. The global incidence identity gives

$$\log |\operatorname{N} \det(X, e)| = 3h(X) + 3h(e) - \Lambda_e(X).$$

All four endpoint heights equal $h(\alpha)$, because changing sign or inverting an algebraic number does not change its projective height. Taking the two minima and then their maximum gives (41); rearrangement gives (42). □

This identity locates the remaining Hodge–thin-group statement. The pairs $\{e_A, e_B\}$ and $\{e_C, e_D\}$ are the endpoint pairs of the two adjacent special axes. A high state must enter, in total three-place proximity, one endpoint tube from each pair. Their four possible intersections are exactly the four unit loci in Proposition 6.4. Hodge theory must supply the proximity growth; the consecutive thin recurrence then supplies the fixed S-unit equation. Neither statement is being substituted for the other.

Corollary 6.8 (Why the split target is weaker than a height barrier)**.** *For every primitive state X,*

$$\log \max\{\min(|\operatorname{N} A|, |\operatorname{N} B|), \min(|\operatorname{N} C|, |\operatorname{N} D|)\} \leq 3h(X) + O_7(1). \tag{43}$$

Consequently the quadratic excursion estimate $h(X_j) \leq 2h(x) + O_7(1)$ implies (40). The converse fails at the level of height geometry: bounded split minima place X_j in the four unit corridors of Proposition 6.4, and the height can diverge along such a corridor.

Proof. For each of the four fixed linear factors L, the triangle inequality at the three real places gives

$$\log |\operatorname{N} L(X)| \leq 3h(X) + O_L(1).$$

Taking minima and a maximum proves (43). The last assertion follows from the explicit parametrizations $X = \Phi(c\epsilon)$, for which $h(X) = h(\epsilon) + O(1)$ while ϵ ranges through the rank-two unit group. □

Thus one need not prove that the Rosen orbit remains in a Northcott bounded set. It is enough to prove that every high excursion lies in one of four arithmetically explicit rank-two corridors; Theorem 6.6 then uses consecutive states to prevent indefinite travel through those corridors unless the orbit closes.

What remains if the split bound is known only at records. Choose infinitely many record states in one component of Proposition 6.4 and write them as $x_\nu = \Phi(\epsilon_\nu)$. If G_ν is the intervening Rosen word and

$$\Phi^{-1} G_\nu \Phi = \begin{bmatrix} a_\nu & b_\nu \\ c_\nu & d_\nu \end{bmatrix},$$

then

$$c_\nu \epsilon_\nu \epsilon_{\nu+1} + d_\nu \epsilon_{\nu+1} - a_\nu \epsilon_\nu - b_\nu = 0. \tag{44}$$

Assume the standard four-summand number-field *abc* conjecture and the slow-word inequality

$$h([a_\nu : b_\nu : c_\nu : d_\nu]) = o(h([\epsilon_\nu \epsilon_{\nu+1} : \epsilon_{\nu+1} : \epsilon_\nu : 1])). \tag{45}$$

Generalized *abc* makes (44) degenerate. If all four coefficients are nonzero, the only determinant-compatible pairing gives

$$c_\nu \epsilon_\nu \epsilon_{\nu+1} = b_\nu, \qquad d_\nu \epsilon_{\nu+1} = a_\nu \epsilon_\nu,$$

and hence $\epsilon_\nu^2 = b_\nu d_\nu / (a_\nu c_\nu)$, contradicting (45). A three-term degeneration is impossible. Thus the conjugated word is diagonal or antidiagonal: the intervening thin-group element normalizes one fixed K-split torus. Its intersection with the discrete Fuchsian group is finite or virtually cyclic. Infinitely many Rosen prefixes in one such coset either repeat or converge to an endpoint of its axis. In both cases the original point is a cusp or a translate of a K-rational hyperbolic eigenline.

This proves the classification on the slow-word branch. It is also a precise account of what generalized *abc*, stronger than ordinary Roth, contributes: it selects the torus normalizer once (45) is known. The remaining heptagonal task is to obtain that inequality, or a sufficient positive-margin version, from the stopped Hodge defect. The absolute factor norm does not imply it. In fact `rosen_split_sunit_return_audit.py` finds ratios

$$0.9059, \qquad 0.9794, \qquad 0.9852$$

between the intervening-word height and the four-term unit height on the three repeated-type record pairs of the exponent-80 adversarial unit ray. The sharp family approaches the critical ratio one, not the little-oh regime. Generalized *abc* therefore proves a genuine slow branch, while Theorem 6.6 is the more promising way to avoid moving coefficients altogether.

The relevant codifferent calculation is more informative in the whole prime family. Besides clarifying the role of 11, it identifies the natural higher-rank analogue of the H_4 quotient without assuming that the corresponding translations already lie in the triangle group.

Proposition 6.9 (The prime-family codifferent hyperplane)**.** *Let $p = 2g + 1 \geq 7$ be an odd prime,*

$$K_p = \mathbf{Q}(\lambda_p), \qquad \lambda_p = 2\cos(\pi/p),$$

and let f_p be the minimal polynomial of λ_p. Put $\vartheta_p = \lambda_p / f_p'(\lambda_p)$. Then

$$\mathfrak{D}_{K_p}^{-1} = \vartheta_p \mathcal{O}_{K_p}, \qquad \mathrm{Tr}(\vartheta_p) = 0.$$

If

$$\kappa = a_0 + a_1 \lambda_p + \cdots + a_{g-1} \lambda_p^{g-1},$$

then the primitive integral functional

$$d_p(\kappa) := \operatorname{Tr}(\vartheta_p \kappa) = a_{g-2} + a_{g-1}. \tag{46}$$

For every $N \geq 1$, the rank-$(g-1)$ module

$$M_{p,N} = \mathbf{Z} + N\big(\mathbf{Z}\lambda_p + \cdots + \mathbf{Z}\lambda_p^{g-3} + \mathbf{Z}(\lambda_p^{g-2} - \lambda_p^{g-1})\big) \subseteq \ker d_p \tag{47}$$

(with the middle sum empty when $g = 3$) satisfies

$$\mathcal{O}_{K_p}/M_{p,N} \simeq \mathbf{Z} \oplus (\mathbf{Z}/N\mathbf{Z})^{g-2}. \tag{48}$$

Consequently normalization by $M_{p,N}/\mathbf{Z}$ reduces the original rank-$(g-1)$ correction lattice $\mathcal{O}_{K_p}/\mathbf{Z}$ to one unbounded coordinate d_p and $g-2$ finite labels.

Order the real embeddings by

$$\lambda_{p,j} = 2\cos((2j-1)\pi/p), \qquad 1 \leq j \leq g.$$

Then

$$\operatorname{sgn} \sigma_j(\vartheta_p) = \begin{cases} (-1)^{j-1}, & j \leq \lceil g/2 \rceil, \\ (-1)^{j}, & j > \lceil g/2 \rceil. \end{cases} \tag{49}$$

Among $g \geq 3$, the pattern consists of one positive sign followed only by negative signs if and only if $g = 3$, equivalently $p = 7$.

Proof. The maximal real cyclotomic order is $\mathcal{O}_{K_p} = \mathbf{Z}[\lambda_p]$. If $\zeta = e^{\pi i/p}$, then $\lambda_p = \zeta^{-1}(1+\zeta^2)$; since $\mathrm{N}_{\mathbf{Q}(\zeta^2)/\mathbf{Q}}(1+\zeta^2) = \Phi_p(-1) = 1$, λ_p is a unit. Therefore the usual monogenic formula $\mathfrak{D}^{-1} = f_p'(\lambda_p)^{-1}\mathcal{O}_{K_p}$ gives the first identity. The conjugates of λ_p are indexed by the odd nonmultiples of p modulo $2p$. Their sum is 1: they are the roots of $X^p + 1$ other than -1, grouped into inverse pairs. Hence the coefficient of X^{g-1} in f_p is -1.

For a monic polynomial of degree g, the residue identities give

$$\operatorname{Tr} \frac{\lambda_p^j}{f_p'(\lambda_p)} = 0 \quad (0 \leq j \leq g-2), \qquad \operatorname{Tr} \frac{\lambda_p^{g-1}}{f_p'(\lambda_p)} = 1.$$

Reducing λ_p^g by f_p, its λ_p^{g-1}-coefficient is 1. Applying the displayed identities to $\lambda_p\kappa/f_p'(\lambda_p)$ proves (46), including $\operatorname{Tr}(\vartheta_p) = 0$. The displayed generators in (47) form N times a primitive basis of $\ker d_p/\mathbf{Z}$. Smith normal form gives (48).

The roots $\lambda_{p,j}$ are strictly decreasing. Therefore $\operatorname{sgn} f_p'(\lambda_{p,j}) = (-1)^{j-1}$. Exactly $\lceil g/2 \rceil$ of the roots are positive. Division gives (49). For $g = 3$ the signs are $(+,-,-)$; for even $g \geq 4$ the last sign is positive, and for odd $g \geq 5$ the third sign is positive. This proves the last assertion. □

For $p = 7, N = 4$, $M_{p,N}$ is exactly the module M_4 in (99). For $p = 11$, it is

$$M_{11,N} = \mathbf{Z} + N\mathbf{Z}\lambda_{11} + N\mathbf{Z}\lambda_{11}^2 + N\mathbf{Z}(\lambda_{11}^3 - \lambda_{11}^4), \qquad \mathcal{O}_{K_{11}}/M_{11,N} \simeq \mathbf{Z} \oplus (\mathbf{Z}/N\mathbf{Z})^3.$$

This is the natural degree-five intermediate *reduction module*: it removes three of the four correction directions and leaves

$$d_{11} = a_3 + a_4, \qquad \big(\operatorname{sgn} \sigma_j(\vartheta_{11})\big)_{j=1}^5 = (+,-,+,+,-).$$

The analogous $p = 7$ signs are $(+, -, -)$. Thus in degree three the one positive trace direction is exactly the uniformizing direction, whereas in degree five the canonical trace coordinate mixes two positive and two negative outer signs.

There is a corresponding candidate cusp enlargement generated by G_{11} and the translations $T(\lambda_{11} M_{11,N})$. Unlike the $p = 7, N = 4$ case, the present manuscript has not identified a level N for which this enlargement is congruence-neutral, nor proved its exact cusp intersection. Even if that group-theoretic step is supplied, equation (49) shows that its remaining scalar direction is not the one-versus-four Hodge direction. The obstruction at 11 is therefore not the absence of a canonical higher-rank module; it is the failure of real Hodge alignment after the module is constructed.

6.2 The five Hodge weights

There is a short triangle-group derivation which works uniformly for odd q. Let

$$\mathcal{M}_q = \{m : 1 \le m < q/2,\ \gcd(m, q) = 1\}.$$

The Galois embeddings of the invariant trace field are indexed projectively by $m \in \mathcal{M}_q$. In the corresponding real representation, the elliptic generators have rotation numbers $1/2$ and m/q. With the cusp orientation fixed by the positive parabolic, the relative Euler number is therefore

$$e_m = 1 - \frac{1}{2} - \frac{m}{q} = \frac{q - 2m}{2q}. \tag{50}$$

The factor $m = 1$ is uniformizing. Dividing by its Euler number gives the normalized Hodge, or Lyapunov, weight

$$\chi_m = \frac{e_m}{e_1} = \frac{q - 2m}{q - 2}. \tag{51}$$

This is the same degree calculation used in (12); it is not an assumption about digit frequencies.

For $q = 11$, the projective indices are $m = 1, 2, 3, 4, 5$, and (51) gives

$$(\chi_1, \dots, \chi_5) = \left(1, \frac{7}{9}, \frac{5}{9}, \frac{3}{9}, \frac{1}{9}\right). \tag{52}$$

Point a geodesic at a field endpoint and form the full integral Asai tensor of its five conjugate endpoint vectors. The central factor has formal exponent -1, while an Oseledets-regular outer factor has top exponent χ_m. The predicted exponent of the tensor is consequently

$$-1 + \frac{7}{9} + \frac{5}{9} + \frac{3}{9} + \frac{1}{9} = \frac{7}{9}. \tag{53}$$

At $p = 7$ the identical calculation gives $-1/5$. The sign, rather than the number of terms, is the important difference.

It is convenient to name this sign. Define the *integral Hodge defect*

$$\Delta_q = 1 - \sum_{\substack{m \in \mathcal{M}_q \\ m \neq 1}} \chi_m. \tag{54}$$

The ordinary integral Asai tensor has formal exponent $-\Delta_q$.

Proposition 6.10 (The scalar threshold occurs exactly between 9 and 11)**.** *For odd $q \ge 3$,*

$$\Delta_q > 0 \quad \Longleftrightarrow \quad q \in \{3, 5, 7, 9\}. \tag{55}$$

In particular, $q = 11$ is the first odd parameter with an expanding full integral Asai endpoint tensor.

Proof. For $q = 3, 5, 7, 9$, the sets of Hodge indices are respectively

$$\{1\}, \qquad \{1, 2\}, \qquad \{1, 2, 3\}, \qquad \{1, 2, 4\},$$

and direct substitution in (54) gives

$$\Delta_3 = 1, \qquad \Delta_5 = \frac{2}{3}, \qquad \Delta_7 = \frac{1}{5}, \qquad \Delta_9 = \frac{1}{7}.$$

Now let $q \geq 11$. The index 2 always occurs. If $3 \nmid q$, the index 3 also occurs, and the two corresponding outer weights already have sum

$$\frac{q-4}{q-2} + \frac{q-6}{q-2} = \frac{2q-10}{q-2} > 1.$$

If $3 \mid q$, then $q \geq 15$, the index 4 occurs, and the weights with indices $2, 4$ have sum

$$\frac{q-4}{q-2} + \frac{q-8}{q-2} = \frac{2q-12}{q-2} > 1.$$

Thus in either case the sum of all outer weights is greater than one, so $\Delta_q < 0$. □

The first cases around the threshold are

q	$[K_q : \mathbf{Q}]$	normalized Hodge weights	Δ_q	Asai exponent
7	3	$1, 3/5, 1/5$	$1/5$	$-1/5$
9	3	$1, 5/7, 1/7$	$1/7$	$-1/7$
11	5	$1, 7/9, 5/9, 3/9, 1/9$	$-7/9$	$7/9$

For an odd prime $p = 2g + 1$, one has $\mathcal{M}_p = \{1, \ldots, g\}$, and summing the odd numerators gives

$$\sum_{m=2}^{g} \chi_m = \frac{(g-1)^2}{2g-1}, \qquad \Delta_p = \frac{-g^2 + 4g - 2}{2g-1}. \tag{56}$$

Thus $p = 7$ is the last odd prime for which the *one-line* full endpoint Asai tensor has negative formal exponent. The next prime, $p = 11$, is already on the other side of this scalar threshold. This is a limitation of that particular height, not yet a prediction about the Rosen orbit itself.

> **Deep input and elementary consequence.** That the conjugate representations occur as polarized rank-two Hodge systems is geometric. Once this is known, their degrees are the relative Euler numbers in (50). The summation leading to (56) is elementary, but its sign decides whether the product formula and generic Hodge contraction reinforce or oppose each other in the full scalar Asai line.

6.3 The threshold is an average, not a uniform contraction

The sign of Δ_p must not be read as a pointwise statement about every geodesic of the triangle orbifold. The following exact calculation is the reason that the comparison with $p = 11$ leads to an arithmetic selection conjecture rather than to a uniform hyperbolicity argument.

For a word $W \in G_7$ which is hyperbolic at all three real embeddings, write

$$a_i(W) = \operatorname{arcosh} \frac{|\sigma_i(\operatorname{tr} W)|}{2}, \qquad \chi_{\mathrm{sc}}(W) = -a_0(W) + a_1(W) + a_2(W).$$

Thus χ_{sc} is the period of the scalar endpoint-Asai potential around the closed geodesic defined by W.

Proposition 6.11 (The heptagonal scalar defect is not uniform)**.** *With the six induced matrices $A_j = -SR^j$ defined above, put*

$$W_8 = A_1^8 A_6 A_5 A_6 A_1 A_2.$$

Then W_8 is hyperbolic at all three embeddings and

$$\chi_{\mathrm{sc}}(W_8) = 0.0638765195\ldots > 0. \tag{57}$$

Its two fixed points lie in a quadratic extension of K_7, not in $\mathbf{P}^1(K_7)$. Consequently no continuous transfer function can turn the scalar endpoint-Asai potential into a nonpositive potential on every arrow of the full compactly induced geodesic coding.

Proof. Direct multiplication gives

$$\operatorname{tr} W_8(x) = 72x^4 + 2.$$

For the three roots $x_i = \sigma_i(\lambda)$, put $u_i = 36x_i^4 + 1$, so that $a_i = \operatorname{arcosh}(u_i)$. Eliminating x with $x^3 - x^2 - 2x + 1 = 0$ gives

$$\prod_{i=0}^{2}(U - u_i) = U^3 - 471U^2 + 34635U - 80821.$$

In particular,

$$\sum_i u_i^2 - 2u_0u_1u_2 - 1 = 471^2 - 2(34635) - 2(80821) - 1 = -9072. \tag{58}$$

Moreover $u_0 > u_1u_2$. Indeed $u_0u_1u_2 = 80821$, and reduction in K_7 gives

$$u_0^2 - 80821 = 36(1014\lambda^2 + 686\lambda - 2751) > 0,$$

where the last inequality already follows from $\lambda > 9/5$. Equation (58) is equivalent to

$$(u_1^2 - 1)(u_2^2 - 1) - (u_0 - u_1u_2)^2 = 9072 > 0.$$

Hence

$$\cosh(a_1 + a_2) = u_1u_2 + \sqrt{(u_1^2 - 1)(u_2^2 - 1)} > u_0 = \cosh a_0,$$

which proves (57). The discriminant of the fixed-point equation is $(\operatorname{tr} W_8)^2 - 4$, whose norm is

$$346159153152 = 2^{15}3^6 \cdot 43 \cdot 337.$$

It is not a rational square, so the discriminant is not a square in K_7. The fixed points are therefore not K_7-rational.

Finally, if $\varphi + u - u \circ T \le 0$ for a continuous transfer function u, summing around the periodic orbit of W_8 would give $\chi_{\mathrm{sc}}(W_8) \le 0$, a contradiction. □

The proposition separates the geometric and arithmetic assertions. The number $-\Delta_7 = -1/5$ is the scalar exponent for the natural invariant measure; it is not the maximum exponent over invariant measures. Wrong-sign periodic measures already occur at $p = 7$, but their endpoints need not be rational over the trace field. A K_7-rational point can shadow such a quadratic closed geodesic for a long finite time. The required theorem must show that it cannot shadow successive wrong-sign regimes forever unless the complete Rosen state actually repeats. Compactness and ordinary

Livšic theory cannot make that deduction because $\mathbf{P}^1(K_7)$ is dense and is not a closed topological subsystem. A height, product-formula, or Arakelov input is therefore logically indispensable.

This also states exactly what $p = 11$ tests. At $p = 7$, positive scalar drift is exceptional in the invariant-measure space and must be arithmetically excluded for nonperiodic trace-field endpoints. At $p = 11$, positive scalar drift is the natural Hodge average itself, so there is no analogous scalar contradiction to exclude. The rank-four pairwise bundle remains negatively curved, but using it would require the stronger changing-line transport theorem formulated below.

6.4 What degree five reveals: the natural rank is the unit rank

The scalar calculation suppresses the relative sizes of the conjugate factors. Those relative sizes form a vector of dimension $d - 1$, where $d = [K_q : \mathbf{Q}]$; this is also the rank of the unit group. In degree three there are two such coordinates. They are the two summands of the rank-two pairwise bundle in Section 17.1. For $p = 11$ there are four, not two.

For an odd prime $p = 2g + 1$, let $\mathcal{L}_m$, $1 \le m \le g$, be the Hodge lines indexed as in (51); thus $\mathcal{L}_1$ is the uniformizing line. Put

$$\mathcal{N}_m = \mathcal{L}_m \otimes \mathcal{L}_1^{-1} \quad (2 \le m \le g), \qquad \mathcal{E}_p^{\mathrm{pair}} = \bigoplus_{m=2}^{g} \mathcal{N}_m. \tag{59}$$

Proposition 6.12 (The pairwise bundle has the right rank and sign)**.** *The bundle $\mathcal{E}_p^{\mathrm{pair}}$ has rank $g - 1$. Its summands satisfy*

$$\deg \mathcal{N}_m = -\frac{m-1}{2p}, \qquad \frac{-\deg \mathcal{N}_m}{\deg \mathcal{L}_1} = 1 - \chi_m = \frac{2(m-1)}{p-2} > 0. \tag{60}$$

Along a unit-speed geodesic in the uniformizing factor, their logarithmic norms have pointwise derivative bounded above by

$$-1 + \|DF_m\|_{\mathrm{hyp}} < 0. \tag{61}$$

Moreover, if P_p is the order-p orbifold point and $\Xi_m = \theta_m/\theta_1$, then

$$\operatorname{div}(\Xi_m) = (m-1)P_p, \qquad (\det \mathcal{E}_p^{\mathrm{pair}})^{-2} \simeq \mathcal{O}\left(\binom{g}{2} P_p\right). \tag{62}$$

In normalized units the total pairwise Hodge gap is

$$\sum_{m=2}^{g} (1 - \chi_m) = \frac{g(g-1)}{2g-1} > 0. \tag{63}$$

Thus there is no sign transition for the full pairwise bundle.

Proof. The relative Euler number is $e_m = (p - 2m)/(2p)$, and a polarized rank-two Hodge system has $\deg \mathcal{L}_m = e_m/2$. Hence

$$\deg \mathcal{N}_m = \frac{e_m - e_1}{2} = -\frac{m-1}{2p}.$$

Division by $\deg \mathcal{L}_1 = (p-2)/(4p)$ gives (60). Schwarz–Pick gives (61). At the order-p point the m-th Schwarz map has local degree m, so its derivative relative to the uniformizing derivative vanishes to order $m - 1$. This proves the first identity in (62); summing over m proves the second. Finally,

$$\sum_{m=2}^{g} \frac{2(m-1)}{2g-1} = \frac{g(g-1)}{2g-1}.$$

□

The rank $g-1$ is forced arithmetically as well. If $t_m = \log|\sigma_m(\varepsilon)|$ for a unit, then $\sum_m t_m = 0$. The map

$$(t_1, \dots, t_g) \longmapsto (t_2 - t_1, \dots, t_g - t_1) \tag{64}$$

is an isomorphism from the hyperplane $\sum t_m = 0$ to $\mathbf{R}^{g-1}$. Indeed, if $u_m = t_m - t_1$, then

$$t_1 = -\frac{1}{g}\sum_{m=2}^{g} u_m, \qquad t_m = u_m - \frac{1}{g}\sum_{j=2}^{g} u_j. \tag{65}$$

Consequently all $g-1$ pairwise coordinates control the complete unit shape, whereas their scalar sum has a kernel of dimension $g-2$. The passage from rank two at $p=7$ to rank four at $p=11$ is therefore not cosmetic: it is exactly Dirichlet's unit theorem.

There is an exact integral carrier for these coordinates. It is important not to take only the determinant of the pairwise bundle. That determinant puts $g-1$ central factors and one copy of every outer factor into one tensor; making it homogeneous introduces the moving incidence to the power $g-2$. For $g=5$ this is a cubic pole. The correct construction keeps the pairwise coordinates projectively distinct.

Proposition 6.13 (Odd-degree Asai wedge: one radial divisor and all shape coordinates)**.** *Let $K/\mathbf{Q}$ be totally real Galois of odd degree g, let $V_j \simeq \mathbf{R}^2$ be its real factors, and put $W = \bigotimes_{j=1}^{g} V_j$. For primitive columns $P, c \in \mathcal{O}_K^2$, write $r = \det(P, c)$ and*

$$Y(P) = \bigotimes_{j=1}^{g} \sigma_j(P), \qquad Y(c) = \bigotimes_{j=1}^{g} \sigma_j(c).$$

There is a canonical decomposition

$$\bigwedge^2 W = \bigoplus_{\substack{S\subset\{1,\dots,g\}\\ |S| \text{ odd}}} \left(\bigotimes_{j\in S} \bigwedge^2 V_j\right) \otimes \left(\bigotimes_{j\notin S} \operatorname{Sym}^2 V_j\right). \tag{66}$$

If $\mathcal{Z}_S(P,c)$ is the S-component of $Y(P)\wedge Y(c)$, then, for the saturated component lattice, its coefficient ideal has the form

$$I(\mathcal{Z}_S) = \mathfrak{c}_S \prod_{j\in S} (\sigma_j r), \tag{67}$$

where $\mathfrak{c}_S$ is fixed, independent of P, c, r, and supported only above 2. Away from 2, $\mathfrak{c}_S = (1)$. The all-factor projection can be normalized so that its component is the scalar

$$\mathcal{Z}_{\{1,\dots,g\}} = \mathrm{N}_{K/\mathbf{Q}}(r). \tag{68}$$

Thus the scalar component records the radial norm once, and the projective tuple of singleton components records the $g-1$ ratios $(\sigma_j r)/(\sigma_1 r)$, hence every unit-shape direction.

Proof. For two vector spaces one has

$$\bigwedge^2 (A\otimes B) \simeq \left(\bigwedge^2 A \otimes \operatorname{Sym}^2 B\right) \oplus \left(\operatorname{Sym}^2 A \otimes \bigwedge^2 B\right).$$

Iteration gives (66): swapping the two decomposable tensors contributes a minus sign precisely when an odd number of factors are alternating. In an alternating factor,

$$\sigma_j(P) \wedge \sigma_j(c) = \sigma_j(r).$$

After these determinant factors are removed, the remaining coordinates are coefficients of products of two primitive linear forms. In odd residue characteristic that product is nonzero, so the residual coefficient vector is primitive. The symmetric/alternating splitting is integral after inverting 2; passing to the saturated integral component lattice can therefore change the content only by a fixed ideal supported above 2. This proves (67) with the stated $\mathfrak{c}_S$. Because g is odd, the all-factor subset occurs in the exterior square, and multiplying its determinant factors, with the projection normalized once and for all, gives (68). □

This proposition constructs, for both $p = 7$ and $p = 11$, the finite object required by the phrase "one divisor." The radial scalar uses each conjugate of (r) once. The fixed two-primary factors contribute only a bounded Arakelov constant. The four relative coordinates at $p = 11$ are retained by projectivizing the five singleton components, not by tensoring them into one line. What is not supplied by the algebraic decomposition is the stopped metric comparison which must force those components into a bounded projective region along a consecutive thin orbit.

For $p = 11$, the complete ledger is

m	χ_m	pairwise gap $1 - \chi_m$	$\mathrm{div}(\Xi_m)$
2	7/9	2/9	P_{11}
3	5/9	4/9	$2P_{11}$
4	3/9	6/9	$3P_{11}$
5	1/9	8/9	$4P_{11}$

Thus $\mathcal{E}_{11}^{\mathrm{pair}}$ has rank four, total normalized Hodge gap 20/9, and determinant Higgs divisor $10P_{11}$. The one-line Asai exponent 7/9 and the rank-four pairwise gap 20/9 answer different questions. The former says that one scalar integral tensor does not contract generically. The latter says that Hodge theory still contracts every relative unit direction.

6.5 Why this does not contradict Galois averaging

One might try to put several copies of the central factor into the tensor and fewer copies of the outer factors. Over $\mathbf{R}$ this would create any desired negative weight. It is not, however, a projective arithmetic height over $\mathbf{Q}$: Galois descent cycles the privileged multiplicity through all five factors.

The following elementary calculation makes that obstruction precise. Index the embeddings by the Galois group and let $\mathbf{a} = (a_\sigma)$ be a tuple of nonnegative integers. The associated tensor monomial is

$$W_{\mathbf{a}} = \bigotimes_{\sigma} \mathrm{Sym}^{a_\sigma} V_\sigma.$$

Its formal endpoint weight along the distinguished geodesic is

$$\ell(\mathbf{a}) = -a_1 + \sum_{\sigma \neq 1} a_\sigma \chi_\sigma.$$

Proposition 6.14 (Galois averaging preserves the Asai sign)**.** *Let $A = \sum_\sigma a_\sigma$. Summing the formal weights of all Galois conjugates of $W_{\mathbf{a}}$ gives*

$$\sum_{\tau \in \mathrm{Gal}(K_q/\mathbf{Q})} \ell(\tau \mathbf{a}) = -A\Delta_q. \tag{69}$$

If repeated conjugates are removed, both sides are divided by the order of the stabilizer of $\mathbf{a}$. Consequently, when $\Delta_q < 0$, at least one summand in every such Galois descent has positive formal weight.

Proof. Galois acts transitively on the embeddings. Hence, for every fixed embedding σ, the sum of its multiplicities over all translates of $\mathbf{a}$ is A. The central coefficient therefore contributes $-A$, and the outer coefficients contribute $A\sum_{\sigma\neq 1}\chi_\sigma$. Their sum is $-A\Delta_q$. If this is positive, not every conjugate summand can have negative weight. □

This proposition is a no-go statement for one-dimensional flat integral heights built functorially from one endpoint line. The pairwise bundle (59) is not such a line. It keeps the full vector of $g-1$ relative Hodge coordinates and couples that vector to the affine determinant of two consecutive states. After Galois completion some real components expand, exactly as the proposition requires; the proposed argument instead uses the whole vector, whose unit-shape map (64) is proper. Its missing assertion is a stopped maximal-slope statement, not simultaneous contraction of every Galois conjugate scalar.

Thus the rigorous conclusion of the averaging proposition is narrower than the earlier draft suggested. At $p = 11$ a flat scalar Asai proof is unavailable. A rank-four pairwise proof remains Hodge-favorable, but it must show that all four channels share the one determinant ideal of the same two consecutive frames. Counting an independent finite divisor in each channel would merely pay four times for four contractions and would prove nothing.

6.6 Exact Rosen experiments

The numerical comparison uses Winsor's convention

$$x = [a_0, a_1, \ldots] = a_0\lambda - \frac{1}{a_1\lambda - \cdots}, \qquad x - a_0\lambda \in (-\lambda/2, \lambda/2].$$

If $x_j = p_j/q_j$, one step is the exact recurrence

$$a_j = \operatorname{nint}(x_j/\lambda), \qquad (p_{j+1}, q_{j+1}) = (-q_j, p_j - a_j\lambda q_j). \tag{70}$$

All pairs were propagated in $\mathbf{Z}[\lambda_{11}]^2$ using (22). The branch choice and every archimedean diagnostic were evaluated at five times the decimal length of the exact coefficients plus one hundred guard digits. Every claimed repetition was tested by the exact identity $p_iq_j = p_jq_i$ in the quintic field.

The results are sharply different even in the smallest rational sample. For $p = 7$, every integer $1 \le x \le 100$ terminates within 61 steps. For $p = 11$, at cutoff 500, only $x = 1$ terminates; none of $x = 2, \ldots, 100$ terminates or repeats. The particular orbit $x = 2$ was continued for 3000 steps. It had no exact repetition with period at most 500, its largest exact coefficient had 1024 decimal digits, and the minimum observed distance from a half-integer branch boundary was $8.59 \cdot 10^{-5}$.

A small-height field census shows that this is not peculiar to rational integers. Write $x = a_0 + a_1\lambda_{11} + \cdots + a_4\lambda_{11}^4$, with each $a_i \in \{-1, 0, 1\}$, and omit $x = 0$. Among all 242 resulting elements, at cutoff 300, exactly

$$\pm 1, \qquad \pm\lambda_{11}, \qquad \pm(1 + \lambda_{11}), \qquad \pm(\lambda_{11} - 1)$$

terminate, none repeats with period at most 250, and the remaining 234 are unresolved. All states and all repetition tests are exact in K_{11}. The cutoff does not certify that the unresolved expansions are infinite; it does show that the behavior of $x = 2$ is dominant rather than exceptional in the smallest coefficient box.

For Winsor's projective height

$$h_\Sigma([p : q]) = \prod_{\sigma : K_{11} \hookrightarrow \mathbf{R}} (|\sigma(p)| + |\sigma(q)|),$$

applied to x/λ, the logarithm had the values

j	0	10	50	100	200	500	1499	2999
$\log h_\Sigma(x_j/\lambda)$	5.733	20.256	54.797	88.723	149.995	409.382	1147.808	2270.282.

Its least-squares slope on the second half of the run was 0.726271 per digit.

The rank-four diagnostic is more informative. Number the embeddings by the Hodge index $m = 1, \dots, 5$, with $m = 1$ uniformizing, and define

$$u_{m,j} = \log|\sigma_m(q_j)| - \log|\sigma_1(q_j)|, \qquad m = 2, 3, 4, 5, \tag{71}$$

$$\eta_{m,j} = \log \frac{\max\{|\sigma_m(p_j)|, |\sigma_m(q_j)|\}}{|\sigma_m(q_j)|}. \tag{72}$$

The vector $(u_{2,j}, \dots, u_{5,j})$ is the complete denominator unit shape. At the midpoint, step 1499, it is

$$(2701.373,\ 2346.897,\ 2013.760,\ 1689.370), \tag{73}$$

and by step 2999 its maximum absolute coordinate is 5406.095. Its four least-squares slopes on the second half are

$$(1.78720,\ 1.56499,\ 1.32739,\ 1.10638). \tag{74}$$

These numbers have a direct Hodge interpretation. A generic denominator is in the stable direction at the uniformizing place, with exponent -1, and in the top direction at the four other places, with exponents χ_m. Its shape ratios should therefore have exponents

$$(1 + \chi_2, 1 + \chi_3, 1 + \chi_4, 1 + \chi_5) = \left(\frac{16}{9}, \frac{14}{9}, \frac{12}{9}, \frac{10}{9}\right), \tag{75}$$

which are $(1.77778, 1.55556, 1.33333, 1.11111)$. The observed slopes in (74) are close to these values. The same model predicts the radial norm exponent $-1 + \sum_{m=2}^{5} \chi_m = 7/9$, exactly the scalar Asai exponent in (53).

One can recover the five individual slopes from the four differences and the radial norm. They are

$$(-1.01195,\ 0.77525,\ 0.55305,\ 0.31544,\ 0.09443).$$

After normalizing the first to -1, the four outer values are $(0.76610, 0.54652, 0.31172, 0.09331)$, to be compared with $(7/9, 5/9, 3/9, 1/9)$. The normalized radial slope is 0.71765, to be compared with $7/9$. The discrepancies are of the size expected from a finite digit run; the ordering and all five signs agree with the Hodge model.

This sign check is the conceptual point. The actual denominator compares a central *stable* vector with outer *unstable* vectors and hence sees $1 + \chi_m$. The negative relative Hodge line compares unstable directions and sees $-1 + \chi_m$. Passing from the former to the latter is not a change of notation; it is the entire missing transport theorem. By contrast, the projective differences $\eta_{1,j} - \eta_{m,j}$, which are the differences of the singleton Asai-wedge proximity terms up to bounded cross-section constants, remain at the scale of individual Rosen excursions: their maximum radius in the run is 9.051, and their radius is 2.093 at step 2999. The same unit escape is visible for $x = 3, 4, 5$: after 300 steps the largest shape coordinate is between 542 and 547, whereas $\log|\mathrm{N}(q_j)|$ is between 219 and 255.

This audit corrected a serious numerical error in the preceding version. The exact recurrence and its branch decisions had been evaluated at high precision, but the returned algebraic roots

were later used for the height calculation under the default 15-digit context. Quintic conjugates of the 512-digit exact coefficients undergo hundreds of digits of cancellation; the resulting false value at step 1499 was 5745.916. Re-evaluating inside the original precision context gives the values above. The failed calculation and its diagnosis are retained in `explore.tex`.

The experiment certifies exactly the finite statements just made. It does not convert 3000 nonrepeating steps into an infinite-orbit theorem. The corrected experiment shows both rising scalar height and linear escape in all four unit-shape coordinates. This is compatible with the positive scalar exponent in (53). More importantly, it shows that negative degrees of the abstract relative Hodge lines do not by themselves control the changing denominator selected by the Rosen algorithm. The program is `rosen_p11_comparison.py`.

6.7 The $q = 9$ control

The sign threshold in Proposition 6.10 would be less persuasive if only 7 and 11 had been computed. The nearest odd case on the contracting side is $q = 9$. It has the same field degree and unit rank as $q = 7$, but is composite and has different elliptic indices:

$$K_9 = \mathbf{Q}(\lambda_9), \qquad \lambda_9 = 2\cos(\pi/9), \qquad \lambda_9^3 - 3\lambda_9 - 1 = 0,$$

$$(\chi_1, \chi_2, \chi_4) = \left(1, \frac{5}{7}, \frac{1}{7}\right), \qquad \Delta_9 = \frac{1}{7}.$$

The exact recurrence in $\mathbf{Z}[\lambda_9]^2$ was run on every integer $1 \le x \le 100$, with the same Rosen convention and exact projective equality test as at 11. Sixty-six orbits terminated and thirty-four became projectively periodic; none remained unresolved after 500 steps. The largest terminating time was 109, for $x = 55$, and the latest detected periodic return occurred by step 107.

The smallest comparison is already explicit. For $x = 2$, the first nine digits are

$$1, -4, 1, 1, 3, -4, 1, 1, 3.$$

If (p_j, q_j) denotes the exact homogeneous state, then

$$(p_9, q_9) = \varepsilon(p_5, q_5), \qquad \varepsilon = 9 + 14\lambda_9 - 10\lambda_9^2, \qquad \mathrm{N}_{K_9/\mathbf{Q}}(\varepsilon) = 1. \tag{76}$$

Thus the expansion has preperiod five and period four. This identity is an equality in the cubic field, not a floating-point recognition. The calculation is reproduced by `rosen_q9_control.py`.

The control does not by itself prove that the defect sign governs every algebraic endpoint, and degree and sign still covary in these smallest examples. It does show that the quick closure seen at 7 persists for the composite case immediately below the exact sign threshold, whereas the first case above that threshold behaves differently. This is the precise, limited evidence supplied by $q = 9$.

The comparison suggests one unifying conjecture at exactly the strength used by the scalar argument.

Conjecture 6.15 (Trace-field Oseledets alternative)**.** *Let q be odd and $x \in \mathbf{P}^1(K_q)$. Then either the Rosen state of x terminates or eventually repeats, or, along its recurrent compact returns, the uniformizing denominator line has exponent -1 and its Galois-conjugate lines have exponents χ_m. Equivalently, the scalar integral endpoint-Asai tensor has exponent $-\Delta_q$.*

Here "has exponent" means only that the logarithmic growth rate divided by central hyperbolic time exists with the displayed value. This is much weaker than asking that every digit block have

its measure-theoretic frequency. Proposition 6.11 shows why the trace-field hypothesis is essential: the same conclusion is false for arbitrary real or quadratic-over-K_q endpoints. The conjecture is an arithmetic selection rule saying that an aperiodic trace-field ray cannot remain forever in a wrong-sign exceptional invariant regime.

The positive-defect consequence is immediate. If $\Delta_q > 0$, the third alternative would make a nonzero normalized integral tensor tend to zero at compact returns, contradicting lattice discreteness. Thus every trace-field endpoint closes. When $\Delta_q < 0$, the same tensor instead grows, so integrality permits the third alternative. To predict a phase transition one adds the existence assertion that at least one trace-field endpoint realizes it on the expanding side.

This gives the clean phase-transition conjecture.

Conjecture 6.16 (Integral Hodge-defect phase transition)**.** *Let q be odd and let $K_q = \mathbf{Q}(2\cos(\pi/q))$. If $\Delta_q > 0$, every point of $\mathbf{P}^1(K_q)$ has a terminating or eventually periodic Rosen expansion. If $\Delta_q < 0$, there is a point of $\mathbf{P}^1(K_q)$ with an infinite aperiodic Rosen expansion. For $q = 11$, the first candidate is $x = 2$.*

There is a precise conditional statement behind the first half.

Proposition 6.17 (Defect principle under algebraic Rosen genericity)**.** *Fix an odd q for which the Galois conjugate rank-two systems occur in the Hodge bundle with normalized exponents χ_m. Suppose every nonterminating, nonperiodic point of $\mathbf{P}^1(K_q)$ is generic for the Bowen–Margulis measure of the Rosen natural extension. If $\Delta_q > 0$, every point of $\mathbf{P}^1(K_q)$ terminates or is eventually periodic. If $\Delta_q < 0$, the same calculation gives no lattice contradiction: a generic denominator has scalar Asai exponent $-\Delta_q > 0$.*

Proof. At the uniformizing place the selected denominator is stable with exponent -1; at every outer place its upper exponent is at most χ_m. The integral Asai vector therefore has upper exponent at most $-1 + \sum_{m\neq 1} \chi_m = -\Delta_q$. If $\Delta_q > 0$, it tends to zero at compact returns, contradicting discreteness of the nonzero integral lattice. Hence the orbit closes. When $\Delta_q < 0$, the displayed upper bound is positive and the contradiction disappears. □

Bowen–Margulis genericity is a convenient sufficient hypothesis for Conjecture 6.15, but it is not the actual target. Only the signs and values of the finitely many conjugate growth exponents are needed. This distinction avoids replacing the heptagonal theorem by an unnecessarily strong normality conjecture for all Rosen digit blocks.

The two assertions in Conjecture 6.16 have different strengths. The positive-defect half is exactly the sort of exceptional-ray rigidity needed for the heptagonal theorem. The negative-defect half asks only for one escaping algebraic endpoint. The sign calculation proves neither assertion, but it gives a uniform quantity with the right observed behavior at $q = 5, 7, 9, 11$.

There is a stronger competing conjecture. It is stated as a structural assertion because that is the level at which the $p = 7$ gap has actually been isolated.

Conjecture 6.18 (One-divisor pairwise fusion)**.** *Let $p = 2g + 1 \equiv 3 \pmod 4$ be prime. For two consecutive reduced frames in a Rosen orbit in $\mathbf{P}^1(K_p)$, form their determinant Δ, the $g-1$ pairwise Hodge channels $\mathcal{N}_m$, and the explicit Asai-wedge carrier of Proposition 6.13. Give its singleton components the Hodge metrics induced from the $\mathcal{N}_m$, and give its scalar component the norm lattice (68). Before an incidence vanishes or the complete affine frame repeats, the stopped maximal-slope vector of this carrier retains a fixed positive proportion of the sum of the $g-1$ contractions in (61).*

An incidence zero is the terminating alternative. Repetition of the complete affine frame is the periodic, special-hyperbolic alternative. The usual lattice-discreteness deduction would therefore

turn Conjecture 6.18 into a termination-or-periodicity theorem for K_p-rational Rosen orbits. At $p = 7$ it is precisely the rank-two stopped fusion problem carried by this manuscript. At $p = 11$ it is a rank-four problem with common Hodge divisor $10P_{11}$.

There are consequently two genuinely different uniform predictions, and the present calculation should not blur them.

1. The *scalar-threshold prediction* treats the sign of Δ_p as decisive and predicts a nonclosing orbit at $p = 11$, with $x = 2$ the first candidate.
2. The *pairwise-fusion prediction* says that the relevant rank is always the unit rank $g - 1$, and predicts termination or periodicity also at $p = 11$; the observed orbit of 2 is then merely a long rank-four test case.

The first prediction is aligned with the raw denominator computation: all four arithmetic shape coordinates drift linearly, despite the negative degree of every relative Hodge line. The second prediction remains logically possible through a genuinely stopped and nonlocal transport theorem; a pointwise identification of the denominator with a decreasing Hodge line is numerically false. Proposition 18.3 supplies one concrete non-scalar mechanism which survives this test: the all-place resultant is controlled by one central commutator trace. The $p = 11$ comparison therefore does not replace the scalar sign barrier by a routine higher-rank argument. It isolates central two-sided axis alignment as the common geometric issue which could still support a uniform theorem.

7 The same question as a thin-group stabilizer problem

There is a group-theoretic reformulation which is elementary, but which changes how the missing estimate should be understood. Let

$$\Gamma_K = \mathrm{SL}_2(\mathcal{O}_K),$$

and let $B \subset \mathrm{SL}_2$ be the upper triangular subgroup. The field K has class number one. One quick proof is Minkowski's: every ideal class contains an integral ideal of norm at most

$$\frac{3!}{3^3}\sqrt{49} = \frac{14}{9} < 2,$$

and hence of norm one.

Take $x = [p : q] \in \mathbf{P}^1(K)$, with (p, q) primitive. There are $a, b \in \mathcal{O}_K$ for which $pb - qa = 1$, so

$$H_x = \begin{pmatrix} p & a \\ q & b \end{pmatrix} \in \Gamma_K, \qquad H_x\infty = x. \tag{77}$$

The conjugate $B_x = H_x B H_x^{-1}$ is the K-rational Borel subgroup which fixes x. Changing the second column of H_x changes the frame but not B_x.

Proposition 7.1 (Borel form of the classification)**.** *Conjecture 1.1 is equivalent to*

$$\big(G_7 \cap B_x(K)\big)/\{\pm 1\} \neq \{1\} \qquad \textit{for every } x \in \mathbf{P}^1(K). \tag{78}$$

Equivalently, for every $H \in \Gamma_K$,

$$H^{-1}G_7H \cap B(K) \quad \textit{contains a projectively nonidentity element.} \tag{79}$$

Proof. An element fixes x precisely when it belongs to $B_x(K)$, which proves the equivalence of the two displayed formulations. If $\gamma \in G_7 \cap B_x(K)$, then integrality of H_x, H_x^{-1}, and γ gives

$$H_x^{-1}\gamma H_x = \begin{pmatrix} \varepsilon & u \\ 0 & \varepsilon^{-1} \end{pmatrix}, \qquad \varepsilon \in \mathcal{O}_K^\times, \quad u \in \mathcal{O}_K. \tag{80}$$

At the distinguished real embedding this matrix is either hyperbolic or has $|\sigma_0(\varepsilon)| = 1$. In the latter case $\varepsilon = \pm 1$, since σ_0 is injective. A projectively nonidentity matrix of this form is then unipotent and hence parabolic. Its fixed point is a cusp. In the former case it is hyperbolic, fixes x, and its two eigenvalues $\varepsilon, \varepsilon^{-1}$ lie in K. This is the special alternative. Conversely, either alternative supplies a projectively nonidentity element of the stabilizer. □

This proposition removes one possible source of confusion. "Special" does not mean that an unrelated exceptional set has been added to the cusps. Once a nontrivial stabilizer of a K-rational boundary point has been found, its diagonal entry is automatically a unit. The two outcomes are exactly the two possible parts of a rational Borel: its unipotent direction gives a cusp, and its unit direction gives a hyperbolic period.

The same statement has a useful geometric interpretation. Under the three real embeddings, Γ_K is the Hilbert modular group acting on $\mathbb{H}^3$. The quotient

$$X_K = \Gamma_K \backslash \mathbb{H}^3$$

is the Hilbert modular threefold for K, and the modular embedding (7) realizes the triangle quotient as a curve in X_K. A point $x \in \mathbf{P}^1(K)$ specifies a rational boundary flag of the threefold. Its upper triangular stabilizer contains the unit lattice

$$\left\{ \begin{pmatrix} \varepsilon & 0 \\ 0 & \varepsilon^{-1} \end{pmatrix} : \varepsilon \in \mathcal{O}_K^\times \right\}.$$

The logarithms of the three absolute values of ε lie in the plane $t_0 + t_1 + t_2 = 0$ and form a rank-two lattice there. Thus a rational Hilbert-modular cusp has two compact unit-shape directions in addition to its ordinary height direction.

In this language, the conjecture says the following. Every rational boundary flag is either actually reached by a cusp of the Teichmüller curve, or the corresponding lift of that curve is periodic in a nonzero unit direction. The second possibility is precisely a special closed geodesic. This is why the rank-two unit ambiguity cannot be removed as a harmless normalization: motion in that ambiguity is one of the desired answers.

There is also an exact link with the remainders. Let e be the primitive standard cusp column and let $c_g = g^{-1}e$. If $\ell_x = (-q, p)$, then in the Borel frame (77)

$$H_x^{-1}c_g = \begin{pmatrix} a_g \\ r_g \end{pmatrix}, \qquad r_g = \ell_x(c_g) = \det((p, q)^t, c_g). \tag{81}$$

Thus a remainder is the lower Bruhat coordinate of the moving cusp vector relative to the rational Borel B_x. The equality $r_g = 0$ says that the moving cusp has entered B_x. Equality of two complete reduced states says that the quotient of the two prefixes lies in B_x, and hence gives the unit or unipotent stabilizer in Proposition 7.1.

The negative Hodge weight also has a transparent meaning on this Borel frame. The real character underlying $\mathcal{N} = \mathcal{L}_1\mathcal{L}_2\mathcal{L}_0^{-2}$ is

$$(t_0, t_1, t_2) \longmapsto t_1 + t_2 - 2t_0.$$

On a unit, the product formula turns this into $-3t_0$. Hence $\mathcal{N}$ is not invariant under the unit-frame lattice: it measures motion along that lattice. This is not a defect to be averaged away. It says that the contracting Hodge quantity lives on the Borel-frame torsor, whereas the ordinary projective height lives after quotienting by that torsor. A proof must show that a first norm record cannot escape in the height direction while continually changing its complete Borel frame. If that frame repeats, Proposition 7.1 gives a stabilizer; its diagonal entry then distinguishes the special alternative from the parabolic one.

7.1 The proof spine

The cleanest unconditional result in this paper is not the congruence calculation below. It is the integral Asai obstruction of Section 8. Here is the resulting proof spine, with the one missing implication stated without euphemism.

1. Proposition 7.1 turns the classification into the assertion that every rational line $[P]$ has a nonidentity stabilizer in the full group.
2. A rational line defines the nonzero integral Asai vector
$$Y(P) = \sigma_0(P) \otimes \sigma_1(P) \otimes \sigma_2(P).$$
If the ray ending at $[P]$ is not a cusp, it returns to a compact part of the triangle surface. Lattice discreteness at those returns forces the sum of its two outer Hodge growth rates to be at least 1; this is Proposition 8.1.
3. Proposition 8.3 strengthens the exponent statement, while Proposition 8.2 identifies the excess with $6h([gP]) + O(1)$. If the outer growth exceeds the central contraction by only a bounded amount along infinitely many compact returns, Northcott gives repetition. Thus a counterexample must escape through rational orbit points of unbounded height even when that height is sublinear in the return time. Proposition 8.4 makes the residual gap exact: among the first N compact returns, a nonperiodic orbit must have compensation at least $\log N - O(1)$.
4. The two outer Lyapunov exponents for the Hodge variation are $1/5$ and $3/5$. Their sum is $4/5$. A noncuspidal rational endpoint is therefore separated from the ordinary Hodge behavior by the definite gap $1/5$. This conclusion uses the complete Galois representation of the group and no choice of generators.
5. What remains is an exceptional-ray theorem: a trace-field ray whose outer growth reaches or exceeds 1 must be periodic. Proposition 8.3 sharpens the boundary case: when the exponent is exactly 1, bounded excess already closes the orbit, so only unbounded sublinear excess remains at that boundary. Periodicity gives a nonidentity stabilizer, and Proposition 7.1 then gives exactly the special hyperbolic alternative. Theorem 8.37 proves this conditional implication.

The first four assertions are proved. The fifth is the precise place where the proof stops. It is substantially narrower than the original classification: the candidate counterexamples have already been forced into one Hodge large-deviation stratum, and within its boundary they must escape every bounded integral-height set.

Three further unconditional statements describe the boundary of that stratum. Proposition 4.3 proves the doubled-central singular-value gap for every group prefix, with size proportional to thick-part displacement. Proposition 8.7 proves that sublinear state-height escape is exactly sublinear

all-place tracking of the rational Borel B_P; Proposition 8.8 makes this comparison two-sided and proves that the quotient height is a proper height of the complete rational state. Proposition 17.72 proves that if even one outer modular embedding has the rational Galois boundary value with finite angular derivative, then the endpoint is already a cusp. Thus the unresolved boundary case is not arbitrary slow growth: it is sublinear tracking of a rational Borel while both outer rational flags remain in the double-horoball-mismatch channel. Proposition 17.73 then identifies that channel with one explicit cone in the logarithmic unit lattice; closed special orbits give the rational rays in this cone described by Corollary 17.74.

The active mechanism for proving the fifth implication has two stages, both packaged in Conjecture 17.52. The inequality stage is a stopped Arakelov maximum principle for the integral Asai line. It must prevent the positive Galois component from overcoming the 1/5 Hodge gap at every first record. The equality stage must make a complete reduced state recur. Proposition 8.9 identifies this second demand as recurrence of the arithmetic Borel coset, and Proposition 17.73 describes only the diagonal, or Levi, coordinate after that recurrence. Thus the conjecture is an equality-case extension of Filip's integral-vector theorem, but its neutral conclusion is frame recurrence rather than convergence of a unit direction. This is one geometric route because the normalized Asai state already supplies the integral lattice, the reciprocal remainder divisor, and the correct equality locus.

The later analysis makes this route more concrete. The dynamically correct future–cusp frame and the height-reduced frame differ by one affine coordinate s_n, satisfying $s_{n+1} = d_n^2 s_n + d_n \beta_n$. Balancing turns this into the additive equation $z_{n+1} - z_n = \xi_n$, so the unit character itself is not the uncontrolled variable. Proposition 8.25 identifies that recurrence with the canonical extension

$$0 \longrightarrow \mathscr{L}^2 \longrightarrow \mathscr{L} \otimes V \longrightarrow \mathcal{O} \longrightarrow 0$$

on the projective flat bundle, and Corollary 8.26 identifies its logarithmic cotangent with the transverse Wronskian. Thus the integral rank-two object and its first jet are constructed.

There is now an exact description of the remaining boundary value. For an interval quotient $k_{m,n} = g_m^{-1} g_n$, its two coordinates in the future–cusp frame are

$$(a_{m,n}, b_{m,n}) = \widehat{H}_P^{-1} k_{m,n} P, \qquad b_{m,n} = \det(P, k_{m,n} P).$$

The second coordinate is simultaneously the old integral fixed-point form and, up to a unit, the two-frame Wronskian minor. The dynamic and Wronskian frames differ by the single rational shear $-a_{m,n}/b_{m,n}$. Schwarz–Pick accumulates the full Hodge defect on the actual geodesic; no path-shadowing assertion remains. Proposition 8.16 proves that the entire loss in passing to the rational Wronskian frame is

$$2 \log |\,\mathrm{CR}_0\,| - \log |\,\mathrm{CR}_1\,| - \log |\,\mathrm{CR}_2\,|,$$

where the two algebraic endpoint lines are exactly $[E]$ and $[k_{m,n}P]$. Proposition 8.19 shows that one sufficient scalar assertion is

$$2\alpha_0 - \alpha_1 - \alpha_2 \geq \eta_P \mathfrak{C}(m,n) - O_P(1).$$

That route asks central visual misalignment to dominate both outer backtracking defects. Proposition 8.20 proves that this comparison is not necessary: without any arithmetic hypothesis, at least one third of the defect appears either in the central Green drop or in the controlled two-future Hodge potential. Problem 17.17 is therefore the preferred remaining assertion. It asks a stopped

Hodge-selected slope to retain the larger coordinate while charging the common divisor $(\Delta_{m,n})$ only once. Proposition 17.20 proves that the ordinary static maximal slope cannot do this: it sees only the symmetric Green sum. Proposition 17.23 constructs the correct two-coordinate carrier and shows that its static global degree difference is zero, so the required strict term must be the stopped Hodge transport. This is the exact point at which the full pointed thin orbit, rather than an ambient primitive pair, must enter. Proposition 17.18 constructs the framed lattice explicitly as

$$0 \longrightarrow \mathcal{O}_K Q_m \longrightarrow \mathcal{O}_K Q_m + \mathcal{O}_K Q_n \longrightarrow (\Delta_{m,n}) \longrightarrow 0,$$

and Proposition 17.19 proves that its ordinary maximal slope is simultaneous approximation, by one $s \in K$, to the three moving Hodge splitting targets. The dynamic choice $s = -A_{m,n}/\Delta_{m,n}$ uses the full determinant denominator and is the actual-path line. The exact degree formula (345) then proves that this static slope is not the desired cyclic metric. Thus the algebraic objects and the precise failure of ordinary slope are both specified. Proposition 8.22 identifies the neutral recurrent affine splitting with the second eigenline of a hyperbolic stabilizer. Thus the equality alternatives are already the cusp and special alternatives of the theorem, not an unspecified degenerate set.

An earlier finite-quotient experiment is retained below as an audit of a failed route, not as a second proof strategy. Its required pointed inverse theorem already contains the same rational-Borel rigidity as step 5, so finite-level expansion does not advance the missing implication. In The active path treats the finite part as one Arakelov divisor and the archimedean part as the pointed thin-orbit cross ratio above.

7.2 Which structure of the thin group is relevant

The embedding of G_7 in the Hilbert modular group is genuinely thin. In fact its diagonal Galois representation

$$\rho_{\mathrm{Gal}} : G_7 \longrightarrow \mathrm{SL}_2(\mathbf{R})^3, \qquad g \longmapsto (\sigma_0(g), \sigma_1(g), \sigma_2(g)),$$

is Zariski dense.

Proposition 7.2 (Zariski closure of the Galois representation)**.** *The real Zariski closure of* $\rho_{\mathrm{Gal}}(G_7)$ *is* SL_2^3.

Proof. In each factor the image contains the nontrivial upper unipotent $\sigma_i(T)$ and its lower-unipotent conjugate by S. Its projection to each copy of SL_2 is therefore Zariski dense.

The identity component of the full closure is consequently a semisimple subdirect product of three copies of SL_2. If it were proper, the algebraic Goursat lemma would identify two factors through an automorphism of PSL_2. Such an automorphism preserves the squared trace. For the single group element ST, however, the three squared traces are

$$\lambda_0^2, \qquad \lambda_1^2, \qquad \lambda_2^2,$$

which are pairwise distinct. No two factors can be identified, so the subdirect product is the full product. □

Zariski density is only an algebraic statement. In the two outer real factors one can say more, and the stronger statement makes the precise meaning of thinness visible.

Proposition 7.3 (Outer adelic density of the thin group)**.** *After passing to projective groups, the joint outer image*

$$G_7 \longrightarrow \mathrm{PSL}_2(\mathbb{R}) \times \mathrm{PSL}_2(\mathbb{R}), \qquad g \longmapsto (\sigma_1(g), \sigma_2(g)),$$

is topologically dense. The same is true for every finite-index subgroup of G_7.

Consequently, at every ideal $\mathfrak{a}$ for which reduction is surjective, each prescribed class $\bar{g} \in \mathrm{PSL}_2(\mathcal{O}_K/\mathfrak{a})$ has lifts whose two outer images are dense in $\mathrm{PSL}_2(\mathbb{R})^2$. In particular, finite congruence conditions and open conditions in both outer real factors can be imposed simultaneously.

Proof. First consider either outer factor separately. The element

$$A_3 = \begin{pmatrix} \alpha & \alpha \\ \lambda & \alpha \end{pmatrix}, \qquad \alpha = \lambda^2 - 1,$$

has trace 2α. It is hyperbolic at σ_0, whereas

$$|2\sigma_1(\alpha)| < 2, \qquad |2\sigma_2(\alpha)| < 2.$$

Thus its two outer images are elliptic. They have infinite order: an outer image of finite projective order would give $\sigma_i(A_3^n) = \pm I$ for some n, and injectivity of the field embedding would give $A_3^n = \pm I$, contrary to its central hyperbolicity. Hence each outer projection of G_7 is nondiscrete. It is also nonelementary, since it contains the upper unipotent T and its lower-unipotent conjugate by S. The classification of closed subgroups of $\mathrm{PSL}_2(\mathbb{R})$ now makes the closure of each projection the full projective group.

Let H be the closure of the joint outer image. The powers of the same element A_3 lie in a compact product of two elliptic circles. They form an infinite set, so differences of two powers accumulate at the identity. Thus H is not discrete. Its two projections are surjective. The identity component H° is normal in H, so each of its projections is either trivial or the full simple group. If one is trivial, the presence of the other full factor and surjectivity of H force H to be the full product. If both are full, the topological Goursat lemma says that H° is either the full product or the graph of an automorphism. The graph case is impossible: an automorphism graph is self-normalizing in the product, so the normality of H° would put all of H in that graph; but an automorphism of $\mathrm{PSL}_2(\mathbb{R})$ preserves squared trace, while the squared traces of ST in the two outer factors are $\lambda_1^2 \neq \lambda_2^2$. Therefore H is the full product.

If $G' \subset G_7$ has finite index, the connected group $\mathrm{PSL}_2(\mathbb{R})^2$ is a finite union of translates of the closure of the outer image of G'. Such a closure has finite index and is therefore open; connectedness makes it the whole group. Apply this to the kernel of reduction modulo $\mathfrak{a}$. Strong approximation supplies one lift of $\bar{g}$, and multiplication by the kernel makes the outer images of that entire residue-class fibre dense. □

Remark 7.4 (What this does and does not prove). This proposition is a genuine statement about the complete thin group, not its generators. It says that there is no qualitative obstruction to balancing both outer Hodge directions inside any admissible congruence class. It does not bound the distinguished hyperbolic displacement of the required lift. The classification needs the quantitative form at the entropy scale $R = (1 + o(1)) \log \mathrm{Nm}\,\mathfrak{a}$, together with the Hodge window $a_1 + a_2 \leq (2/5 + o(1))R$. Thus the remaining assertion is an optimal quantitative weak-approximation theorem for this thin diagonal embedding; ordinary strong approximation records only its finite, qualitative shadow.

7.3 There is no larger discrete group to interpolate through

It is tempting to enlarge G_7 to a group which is still thin but has more arithmetic structure, prove the height inequality there, and then descend. For the heptagonal group this suggestion has a

precise and rather restrictive answer. There is a natural arithmetic comparison group, but there is no larger discrete comparison group.

Put $\Gamma = \mathrm{PSL}_2(\mathcal{O}_K)$, and always use the distinguished embedding σ_0 when discussing discreteness.

Proposition 7.5 (Dichotomy for intermediate groups)**.** *Let*

$$G_7 \subseteq H \subseteq \Gamma.$$

Then exactly one of the following holds.

1. *The image $\sigma_0(H)$ is discrete, in which case $H = G_7$.*
2. *The image $\sigma_0(H)$ is dense in $\mathrm{PSL}_2(\mathbb{R})$.*

Thus every proper intermediate group immediately loses the Fuchsian orbifold and the Rosen cross-section used in this paper.

Proof. Suppose first that $\sigma_0(H)$ is discrete. A discrete group containing the finite-covolume lattice G_7 contains it with finite index d. The orbifold Euler characteristic of the triangle orbifold of signature $(2, 7, \infty)$ is

$$\chi(G_7\backslash\mathbb{H}) = \frac{1}{2} + \frac{1}{7} - 1 = -\frac{5}{14}.$$

The overgroup still contains a parabolic, so its quotient is a noncompact orientable hyperbolic orbifold. Such an orbifold has $-\chi \geq 1/6$, with equality for signature $(2, 3, \infty)$. Consequently $5/(14d) \geq 1/6$, and hence $d \leq 2$. If $d = 2$, the putative quotient would have Euler characteristic $-5/28$. Write its signature as $(g; m_1, \ldots, m_r; s)$, where $s \geq 1$ is the number of cusps. The formula

$$-\chi = 2g - 2 + s + \sum_{j=1}^{r}(1 - 1/m_j)$$

reduces the only possible case to $g = 0, s = 1, r = 2$, and would require $1/m_1 + 1/m_2 = 23/28$. For $2 \leq m_1 \leq m_2$, this has no solution: $m_1 = 2$ would give $1/m_2 = 9/28$, while $m_1 \geq 3$ makes the sum at most $2/3$. Hence $d = 1$.

Now suppose that $\sigma_0(H)$ is nondiscrete and let L be its topological closure. The identity component L° is nontrivial and is normalized by G_7. Its Lie algebra is therefore invariant under the adjoint action of the Zariski-dense group G_7. Since $\mathfrak{sl}_2$ is simple, that Lie algebra is all of $\mathfrak{sl}_2(\mathbb{R})$. Thus $L = \mathrm{PSL}_2(\mathbb{R})$. □

The proposition also identifies the maximal-discrete-overgroup idea with no ambiguity. The group G_7 is nonarithmetic as a Fuchsian lattice, and the commensurator theorem makes its commensurator discrete [39]. Proposition 7.5 then gives

$$\mathrm{Comm}_{\mathrm{PSL}_2(\mathbb{R})}(G_7) = G_7.$$

Up to the central sign in SL_2, there is not even a finite discrete enlargement to exploit.

There is nevertheless a canonical arithmetic comparison group. Let $\widehat{\Gamma} = \mathrm{SL}_2(\widehat{\mathcal{O}}_K)$, let $\widehat{G}_7$ be the finite-adelic closure of a lift of G_7, and define its *congruence hull*

$$\Gamma_{\mathrm{cong}}(G_7) := \mathrm{SL}_2(\mathcal{O}_K) \cap \widehat{G}_7. \tag{82}$$

Calegari's exact closure calculation gives

$$[\mathrm{SL}_2(\mathcal{O}_K) : \Gamma_{\mathrm{cong}}(G_7)] = 144. \tag{83}$$

Indeed his integral model is obtained from the present one by conjugating with $\mathrm{diag}(\lambda, 1)$: the upper parameter λ becomes $\lambda^2 = 2 + 2\cos(2\pi/7)$, and the opposite parabolic parameter $-\lambda$ becomes -1. Since λ is a unit, this conjugation preserves $\mathrm{SL}_2(\mathcal{O}_K)$ and its congruence quotients. The two opposite parabolics generate the full triangle group: if $r = ST$, their product in one order is r^2; since r has order seven, this recovers r, and then S. Thus the two integral models have the same congruence closure. The index 144 is the $n = 7$ entry of [22, Theorem 1.6].

This hull clarifies exactly what congruence information can do. The closures of G_7 and $\Gamma_{\mathrm{cong}}(G_7)$ at all finite places are the same. On the other hand, every K-rational line P has a nonidentity unipotent stabilizer in $\Gamma_{\mathrm{cong}}(G_7)$. To see this, choose a nonzero nilpotent $N_P \in M_2(K)$ with kernel P. After clearing denominators, $I + aN_P \in \mathrm{SL}_2(\mathcal{O}_K)$ for all a in a nonzero ideal. Some positive power lies in the finite-index group $\Gamma_{\mathrm{cong}}(G_7)$. Thus the desired Borel incidence is automatic in the hull and difficult only in G_7, even though no finite congruence quotient distinguishes the two groups.

There is a stronger reason that this ambient stabilizer cannot be added as a harmless intermediate symmetry. Let U_P be the one-dimensional unipotent K-group fixing the line P, and put

$$\Lambda_P = U_P(K) \cap \Gamma_{\mathrm{cong}}(G_7), \qquad H_P = \langle G_7, \Lambda_P \rangle.$$

Proposition 7.6 (Full cusp saturation is already arithmetic)**.** *The group H_P has finite index in $\mathrm{SL}_2(\mathcal{O}_K)$. In particular, the most natural enlargement which adds the complete integral unipotent stabilizer of P is not an intermediate thin group.*

Proof. Use the three real embeddings to regard

$$\mathrm{SL}_2(\mathcal{O}_K) \subset \mathbf{G} := \mathrm{SL}_2(\mathbb{R})^3$$

as an arithmetic lattice. The group H_P, being a subgroup of that lattice, is discrete in $\mathbf{G}$. The $\mathbb{Q}$-Zariski closure of G_7 is all of $\mathrm{Res}_{K/\mathbb{Q}}\,\mathrm{SL}_2$. One quick way to check this in the present cubic case is to use Goursat's lemma. Every real projection is Zariski dense in SL_2, while a diagonal identification of two factors would make their squared trace characters agree. They do not: the order-seven element has trace $\sigma_i(\lambda)$, whose three squares are distinct. Thus H_P is Zariski dense in $\mathbf{G}$.

The group Λ_P is commensurable, after choosing a K-rational coordinate on U_P, with the additive group of a nonzero fractional ideal of K. Under the Minkowski embedding it is therefore a lattice in

$$U_P(\mathbb{R})^3 \cong \mathbb{R}^3.$$

It is irreducible in the precise sense used by Benoist–Miquel: its intersection with every proper normal subproduct of $\mathbf{G}$ is trivial. Indeed a nonzero element of a fractional ideal has a nonzero image under every field embedding. The arithmeticity theorem of Benoist–Miquel now applies: a discrete Zariski-dense subgroup of a real semisimple group of rank at least two which contains an irreducible lattice in a horospherical unipotent group is an arithmetic lattice [12]. Hence H_P is a lattice in $\mathbf{G}$. Since it is contained in the lattice $\mathrm{SL}_2(\mathcal{O}_K)$, its index there is finite. □

The qualification "complete" is essential. If one adjoins only one unipotent $u \in \Lambda_P$, the theorem does not apply: the cyclic group $\langle u \rangle$ has rank one, whereas Λ_P has additive rank three. The group $\langle G_7, u \rangle$ may therefore be a proper infinite-index subgroup or may already be arithmetic, depending on u; in either case its distinguished real projection is dense by Proposition 7.5. It has no canonical Fuchsian quotient and no inherited Rosen stopping rule. This exposes a sharp threshold: adding too

little supplies no global cusp lattice, while adding the full lattice supplies so much that arithmeticity is forced.

Even the most canonical rank-two additive correction crosses this threshold. At the standard cusp, the existing stabilizer contains $T(\lambda\mathbb{Z})$, and $\mathrm{Tr}_{K/\mathbb{Q}}(\lambda) = 1$. Consequently

$$\mathcal{O}_K = \lambda\mathbb{Z} \oplus \ker\big(\mathrm{Tr}_{K/\mathbb{Q}} : \mathcal{O}_K \to \mathbb{Z}\big).$$

Adjoining all translations in the trace-zero rank-two lattice therefore adjoins the complete $T(\mathcal{O}_K)$, so Proposition 7.6 applies. A codimension-one additive enlargement is thus already too large; only a single additional cyclic direction escapes this arithmeticity argument, although the canonical rational direction is settled explicitly below.

The most obvious unit enlargement shows why one cannot simply add the missing neutral direction. Lift G_7 to $\mathrm{SL}_2(\mathcal{O}_K)$ and put

$$D = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix}.$$

The minimal polynomial gives $\lambda^{-1} = -\lambda^2 + \lambda + 2 \in \mathcal{O}_K$, so D is integral. Also

$$D^n T(\lambda) D^{-n} = T(\lambda^{2n+1}).$$

Relative to the integral basis $1, \lambda, \lambda^2$, the three translation parameters

$$\lambda = (0, 1, 0), \qquad \lambda^3 = (-1, 2, 1), \qquad \lambda^5 = (-3, 5, 4)$$

form a $\mathbb{Z}$-basis: the determinant of the displayed coefficient matrix is 1. Hence $\langle G_7, D\rangle$ contains every upper elementary matrix $T(a)$, $a \in \mathcal{O}_K$, and conjugation by S gives every lower elementary matrix. Vaserstein's theorem says that $\mathrm{SL}_2(\mathcal{O}_K)$ is generated by elementary matrices when $\mathcal{O}_K^\times$ is infinite [61]. Therefore

$$\langle G_7, D\rangle = \mathrm{SL}_2(\mathcal{O}_K). \tag{84}$$

The unit which would neutralize the troublesome Galois weight also arithmeticizes the group completely.

In fact one need not adjoin a diagonal element. The smallest canonical additive enlargement already contains it.

Proposition 7.7 (One rational translation arithmeticizes G_7)**.** *With $T(a) = \left(\begin{smallmatrix} 1 & a \\ 0 & 1 \end{smallmatrix}\right)$, one has*

$$\langle G_7, T(1)\rangle = \mathrm{SL}_2(\mathcal{O}_K).$$

Proof. Exact multiplication in $\mathbb{Z}[\lambda]/(\lambda^3 - \lambda^2 - 2\lambda + 1)$ gives

$$\begin{aligned} &T(-1)ST(-\lambda)ST(1)T(-\lambda)S^{-1}T(-\lambda) \\ &\qquad \cdot T(-1)S^{-1}T(\lambda)T(-1)S^{-1}T(1) = \begin{pmatrix} \lambda & 0 \\ 0 & 2 + \lambda - \lambda^2 \end{pmatrix} = D. \end{aligned} \tag{85}$$

The last equality uses $2 + \lambda - \lambda^2 = \lambda^{-1}$. The left side belongs to $\langle G_7, T(1)\rangle$, since $S, T(\lambda) \in G_7$. Equation (84) now proves the result. □

There is nevertheless a genuinely proper additive enlargement. Its two visible translation parameters initially make it look one rank short of arithmeticity; the split order-seven axis creates the missing rank. Put

$$L_2 = 2\mathbb{Z} + \lambda\mathbb{Z} \subset \mathcal{O}_K, \qquad H_2 = \langle G_7, T(2)\rangle.$$

Proposition 7.8 (The apparent rank-two enlargement is arithmetic)**.** *The group H_2 is a proper finite-index subgroup*

$$G_7 \subsetneq H_2 \subsetneq \mathrm{SL}_2(\mathcal{O}_K).$$

Its distinguished real projection is dense, and it contains the rank-two translation group $T(L_2)$. In particular $H_2 \cap U(\mathcal{O}_K)$ has rank three and strictly contains $T(L_2)$.

Proof. The standard cusp stabilizer of G_7 is $T(\lambda\mathbb{Z})$, and $2/\lambda \notin \mathbb{Q}$, so $T(2) \notin G_7$. Reduction modulo 2 gives the other strict containment. The matrix $T(2)$ becomes the identity, so the image of H_2 is the same dihedral group D_7 as the image of G_7, whereas $\mathrm{SL}_2(\mathcal{O}_K/(2)) \cong \mathrm{SL}_2(\mathbb{F}_8)$ has order 504. Thus H_2 is proper. It is Zariski dense because it contains G_7.

The additive subgroup $2\mathbb{Z} + \lambda\mathbb{Z}$ is dense in the distinguished copy of $\mathbb{R}$, since $\lambda/2$ is irrational. The closure of the distinguished projection of H_2 therefore contains the full upper unipotent one-parameter group. Conjugation by S gives the opposite one, and these two groups generate $\mathrm{SL}_2(\mathbb{R})$. This proves density.

It remains to prove finite index. Put

$$\alpha = \lambda^2 - 1, \qquad v_C = (\alpha, 1)^t,$$

and use the nilpotent convention

$$N_v = \begin{pmatrix} -pq & p^2 \\ -q^2 & pq \end{pmatrix} \quad \text{for } v = (p, q)^t.$$

The two matrices

$$W = \begin{pmatrix} 1+\lambda^2 & \lambda \\ \lambda & 1 \end{pmatrix} \in G_7, \qquad h = -LSE^{-1}SL = \begin{pmatrix} 1+2\lambda & 2\alpha^2 \\ 2 & 1+2\lambda \end{pmatrix} \in H_2$$

both fix the line Kv_C. Exact multiplication gives

$$[W, h] = I + 2\lambda^3\alpha^4 N_{v_C}. \tag{86}$$

This is a nonidentity integral unipotent in H_2.

On v_C, the eigenvalue of W is

$$u = \lambda^2 + \lambda.$$

Covariance of N_v under SL_2 therefore gives

$$W^n[W, h]W^{-n} = I + 2\lambda^3\alpha^4 u^{2n} N_{v_C} \qquad (n \in \mathbb{Z}). \tag{87}$$

The cubic field K has no proper intermediate field, and $u^2 \notin \mathbb{Q}$; hence $\mathbb{Q}(u^2) = K$ and $1, u^2, u^4$ are a $\mathbb{Q}$-basis. In the integral basis $1, \lambda, \alpha$, their coefficient determinant is 29. The three parameters in (87) for $n = 0, 1, 2$ consequently span a full additive lattice in K. Thus H_2 contains an irreducible rank-three lattice in the horospherical group fixing Kv_C. Proposition 7.6, equivalently the Benoist–Miquel horospherical arithmeticity theorem, now makes H_2 finite index in $\mathrm{SL}_2(\mathcal{O}_K)$.

A finite-index subgroup meets the standard group $U(\mathcal{O}_K) \cong (\mathcal{O}_K, +)$ in finite index. Its intersection therefore has rank three and cannot equal $T(L_2)$, proving the last assertion. The commutator and parameter identity in (86) are checked independently by `intermediate_cross_form_search.py`. □

Proposition 7.9 (The four split divisors form an H_2-axis square)**.** *Put $L = T(\lambda)$, $E = T(2)$, and*

$$u = \lambda\alpha^2, \qquad v = \lambda^{-1}\alpha^4.$$

The four quadratic forms of Proposition 6.5 are, up to fixed nonzero scalars, the fixed-point forms of the following four elements of H_2:

$$h_{AC} = LSE^{-1}SL = \begin{pmatrix} -1-2\lambda & -2\alpha^2 \\ -2 & -1-2\lambda \end{pmatrix}, \qquad Q_{h_{AC}} = -2Q_{AC}, \tag{88}$$

$$h_{BD} = Sh_{AC}S = \begin{pmatrix} 1+2\lambda & -2 \\ -2\alpha^2 & 1+2\lambda \end{pmatrix}, \qquad Q_{h_{BD}} = -2\alpha^2 Q_{BD}, \tag{89}$$

$$h_{AD} = E^{-1}SE^{-1}L^3SE^{-1}S = \begin{pmatrix} -4+12\lambda & -3+6\lambda \\ 3-6\lambda & 2-3\lambda \end{pmatrix}, \qquad Q_{h_{AD}} = (3-6\lambda)Q_{AD}, \tag{90}$$

$$h_{BC} = ESEL^{-3}SES = \begin{pmatrix} 4-12\lambda & -3+6\lambda \\ 3-6\lambda & -2+3\lambda \end{pmatrix}, \qquad Q_{h_{BC}} = (3-6\lambda)Q_{BC}. \tag{91}$$

Their fixed-line pairs are respectively

$$(e_A, e_C), \quad (e_B, e_D), \quad (e_A, e_D), \quad (e_B, e_C).$$

The first two have eigenvalues $u^{\pm 1}$, and the last two have eigenvalues $v^{\pm 1}$, up to simultaneous signs. Since $\mathcal{O}_K^\times = \{\pm\lambda^m\alpha^n : m, n \in \mathbb{Z}\}$, the two exponent vectors

$$(1, 2), \qquad (-1, 4)$$

span an index-six sublattice of the full rank-two unit lattice.

Modulo four, the image of G_7 has order 1792, while the image of H_2 has order 3584. All four displayed elements lie in the nontrivial coset. In particular none belongs to G_7, and none belongs to the congruence-neutral group H_4.

Proof. Multiplying the displayed words in $\mathbb{Z}[\lambda]/(\lambda^3 - \lambda^2 - 2\lambda + 1)$ gives the four matrices. For

$$Q_M(X, Y) = cX^2 + (d-a)XY - bY^2$$

direct substitution gives (88)–(91). Thus the zero lines are the indicated pairs of A, B, C, D. On the eigenline $(\alpha, 1)^t$, the positive representative of h_{AC} has eigenvalue

$$1 + 2\lambda + 2\alpha = \lambda\alpha^2 = u.$$

On the eigenline $(-\alpha, 1)^t$, h_{AD} has eigenvalue

$$2 - 3\lambda - (3-6\lambda)\alpha = \lambda^{-1}\alpha^4 = v.$$

Determinant one gives the reciprocal eigenvalues; the other two matrices are the opposite edges of the same square. The determinant of $\left(\begin{smallmatrix} 1 & -1 \\ 2 & 4 \end{smallmatrix}\right)$ is six, which proves the unit-lattice index.

For the last assertion, reduce the exact matrices modulo four. The level-four calculation for G_7 gives an image of order 1792; adjoining $T(2)$ doubles it. In this index-two extension, a word has nontrivial coset exactly when it contains an odd number of $E^{\pm 1}$ letters. Each word in (88)–(91) has odd parity. Since H_4 has the same modulo-four image as G_7, none of the four elements can lie in H_4. The exact finite-ring enumeration and all four fixed-form identities are independently checked by `intermediate_cross_form_search.py`. □

Corollary 7.10 (The four walls are adjacent arithmetic cusp heights)**.** *For $X \in \{A, B, C, D\}$, let $e_X \in \mathcal{O}_K^2$ be the corresponding primitive zero line. There is a nonidentity parabolic*

$$k_X = I + t_X N_{e_X} \in H_2, \qquad t_X \in K^\times.$$

After fixing the signs of the primitive columns so that $X(P) = \det(P, e_X)$, every edge $XY \in \{AC, AD, BC, BD\}$ satisfies

$$|\,\mathrm{N}\, Q_{XY}(P)|^2 = \frac{|\,\mathrm{N}\det(P, k_X P)|\,|\,\mathrm{N}\det(P, k_Y P)|}{|\,\mathrm{N}(t_X t_Y)|}. \tag{92}$$

Thus the four-wall minimum is, up to four fixed constants, the minimum geometric mean of the cusp incidences at the endpoints of an edge of the square.

Proof. Finite index of H_2 implies that its intersection with the integral unipotent group fixing any K-rational line has finite index there; this gives k_X. The identity

$$\det(P, (I + tN_e)P) = t \det(P, e)^2$$

follows directly from the definition of N_e. Multiply it for the two endpoints and use $Q_{XY} = XY$. □

Proposition 7.11 (The wall square is an order-seven resonance)**.** *Keep $L = T(\lambda_q)$, $E = T(2)$, but do not yet specialize the Hecke parameter. The shortest wall word has the universal identity*

$$LSE^{-1}SL = \begin{pmatrix} -1 - 2\lambda_q & -2\lambda_q(\lambda_q + 1) \\ -2 & -1 - 2\lambda_q \end{pmatrix}, \qquad Q_{LSE^{-1}SL} = -2\big(X^2 - \lambda_q(\lambda_q + 1)Y^2\big). \tag{93}$$

At $q = 7$,

$$\lambda_7(\lambda_7 + 1) = (\lambda_7^2 - 1)^2 = \alpha^2,$$

so this intermediate-group wall is exactly the AC wall selected by the split adjacent-parabolic norm.

At $q = 11$, by contrast,

$$\mathrm{N}_{K_{11}/\mathbf{Q}}\big(\lambda_{11}(\lambda_{11} + 1)\big) = -1, \qquad \mathrm{N}_{K_{11}/\mathbf{Q}}(\lambda_{11}^2 + 4) = 5741. \tag{94}$$

The first element is therefore not a square in K_{11}. Moreover the two quadratic extensions obtained by adjoining its square root and by adjoining $\sqrt{\lambda_{11}^2 + 4}$ are distinct. Thus the same natural intermediate-group word and the adjacent Rosen torus already select different axes at $q = 11$.

Proof. Multiplying the five matrices gives (93) without using a polynomial relation for λ_q. For $q = 7$, reduction modulo $\lambda_7^3 - \lambda_7^2 - 2\lambda_7 + 1$ gives the displayed square. For $q = 11$, taking resultants with

$$x^5 - x^4 - 4x^3 + 3x^2 + 3x - 1$$

gives (94). A square has positive rational norm. Finally, if the two quadratic extensions were equal, then $\lambda_{11}(\lambda_{11} + 1)/(\lambda_{11}^2 + 4)$ would be a square in K_{11}, whereas its norm is $-1/5741$. This proves the last assertion. The exact symbolic and resultant checks are included in the accompanying intermediate-group script. □

What the intermediate group now contributes. The role of H_2 is no longer merely that it supplies a rank-two translation plane. It supplies the four exact hyperbolic walls selected by the split norm. For a Rosen prefix g_j, covariance gives

$$Q_{h_\mu}(g_j P) = Q_{g_j^{-1} h_\mu g_j}(P), \qquad \mu \in \{AC, AD, BC, BD\}. \tag{95}$$

Thus the all-time target is a nearest-wall estimate for four ordered conjugacy families in the one nontrivial level-four coset of H_2. It is not a prime-by-prime assertion.

This does not allow trace dominance to be applied to all of H_2: its distinguished real projection is dense. The new, narrower proof target is to establish the nearest-wall estimate only for the four families in (95), with g_j restricted to consecutive Rosen prefixes. The two wall holonomies u and v see both directions of the unit lattice, while the one mod-four coset records the finite divisor once. This is exactly the rank-two Arakelov/fusion structure that was absent from the earlier use of H_4.

Proposition 7.12 (One legal neutral move can eject a unit corridor)**.** *Put*

$$\epsilon_N = (\lambda\alpha)^N, \qquad q_N = \frac{\epsilon_N - 1}{\lambda}, \qquad p_N = 1 - \alpha^{-1} q_N.$$

Then (p_N, q_N) is primitive and

$$A(p_N, q_N) = \epsilon_N, \qquad D(p_N, q_N) = 1.$$

Hence both orientation minima have norm one. For every sufficiently large N, the state p_N/q_N lies in the Rosen interval and its next digit is $+1$. At the following state one orientation minimum still has norm one, but the other tends to infinity. More precisely,

$$\begin{aligned} B' &= -\alpha, \\ C' &= \alpha^{-1}\epsilon_N - 2\alpha, \\ D' &= 2(2\alpha - 2\lambda - 1)\epsilon_N + (2 + 2\lambda - 3\alpha), \end{aligned} \tag{96}$$

and

$$\min\{|\,\mathrm{N}\, C'|, |\,\mathrm{N}\, D'|\} \longrightarrow \infty.$$

Thus neither the two-minimum wall height nor its one-divisor quadratic reformulation is a one-step Lyapunov function, even on primitive admissible states and even at digit $+1$.

Proof. The two displayed incidence identities are immediate from the definitions. Any ideal dividing p_N and q_N divides both ϵ_N and 1, so the pair is primitive. At the distinguished embedding,

$$\frac{p_N}{q_N} \longrightarrow -\alpha^{-1}, \qquad -\frac{q_N}{\lambda p_N} \longrightarrow \frac{\alpha}{\lambda} = 1.2469\ldots.$$

The first limit is in the interior of the Rosen interval and the nearest integer to the second is one. This proves the digit assertion for all large N.

For a $+1$ move, Proposition 6.2 gives $B' = -\alpha D = -\alpha$. Direct substitution gives the other two expressions in (96). Now

$$|\sigma_0(\lambda\alpha)| > 1, \qquad |\sigma_i(\lambda\alpha)| < 1 \quad (i = 1, 2).$$

At the distinguished place both C' and D' therefore grow like a nonzero constant times $\sigma_0(\epsilon_N)$; at each outer place they converge to the displayed nonzero constant term. Multiplying the three real absolute values proves that both absolute field norms tend to infinity. □

Problem 7.13 (The remaining H_2 wall-shadowing inequality)**.** *Let*

$$\mathcal{H}_\square = \{h_{AC}, h_{AD}, h_{BC}, h_{BD}\}.$$

For every primitive $P \in \mathcal{O}_K^2$, prove that there is $C(P)$ such that every Rosen prefix g_j of the ray defined by P satisfies

$$\min_{h \in \mathcal{H}_\square} \left|\mathrm{N}\det\left(P, g_j^{-1} h g_j P\right)\right| \le C(P). \tag{97}$$

This is equivalent to the missing classification, but it is not by itself a reduction in difficulty. Indeed

$$\det(P, g_j^{-1} h g_j P) = Q_h(g_j P),$$

and Proposition 7.9 identifies the four right sides with fixed multiples of $Q_{AC}, Q_{AD}, Q_{BC}, Q_{BD}$. Thus (97) is equivalent, up to the fixed norm factors 8 and 27, to the all-time split bound in Theorem 6.6. Proving Problem 7.13 proves the classification. The withdrawn estimate (40) shows that this qualitative boundedness cannot be replaced by the proposed uniform $6h(x)$ budget.

Proposition 7.12 explains why the argument must be stopped and global. At a $+1$ digit the A, B pair is transported from the C, D pair by units, but the new C, D pair contains the transverse shear $2\lambda r_j$; at a -1 digit the roles reverse. One half of the square is controlled at a time. The missing theorem must fuse the two times while charging the transverse unit excursion to the initial Arakelov height. The thin-group stammering theorem handles the case in which one repeated block dominates that height. What remains is precisely the balanced many-block case.

The ejected states in Proposition 7.12 do not immediately make another arbitrary move. They follow alternating (1,-1) digits for a number of steps which tends to infinity with (N). Proposition 6.3 shows that this whole interval is height-neutral up to a fixed constant. Thus the large one-step wall value is not evidence for repeated creation of height: it is a change of coordinate inside one already-funded unit corridor. The stronger junction-height estimate (33) would have been sufficient, but the Pisot-unit family gives the same adverse trend. Any replacement has to use the relation between the two endpoint times and the closing word; the present calculations do not support a scalar height ceiling at the junction.

Proposition 7.8 changes the geometric meaning of this last problem. The four walls do not live in a new thin group: they generate, together with G_7, a finite-index arithmetic subgroup. Hence every K-rational line has a nontrivial unipotent stabilizer in H_2, and the four vertices of the wall square are ordinary H_2-cusps. This does not prove (97). Arithmetic reduction is allowed to insert arbitrary H_2-moves, whereas that inequality permits only the four fixed walls conjugated by the consecutive G_7-prefixes. Forgetting this order replaces the desired G_7-stabilizer by the automatic H_2-stabilizer. The residual theorem is therefore a relative shadowing statement: the Rosen path in the thin subgroup must encounter, with bounded Arakelov cost, the arithmetic cusp walls of H_2.

This square is a $q = 7$-specific proof structure, not the proposed prime-family criterion. At $q = 11$ the adjacent torus is nonsplit, so the four K-rational vertices and the H_2 wall square do not exist; the scalar Hodge defect is also negative. At $q = 5$, by contrast, the torus is nonsplit but the cusp classification holds. The unifying conjecture is therefore still the sign of the Hodge defect. The H_2 square is the extra rational structure which may make the final positive-defect argument accessible at $q = 7$.

For orientation, its finite-adelic closure can also be read off exactly from Calegari's calculation, although this is not the mechanism sought here. In his conjugate model, $T(2)$ becomes $T(2\lambda)$. Modulo four its new tangent vector contributes one of the two scalar directions missing from the triangle-group image in $M_0(\mathcal{O}_K/(2))$; the other remains missing. The mod-two image stays D_7, and the full level-four kernel was already present. Consequently the congruence closure of H_2 has index 72, halfway between the index 144 closure of G_7 and the full Hilbert modular group. This calculation proves properness and describes the finite-place closure, but does not by itself decide finite index. Equation (86) supplies the missing archimedean–algebraic input: the wall square and the special G_7-axis create a complete rank-three cusp lattice at e_C, so Benoist–Miquel proves that H_2 is finite index. In particular the apparently rank-two standard cusp module is not the full intersection:

$$T(2\mathbb{Z} + \lambda\mathbb{Z}) \subsetneq H_2 \cap U(\mathcal{O}_K), \qquad \operatorname{rank}_{\mathbb{Z}}(H_2 \cap U(\mathcal{O}_K)) = 3. \tag{98}$$

This is the cubic counterpart of arithmeticity criteria for groups generated by root unipotents [63], but here the third direction is not inserted by hand: it is forced by the commutator of a split wall with the order-seven special axis.

There is a second rank-two enlargement which is better adapted to the present height problem. It adds no new finite-place information at all. Put

$$L_4 = 4\mathbb{Z} + \lambda\mathbb{Z}, \qquad H_4 = \langle G_7, T(4)\rangle.$$

Proposition 7.14 (The congruence-neutral rank-two enlargement)**.** *One has strict inclusions*

$$G_7 \subsetneq H_4 \subseteq \Gamma_{\mathrm{cong}}(G_7) \subsetneq \mathrm{SL}_2(\mathcal{O}_K).$$

The finite-adelic closures of H_4 and G_7 are equal, the distinguished real projection of H_4 is dense, and $T(L_4) \subseteq H_4 \cap U(\mathcal{O}_K)$. Moreover the following are equivalent:

1. *H_4 has infinite index in $\mathrm{SL}_2(\mathcal{O}_K)$;*
2. *$H_4 \cap U(\mathcal{O}_K) = T(L_4)$.*

Thus H_4 is a thin intermediate group exactly when it creates no upper translation outside its displayed rank-two plane.

Proof. Calegari's level-four calculation says that the closure of G_7 contains the full principal kernel at level four. Hence $T(4) \in \Gamma_{\mathrm{cong}}(G_7)$, and the sandwich

$$G_7 \subset H_4 \subseteq \Gamma_{\mathrm{cong}}(G_7)$$

proves equality of the two finite-adelic closures. The right-hand hull has index (144), so it is proper. The standard cusp stabilizer of G_7 is $T(\lambda\mathbb{Z})$, and $4/\lambda \notin \mathbb{Q}$, so the left inclusion is strict. Density of the distinguished projection follows from Proposition 7.5.

It remains to identify the threshold. Write $a = a_0 + a_1\lambda + a_2\lambda^2 \in \mathcal{O}_K$ and suppose $T(a) \in H_4$. If $a_2 \neq 0$, the three parameters $4, \lambda, a$ are linearly independent over $\mathbb{Q}$. Their integral span is therefore a lattice in the full Minkowski horosphere, and Proposition 7.6 makes H_4 finite index.

Suppose instead that $a_2 = 0$. Combining $T(a)$ with $T(\lambda)$ isolates $T(a_0)$. If $4 \nmid a_0$, Bezout with $T(4)$ produces either $T(1)$ or $T(2)$. The first is impossible because Proposition 7.7 would give the full group, whereas H_4 lies in the proper congruence hull. The second is also impossible: the level-four tangent calculation quoted above shows that $T(2)$ supplies one of the two directions missing from the closure of G_7, and indeed enlarges its congruence closure from index 144 to index 72. Hence $4 \mid a_0$, and $a \in L_4$.

We have proved that either the upper-unipotent intersection is exactly $T(L_4)$, or it has rank three and H_4 has finite index. Conversely, a finite-index subgroup meets $U(\mathcal{O}_K) \cong \mathcal{O}_K$ in rank three, so it cannot have intersection $T(L_4)$. This proves the equivalence. □

The value of H_4 is not that its index has already been decided. Its value is that it is a nontrivial group enlargement which is completely invisible at every finite place. Therefore comparison with H_4 cannot introduce a second ideal divisor: any difference between a consecutive G_7-word and an H_4-normalization is necessarily archimedean or ordered.

There is an even sharper negative answer for an enlargement which acts on the Levi, rather than merely adding translations.

Proposition 7.15 (Every nontrivial unit enlargement is arithmetic)**.** *For every nonzero integer m, the group*

$$\langle G_7, D^m\rangle, \qquad D = \mathrm{diag}(\lambda, \lambda^{-1}),$$

has finite index in $\mathrm{SL}_2(\mathcal{O}_K)$.

Proof. Conjugating the existing translation $T(\lambda)$ by $1, D^m, D^{2m}$ gives the three translation parameters

$$\lambda, \qquad \lambda^{2m+1}, \qquad \lambda^{4m+1}.$$

They are linearly independent over $\mathbb{Q}$. Indeed $\lambda^{2m} \notin \mathbb{Q}$: one real conjugate of λ has absolute value greater than one and another has absolute value less than one, so their nonzero even powers cannot all agree. Since K has prime degree three, $\mathbb{Q}(\lambda^{2m}) = K$, and $1, \lambda^{2m}, \lambda^{4m}$ is a basis. The three displayed parameters consequently span a finite-index additive subgroup of $\mathcal{O}_K$. They give an irreducible lattice in the full Minkowski horosphere, so Benoist–Miquel again implies that the generated group has finite index. □

These two propositions explain which intermediate structure can and cannot address the stopped Arakelov inequality. Left multiplication by $T(\ell)$, $\ell \in L_2$, changes a first column $(A, \Delta)^t$ to $(A + \ell\Delta, \Delta)^t$. It therefore changes the additive splitting $-A/\Delta$ by an integral gauge but leaves $[\bar{A}] \in \mathscr{R}_{(\Delta)}$ unchanged. Thus H_2 acts in the fibre of the Arakelov lift, not on its finite ray–Levi class. Proposition 7.9 now shows that this fibre action is nevertheless exactly the useful one: its four distinguished hyperbolic walls are the four quadratic divisors in the sufficient split estimate. In contrast, D^m acts on the Levi/unit coordinate, but Proposition 7.15 shows that every such enlargement is already arithmetic. The rank-two cusp plane is therefore the correct auxiliary normalization for the split characters, but it does not by itself control the unchanged ray–Levi class or remove the rank-two unit obstruction in Problem 17.17.

Proposition 9.2 gives a second, and more useful, meaning of an intermediate object. Enlarging the permitted continued-fraction quotients from $\mathbb{Z}$ to a rank-two module M, with $\mathbb{Z} \subset M \subset \mathcal{O}_K$, is an intermediate *reduction algorithm*. The choice

$$M_2 = \mathbb{Z} + 2\lambda^{-1}\mathbb{Z}$$

corresponds exactly to the translation module of H_2, because $\lambda M_2 = \lambda\mathbb{Z} + 2\mathbb{Z}$. But forgetting when the extra quotient is used and generating a group from all such moves loses the consecutive Rosen constraint. At the other extreme, allowing every $q \in \mathcal{O}_K$ gives the factor-$1/7$ norm-Euclidean algorithm and, at group level, the full cusp lattice which forces arithmeticity. The potentially exploitable intermediate structure is therefore the *labelled correction process* $\kappa_j = q_j + n_j$, not merely the overgroup generated by its possible translations.

The congruence-neutral group H_4 gives a sharper version of this reduction. Its quotient module is

$$M_4 = \lambda^{-1}L_4 = \mathbb{Z} + 4\lambda^{-1}\mathbb{Z} = \mathbb{Z} + 4(\lambda - \lambda^2)\mathbb{Z}. \tag{99}$$

The last equality uses $\lambda^{-1} = 2 + \lambda - \lambda^2$. Hence

$$\mathcal{O}_K/M_4 \cong \mathbb{Z} \oplus \mathbb{Z}/4\mathbb{Z}. \tag{100}$$

More explicitly, if $\kappa = a + b\lambda + c\lambda^2$, define

$$d_4(\kappa) = b + c = \mathrm{Tr}_{K/\mathbb{Q}}\left(\frac{\lambda^2 + \lambda - 2}{7}\,\kappa\right). \tag{101}$$

Then the quotient map may be written

$$[\kappa]_{M_4} \longmapsto \big(d_4(\kappa),\ b \bmod 4\big), \tag{102}$$

and every class has a representative

$$d\lambda + t(\lambda - \lambda^2), \qquad d \in \mathbb{Z}, \quad t \in \{0, 1, 2, 3\}. \tag{103}$$

Indeed M_4 is generated, modulo $\mathbb{Z}$, by the vector $(4,-4)$ in the coefficient basis (λ,λ^2), whose Smith normal form is $\mathbb{Z}\oplus\mathbb{Z}/4$. The trace formula in (101) follows from

$$\mathrm{Tr}(1)=3,\quad \mathrm{Tr}(\lambda)=1,\quad \mathrm{Tr}(\lambda^2)=5,\quad \mathrm{Tr}(\lambda^3)=4,\quad \mathrm{Tr}(\lambda^4)=13.$$

Thus H_4-normalization removes one entire free direction from the rank-two correction lattice. The only unbounded residual coordinate is the integer $d_4(\kappa)$; the other coordinate has four states. Since $d_4(n)=0$ for $n\in\mathbb{Z}$, one also has

$$d_4(\kappa_j)=d_4(q_j). \tag{104}$$

The growing transverse correction is therefore a function of the ambient Euclidean quotient alone, while the Rosen digit records which consecutive thin move was actually taken.

The functional d_4 is not an arbitrary coordinate chosen to make the quotient rank one. It is the primitive trace coordinate supplied by the codifferent of K.

Proposition 7.16 (The free H_4-coordinate is the codifferent coordinate)**.** *Let $f(X)=X^3-X^2-2X+1$ be the minimal polynomial of λ, and put*

$$\vartheta=\frac{\lambda^2+\lambda-2}{7}=\frac{\lambda}{f'(\lambda)}. \tag{105}$$

Then

$$\mathfrak{D}_K^{-1}=\vartheta\mathcal{O}_K,\qquad \mathrm{Tr}(\vartheta)=0,\qquad \mathrm{N}(\vartheta)=\frac{1}{49}. \tag{106}$$

If $\kappa=a+b\lambda+c\lambda^2$, then

$$d_4(\kappa)=\mathrm{Tr}_{K/\mathbf{Q}}(\vartheta\kappa)=b+c. \tag{107}$$

Moreover, the three conjugates of ϑ, in the order $\sigma_0,\sigma_1,\sigma_2$, have signs

$$(+,-,-). \tag{108}$$

Finally,

$$\ker(d_4:\mathcal{O}_K\to\mathbf{Z})=\mathbf{Z}+\mathbf{Z}(\lambda-\lambda^2),\qquad M_4=\mathbf{Z}+4\mathbf{Z}(\lambda-\lambda^2), \tag{109}$$

so M_4 is the level-four sublattice of the canonical trace hyperplane and $\ker(d_4)/M_4\cong\mathbf{Z}/4\mathbf{Z}$.

Proof. Reduction modulo $f(\lambda)=0$ gives

$$(\lambda^2+\lambda-2)f'(\lambda)=7\lambda.$$

Since $\mathcal{O}_K=\mathbf{Z}[\lambda]$, its codifferent is $f'(\lambda)^{-1}\mathcal{O}_K$. The element λ is a unit, so (105) generates the same fractional ideal. The Newton sums

$$\mathrm{Tr}(1)=3,\quad \mathrm{Tr}(\lambda)=1,\quad \mathrm{Tr}(\lambda^2)=5,\quad \mathrm{Tr}(\lambda^3)=4,\quad \mathrm{Tr}(\lambda^4)=13$$

give

$$\mathrm{Tr}(\vartheta)=0,\qquad \mathrm{Tr}(\vartheta\lambda)=\mathrm{Tr}(\vartheta\lambda^2)=1,$$

which proves (107). The norm of the codifferent is $1/\operatorname{disc}(K)=1/49$, proving the norm assertion. The factorization $\lambda_i^2+\lambda_i-2=(\lambda_i+2)(\lambda_i-1)$, together with $\lambda_0>1$, $0<\lambda_1<1$, and $-2<\lambda_2<-1$, proves (108). The two lattice identities in (109) follow respectively from $b+c=0$ and $\lambda^{-1}=2+\lambda-\lambda^2$. □

The sign pattern in (108) is the same central-versus-two-outer pattern as the heptagonal Hodge defect. This is not an equality between a linear trace functional and logarithmic Hodge weights. Its precise content is that the rank-one quotient selected by the congruence-neutral enlargement is canonically oriented in the same direction as the Hodge contraction. The four-state residual label is exactly the index of M_4 in that canonical hyperplane.

The direction removed from the transverse quotient has not vanished; it has become a separate one-dimensional fibre. There is an exact sequence

$$0 \longrightarrow M_4/\mathbb{Z} \cong \mathbb{Z} \longrightarrow \mathcal{O}_K/\mathbb{Z} \longrightarrow \mathcal{O}_K/M_4 \cong \mathbb{Z} \oplus \mathbb{Z}/4 \longrightarrow 0. \tag{110}$$

The two free directions now have different geometric meanings. To see this without choosing an abstract splitting, take $t \in \{0,1,2,3\}$ with $t \equiv -c \pmod 4$, and put

$$k = -\frac{c+t}{4}, \qquad n = a - 8k, \qquad d = b + c.$$

Then the identity

$$\kappa = n + 4k\lambda^{-1} + d\lambda + t(\lambda - \lambda^2) \tag{111}$$

gives

$$T(\kappa\lambda) = T(n\lambda)\, T(4k)\, T\big(\lambda[d\lambda + t(\lambda - \lambda^2)]\big)\,. \tag{112}$$

The first factor is already in G_7. The second is the single new H_4/G_7 cusp direction and, importantly, has the same real parameter $4k$ at all three embeddings. The last factor is transverse to H_4: apart from the four-state label t, its only unbounded parameter is $d = d_4(\kappa)$. Thus the original rank-two fusion has become two rank-one questions with geometrically distinct roles: a diagonal rational cusp displacement (k), and a Galois-transverse displacement (d).

This algebraic simplification has an exact Hodge interpretation. For a compact return set $\mathcal{K} \Subset \mathbb{H}$, put

$$D_z^{(4)}([\kappa]) = \inf_{m \in M_4} \sum_{i=0}^{2} d_{\mathbb{H}}(F_i(z), T(\sigma_i(\lambda(\kappa - m)))F_i(z))\,.$$

There are constants $c_{\mathcal{K}}, C_{\mathcal{K}} > 0$ such that

$$c_{\mathcal{K}} \log(1 + |d_4(\kappa)|) - C_{\mathcal{K}} \le D_z^{(4)}([\kappa]) \le C_{\mathcal{K}} \log(1 + |d_4(\kappa)|) + C_{\mathcal{K}}. \tag{113}$$

For the upper bound use the four normal-form representatives in (103). For the lower bound, the trace functional in (101) is bounded by a fixed multiple of the largest Minkowski coordinate, while $d_{\mathbb{H}}(w, T(x)w) = 2\,\mathrm{arsinh}(|x|/(2\,\mathrm{Im}\, w))$. Compactness of $\mathcal{K}$ makes the constants uniform.

Equation (113) is a genuine reduction of the missing rank-two correction to a rank-one transverse displacement. It is not yet the stopped inequality: replacing the labelled path by an arbitrary word in H_4 would lose consecutiveness and, by density of the distinguished projection, create uncontrolled shortcuts. The object to study is the original Rosen shift equipped with the scalar cocycle $d_4(\kappa_j)$ and its four-state torsion label.

Proposition 7.17 (The Euclidean extremal line is a square root of the codifferent)**.** *Let*

$$\mathfrak{p} = (\lambda + 2), \qquad y = \frac{f'(\lambda)}{7}, \qquad f(X) = X^3 - X^2 - 2X + 1.$$

Then

$$(7) = \mathfrak{p}^3, \qquad \mathfrak{D}_K = \mathfrak{p}^2, \qquad (y) = \mathfrak{p}^{-1}, \qquad \mathfrak{D}_K^{-1} = \mathfrak{p}^{-2}. \tag{114}$$

At the level of generators one has the exact identity

$$\left(\frac{f'(\lambda)}{7}\right)^2 = (\lambda-1)\frac{\lambda}{f'(\lambda)} = (\lambda-1)\vartheta. \tag{115}$$

The quotient $\mathfrak{p}^{-1}/\mathcal{O}_K$ has seven elements. Its six nonzero classes are

$$\lambda^j y + \mathcal{O}_K, \qquad 0 \le j \le 5,$$

and every one has inhomogeneous Euclidean minimum exactly $1/7$. Consequently these six classes attain the global Euclidean minimum $M(K) = 1/7$.

Proof. Modulo 7,

$$f(X) \equiv (X+2)^3.$$

Also $|\,\mathrm{N}(\lambda+2)| = 7$, so 7 is totally ramified and $(7) = \mathfrak{p}^3$. Since $\mathcal{O}_K = \mathbf{Z}[\lambda]$, the different is $(f'(\lambda))$; its norm is the discriminant 49, hence it is $\mathfrak{p}^2$. Dividing the two ideal identities proves (114). Direct reduction by $f(\lambda) = 0$ gives

$$f'(\lambda)^3 = 49\lambda(\lambda-1),$$

which is (115).

The quotient $\mathfrak{p}^{-1}/\mathcal{O}_K$ is the one-dimensional residue space over $\mathcal{O}_K/\mathfrak{p} \simeq \mathbb{F}_7$. In that residue field $\lambda \equiv -2$, which has order six in $\mathbb{F}_7^\times$; therefore the six displayed classes are exactly the nonzero ones. If ξ belongs to one of them, then every lift $\xi - q$, $q \in \mathcal{O}_K$, has $\mathfrak{p}$-valuation -1. Hence

$$|\,\mathrm{N}(\xi-q)| \in \left\{\frac{m}{7} : m \in \mathbf{Z}_{\ge 1},\ 7 \nmid m\right\},$$

and in particular it is at least $1/7$. The lift $\lambda^j y$ has absolute norm $1/7$, so equality holds. The global upper bound $M(K) = 1/7$ from Proposition 9.2 proves the final assertion. □

For a prime $p = 2g+1$, the same ideal calculation in $\mathbf{Q}(2\cos(\pi/p))$ gives

$$(p) = \mathfrak{p}^g, \qquad \mathfrak{D} = \mathfrak{p}^{g-1}, \qquad \left(\frac{f_p'(\lambda_p)}{p}\right) = \mathfrak{p}^{-1}, \qquad \left(\frac{f_p'(\lambda_p)}{p}\right)^{g-1} = \mathfrak{D}^{-1}$$

as fractional ideals. The square at $p = 7$ becomes a fourth power at $p = 11$. This is the finite-place form of the passage from a rank-two to a rank-four pairwise Hodge carrier.

For $p = 11$, this local picture also exhausts the Euclidean constant. Put $y_{11} = f_{11}'(\lambda_{11})/11$. Since $\lambda_{11} \equiv -2 \pmod{\mathfrak{p}}$ has order five in $\mathbb{F}_{11}^\times$, the ten elements

$$\pm\lambda_{11}^j y_{11} + \mathcal{O}_{K_{11}}, \qquad 0 \le j \le 4,$$

are all the nonzero classes in $\mathfrak{p}^{-1}/\mathcal{O}_{K_{11}}$. The same valuation argument as in the proof above shows that each class has inhomogeneous minimum exactly $1/11$: every lift has norm at least $1/11$, and the displayed lift has norm $1/11$. The known equality $M(K_{11}) = 1/11$ [13] makes them global extremal classes.

> **What the $p = 11$ comparison now rules out.** Both fields have the sharp ramified Euclidean minimum $1/p$, and in both fields its nonzero ramified classes are generated by unit multiples of $f_p'(\lambda_p)/p$. Therefore the contrast between the observed $p = 7$ and $p = 11$ Rosen dynamics cannot be attributed to the failure of norm-Euclidean division in degree five. The first proved structural difference is archimedean: the scalar Hodge defect changes sign and the codifferent normal loses its one-versus-all-outer orientation.

Proposition 7.18 (The three congruence-neutral walls form the Eisenstein A_2 complex)**.** *For $\kappa = a + b\lambda + c\lambda^2$, put*

$$\ell_0(\kappa) = b + c, \qquad \ell_1(\kappa) = b, \qquad \ell_2(\kappa) = c,$$

and define

$$\begin{array}{lll} v_0 = \lambda - \lambda^2, & M_0 = \mathbf{Z} + 4\mathbf{Z}v_0, & H_0 = \langle G_7, T(4\lambda v_0)\rangle, \\ v_1 = \lambda^2, & M_1 = \mathbf{Z} + 4\mathbf{Z}v_1, & H_1 = \langle G_7, T(4\lambda v_1)\rangle, \\ v_2 = \lambda, & M_2 = \mathbf{Z} + 4\mathbf{Z}v_2, & H_2 = \langle G_7, T(4\lambda v_2)\rangle. \end{array}$$

Then M_i is the level-four sublattice of $\ker \ell_i$*, and*

$$\mathcal{O}_K/M_i \simeq \mathbf{Z} \oplus \mathbf{Z}/4\mathbf{Z}. \tag{116}$$

The three modules and functionals are cyclically permuted, up to sign, by $\mathrm{Gal}(K/\mathbf{Q})$*. Moreover*

$$H_0 = H_4, \qquad G_7 \subsetneq H_i \subseteq \Gamma_{\mathrm{cong}}(G_7) \quad (0 \le i \le 2),$$

and every H_i has the same finite-adelic closure as G_7.

For z in a fixed compact return set, let

$$D_z^{(i)}([\kappa]) = \inf_{m \in M_i} \sum_{j=0}^{2} d_{\mathbb{H}}(F_j(z), T(\sigma_j(\lambda(\kappa - m)))F_j(z)).$$

Uniformly in κ,

$$D_z^{(i)}([\kappa]) \asymp_{\mathcal{K}} \log(1 + |\ell_i(\kappa)|) + O_{\mathcal{K}}(1). \tag{117}$$

Finally, the Eisenstein correction norm is exactly

$$Q(b,c) = b^2 + bc + c^2 = \frac{\ell_0(\kappa)^2 + \ell_1(\kappa)^2 + \ell_2(\kappa)^2}{2}. \tag{118}$$

Consequently

$$\log(1 + Q(b,c)) = 2 \max_{0 \le i \le 2} \log(1 + |\ell_i(\kappa)|) + O(1). \tag{119}$$

Proof. Modulo $\mathbf{Z}$, the three vectors v_i have coefficient vectors $(1,-1), (0,1), (1,0)$ in the basis $(\bar{\lambda}, \overline{\lambda^2})$. Their primitive annihilators are respectively $b + c, b, c$. Smith normal form proves (116).

For the generator τ with $\tau(\lambda) = \lambda^2 - \lambda - 1$, the coefficient action modulo $\mathbf{Z}$ is

$$(b,c) \longmapsto (-b - c, b).$$

It sends the three unoriented lines $\mathbf{Z}v_0, \mathbf{Z}v_1, \mathbf{Z}v_2$ cyclically to one another, proving the Galois assertion.

The identity $\lambda(\lambda - \lambda^2) = 1 - 2\lambda$ shows that $H_0 = \langle G_7, T(4)\rangle = H_4$. Every new translation parameter $4\lambda v_i$ belongs to $4\mathcal{O}_K$. Calegari's level-four calculation therefore places it in the congruence closure of G_7, so adjoining it does not change the finite-adelic closure. It is not an integral multiple of λ, whereas $G_7 \cap U(\mathcal{O}_K) = T(\lambda\mathbf{Z})$; hence every left inclusion is strict. The right inclusion is proper because the congruence closure of G_7 is proper.

The proof of (117) is the proof of (113) applied to M_i: the four torsion representatives give the upper bound, and the primitive functional ℓ_i gives the lower bound by finite-dimensional norm equivalence. Equation (118) is immediate. If $M = \max(|b|, |c|, |b + c|)$, then

$$\frac{1}{2}M^2 \le Q(b,c) \le \frac{3}{2}M^2.$$

This proves (119). □

This proposition identifies the exact geometry behind the adaptive maximum in Proposition 8.29. Both the archimedean Hodge problem and the finite correction lattice have the three walls of the A_2 Coxeter complex. The H_4 decomposition chooses the wall $\ell_0 = 0$: $d = \ell_0(\kappa)$ is normal to it, while the unbounded fibre k moves along it. When d is small and k is large, one of the other two wall coordinates b, c is large. Thus the two terms in χ_4 are not unrelated scalar errors; they are the normal and tangential pieces of one Galois-symmetric A_2 wall distance.

There is now a completely scalar cumulative statement which would finish the proof. Use the exact decomposition (111) and define

$$\begin{aligned}\chi_4(\kappa) = {} & \log(1+|d|) + \log(1+|k|) + \mathbf{1}_{t\neq 0}\log 2 \\ & + \mathbf{1}_{d=k=t=0}\log(1+|n|).\end{aligned} \tag{120}$$

The d-term is transverse to H_4, the k-term is the rational cusp fibre H_4/G_7, and t is the four-state label. The last term is used only when the ambient correction already lies on the original integer Rosen line; it is then paid by the central Green channel.

Conjecture 7.19 (Canonical codifferent correction budget)**.** *At every strict norm record choose an ambient quotient $q_j \in \mathcal{O}_K$ satisfying the factor-$1/7$ norm reduction, and put $\kappa_j = q_j + n_j$. The choices can be made so that, for a constant C_7 independent of the initial primitive state,*

$$\sum_{j\in\mathcal{R}} \chi_4(\kappa_j) \le 6h([r_0 : r_{-1}]) + C_7. \tag{121}$$

Here $\mathcal{R}$ is the set of strict norm-record transitions.

Proposition 7.20 (The codifferent correction budget implies the full classification)**.** *Conjecture 7.19 implies Conjecture 1.1.*

Proof. At a strict norm record the ambient correction is nonzero. Its exact decomposition shows that either $d \neq 0$, $k \neq 0$, $t \neq 0$, or all three vanish and $n \neq 0$. Hence $\chi_4(\kappa_j) \ge \log 2$ at every record. The finite right side of (121) permits only finitely many records. All remainder norms are therefore bounded. Proposition 9.10 then makes the Rosen orbit terminate or repeat, and Proposition 9.9 identifies the cusp and special-hyperbolic alternatives. □

This target is more concrete than an unspecified rank-two height. It asks for two ordered rank-one payments. The transverse payment is measured by the canonical codifferent coordinate and the Hodge distance (113); the rational fibre payment is an ordinary central cusp displacement. What is not yet proved is that these payments telescope along the labelled consecutive Rosen path. Treating the same translations as arbitrary generators of H_4 would again lose precisely that order.

The exact program `rosen_euclidean_correction_tests.py` tests (121). In three random batches, with coefficient bounds $10^2, 10^4, 10^6$, it followed 900 primitive states for 75,121 exact steps and found 1,149 record corrections. Every displayed correction satisfies the factor-$1/7$ inequality exactly; none was missed in the searched coefficient box. The largest observed value of $\sum \chi_4 - 6h([r_0 : r_{-1}])$ was 22.988. A separate family of 781 unit and unit-perturbed states gave 652 records and maximum 16.308. The test therefore rules out the version with zero additive constant and supports a bounded-error version. It is not a proof: the search for a Euclidean correction is finite and the orbit samples do not cover arbitrary height or arbitrarily many records.

The charge is not the logarithm of an ordinary continuant. For example, the exact Rosen orbit of $x = 1 - 3\lambda - 3\lambda^2$ has successive strict record corrections

$$\kappa_1 = -7 + 2\lambda + \lambda^2, \qquad \kappa_2 = 10 - 3\lambda^2,$$

for which

$$(d,k,t,n) = (3,-1,3,1), \qquad (-3,0,3,10).$$

Both corresponding ambient reductions have zero remainder, yet the codifferent coordinate reverses sign. Hence the sum in (121) cannot be bounded by a commutative net displacement or the product of positive partial quotients. The intervening G_7-conjugations must remain in the proof. A relative-cocycle or Bass–Serre length could retain that order; the scalar projection alone cannot.

Thus the very natural candidate whose translation module is $\mathbb{Z} + \lambda\mathbb{Z}$ is not thin. The identity (85) was found by a bidirectional exact word search and then verified independently by polynomial reduction; the displayed identity, rather than the search, is the proof.

The useful object is consequently the *pair*

$$G_7 \subset \Gamma_{\mathrm{cong}}(G_7),$$

not an intermediate quotient surface. It supplies an integral Borel witness for every rational line and proves that the missing information is invisible at every finite place. What it does not supply is a word in G_7: a stabilizer found after inserting an arbitrary hull element is only the automatic ambient stabilizer above. The descent must therefore measure, by the stopped Hodge metric, the failure of a consecutive G_7-orbit to realize its finite-adelic closure. This is closely parallel in architecture to arithmetic-holonomy proofs of unbounded denominators, where analytic continuation distinguishes a group from its congruence closure [59]; here the infinite-index and ordered-orbit features require the new rank-two inequality.

There is a useful exact check on this proposed comparison. It shows that the congruence hull supplies an ambient target, but not an additional finite divisor. Retain the cusp-normalized frame $\widehat{H}_P = (P, E)$, where $E = C/\delta_P$ and $\delta_P = \det(P, C)$. For $k \in \mathrm{SL}_2(\mathcal{O}_K)$, write

$$\widehat{H}_P^{-1} k \widehat{H}_P = \begin{pmatrix} a & c \\ b & d \end{pmatrix}. \tag{122}$$

Thus

$$a = \frac{\det(kP, C)}{\delta_P}, \qquad b = \det(P, kP), \qquad ad - bc = 1. \tag{123}$$

Proposition 7.21 (The ambient Borel defect has arithmetic codimension one)**.** *Put*

$$\alpha(k) = \det(kP, C) = \delta_P a, \qquad \beta(k) = \det(P, kP) = b.$$

Then

1. *$k \in B_P(K)$ if and only if $\beta(k) = 0$. Thus congruence to the full Borel is measured by the single lower-left minor $\beta(k)$.*
2. *If a prime ideal divides both $\alpha(k)$ and $\beta(k)$, it divides δ_P, and more precisely*

$$\min\{v_{\mathfrak{p}}(\alpha(k)), v_{\mathfrak{p}}(\beta(k))\} \le v_{\mathfrak{p}}(\delta_P). \tag{124}$$

   Hence the common ideal of the two boundary coordinates is bounded by the fixed ideal (δ_P), independently of k.
3. *For every ideal $\mathfrak{q}$ coprime to δ_P, reduction of k lies in the Borel stabilizing $[P]$ modulo $\mathfrak{q}$ exactly when $\beta(k) \equiv 0 \pmod{\mathfrak{q}}$. The other equality coordinate $\alpha(k) = 0$, which sends $[P]$ to the cusp $[C]$, is a different boundary incidence; it is not a second equation defining the same Borel.*

   Consequently no argument using only congruence proximity to $B_P \cap \Gamma_{\mathrm{cong}}(G_7)$ can contribute a second growing finite-place term to the stopped Arakelov degree.

Proof. The first column of (122) is $(a,b)^t$. It represents kP in the basis (P,E), so its second coordinate vanishes precisely when kP is proportional to P. This proves the first assertion.

The two determinant identities give

$$\delta_P kP = \alpha(k)P + \beta(k)C. \tag{125}$$

Fix a prime $\mathfrak{p}$, put $e = v_{\mathfrak{p}}(\delta_P)$, and work over the corresponding discrete valuation ring. Both P and kP are primitive. If both coefficients on the right of (125) were divisible by $\mathfrak{p}^{e+1}$, then $\delta_P kP$ would be divisible by $\mathfrak{p}^{e+1}$, contradicting primitivity of kP. This is (124).

When $\mathfrak{q}$ is coprime to δ_P, the two columns P, C form a basis modulo $\mathfrak{q}$. In that basis the lower coordinate of kP is a unit multiple of $\beta(k)$. Its vanishing is therefore exactly the claimed Borel congruence. In the same basis, vanishing of $\alpha(k)$ instead puts kP on the other coordinate line $[C]$. The two equations describe the two ends of the affine chart, and their common divisibility is bounded by the preceding calculation. □

This proposition also explains why replacing the Borel by one chosen unipotent would give a misleading surplus. A specified unipotent imposes extra equations such as $a = d = 1$, but the classification must retain the special hyperbolic alternative, whose Levi eigenvalue is not one. The correct target is the whole Borel, and that target has only the one transverse equation $b = 0$.

There is nevertheless a precise relative-descent formulation. In the hard case suppose that $G_7 \cap B_P(K)$ is projectively trivial. Choose a noncentral unipotent $u \in \Gamma_{\mathrm{cong}}(G_7) \cap B_P(K)$. Equality of the finite-adelic closures says that for every ideal $\mathfrak{q}$ there is $k_{\mathfrak{q}} \in G_7$ with $k_{\mathfrak{q}} \equiv u \pmod{\mathfrak{q}}$. Hence

$$\beta(k_{\mathfrak{q}}) \in \mathfrak{q}, \qquad \beta(k_{\mathfrak{q}}) \neq 0. \tag{126}$$

for all $\mathfrak{q}$ outside a fixed finite set: if the second entry vanished, projective triviality would give $k_{\mathfrak{q}} = \pm I$, which cannot be congruent to the fixed noncentral element u for arbitrarily large $\mathfrak{q}$. Finite arithmetic therefore permits arbitrarily deep false contact with the Borel. What has to be proved is that the particular contacts made by consecutive quotients of one algebraic Rosen ray pay strictly more stopped Hodge degree than their determinant ideal supplies. This is the point at which the order of the thin word, rather than another congruence condition, must enter.

There is one further obstruction to turning such a finite Borel contact into an element of the ambient integral Borel. It is elementary, but it is important because it rules out an apparently natural shortcut. For a nonzero ideal $\mathfrak{b}$, put

$$\mathscr{R}_{\mathfrak{b}} = (\mathcal{O}_K/\mathfrak{b})^{\times} / \operatorname{im}(\mathcal{O}_K^{\times}). \tag{127}$$

This finite group records the part of a residue-class diagonal which cannot be changed by a global unit.

Proposition 7.22 (The exact Levi residue of a Borel contact)**.** *In the adapted frame* (122)*, suppose* $b = \beta(k) \neq 0$*. The determinant identity makes* $\bar{a} \in (\mathcal{O}_K/(b))^{\times}$*, and the class*

$$\ell_P(k) = [\bar{a}] \in \mathscr{R}_{(b)} \tag{128}$$

is independent of the integral completion of P*. There is an integral Borel element*

$$u = \begin{pmatrix} \varepsilon & v \\ 0 & \varepsilon^{-1} \end{pmatrix}, \qquad \varepsilon \in \mathcal{O}_K^{\times}, quad v \in \mathcal{O}_K,$$

congruent to the adapted matrix of k *modulo* (b) *if and only if* $\ell_P(k) = 1$*.*

Proof. Modulo (b), the determinant equation reads $ad = 1$, so a is a unit in the quotient. Replacing the second vector of the adapted frame by $E + tP$ conjugates the matrix by $T(t)$ and replaces a by $a - tb$. Hence $\bar{a}$, and therefore (128), does not change.

If an integral Borel lift exists, its upper-left entry gives a global unit $\varepsilon \equiv a \pmod{(b)}$, so the class is trivial. Conversely, given such a unit, choose $v \equiv c \pmod{(b)}$. The congruence $d \equiv a^{-1} \equiv \varepsilon^{-1} \pmod{(b)}$ then shows that the displayed Borel element is congruent to the full adapted matrix. □

The class in (128) does not vanish automatically on the consecutive Rosen language. There is a small exact counterexample. Take

$$P = (6 + 2\lambda - \lambda^2, 1)^t.$$

Its first three Rosen digits are $4, 1, 1$. After the two inversion steps, exact conjugation into the frame $(P, (-1, 0)^t)$ gives first column $(a, b)^t$, where

$$a = 3\lambda^2 - 4\lambda - 2, qquadb = 11\lambda^2 - 9\lambda - 16, qquad|\,\mathrm{N}\, b| = 769. \tag{129}$$

The group $(\mathcal{O}_K/(b))^\times$ is cyclic of order 768. With one cyclic generator, the discrete logarithms of $a, \lambda, \alpha = \lambda^2 - 1, -1$ are respectively

$$100, quad519, quad432, quad384 \pmod{768}.$$

The last three generate the subgroup of multiples of 3, while 100 is not a multiple of 3. Since $\mathcal{O}_K^\times = \langle -1, \lambda, \alpha \rangle$, the Levi class has order three and no integral Borel lift exists. This is an exact finite-ring calculation, not a failure caused by numerical precision.

The same obstruction persists on a long genuine special orbit. The integer 671 has Rosen preperiod–period detection at step 61, with period 23. Among its 62 nontrivial initial-prefix Levi residues, exact ideal-star computations find 30 in the image of the global units and 32 outside it; the largest index of the unit image encountered is 1176. The reproducible calculation is `levi_residue_tests.py`. Thus an argument which first lifts every residue Borel to an integral Borel cannot be repaired by restricting to long Rosen words. A successful stopped inequality must retain the whole affine Borel frame, including this ray-class coordinate, or avoid choosing an integral Borel lift altogether.

The terminal all-place rigidity statement itself is much simpler. The following elementary lemma is the sharp norm threshold which any Hodge estimate must cross.

Lemma 7.23 (All-place ideal vanishing threshold)**.** *Let $0 \neq \mathfrak{q} \subset \mathcal{O}_K$ and let $x \in \mathfrak{q}$. If*

$$\sum_{i=0}^{2} \log^+ |\sigma_i(x)| < \log \mathrm{N}\, \mathfrak{q}, \tag{130}$$

then $x = 0$. The same statement for a fixed fractional lattice differs only by an additive constant depending on that lattice.

Proof. If $x \neq 0$, then $(x) \subseteq \mathfrak{q}$, and hence $|\,\mathrm{N}\, x| \geq \mathrm{N}\, \mathfrak{q}$. On the other hand,

$$\log |\,\mathrm{N}\, x| = \sum_i \log |\sigma_i(x)| \leq \sum_i \log^+ |\sigma_i(x)|,$$

contradicting (130). Multiplication by a fixed denominator reduces the fractional-lattice version to this one. □

For $x = \beta(k)$, this lemma would make a sufficiently short congruence Borel contact an actual Borel incidence. It also shows why the lemma alone makes no progress: choosing $\mathfrak{q} = (\beta(k))$ puts one exactly on the product-formula boundary, with no strict slack. Problem 17.17 asks the consecutive Hodge transport to provide that slack while charging the determinant ideal once. The ray-Levi calculation above shows that it cannot be manufactured by first normalizing every contact with a global unit.

8 A group-wide integral Asai obstruction

There is a second use of the whole thin-group representation which does not mention a Rosen move or a norm record. It turns the classification into a rigidity problem for exceptional Lyapunov behavior.

Let $x = [P] \in \mathbf{P}^1(K)$, with $P \in \mathcal{O}_K^2$ primitive, and put

$$Y(P) = \sigma_0(P) \otimes \sigma_1(P) \otimes \sigma_2(P) \in W = V_0 \otimes V_1 \otimes V_2.$$

This is a nonzero vector in the integral tensor-induced, or Asai, lattice. In this section $\log \|v\|$ always means the logarithm of the *squared* Hodge norm. This is the Busemann normalization in which the top uniformizing exponent is one. Over a compact set it is comparable with twice the logarithm of a Euclidean norm; retaining this factor of two is essential when the Hodge quantity is identified with projective height. Let z_t be the unit-speed geodesic ray in $\mathbb{H}$ with forward endpoint $x_0^\dagger$, and let $\|\cdot\|_{i,t}$ be the Hodge norm in the i-th factor along this ray. The source endpoint line is the contracting line of the uniformizing variation, so

$$\log \|\sigma_0(P)\|_{0,t} = -t + O_P(1). \tag{131}$$

The normalization is the one in which the top uniformizing Lyapunov exponent is one.

Proposition 8.1 (Arithmetic lower bound for the outer exponents)**.** *Suppose that x is not a cusp. Then the ray z_t has a sequence of returns $t_n \to \infty$ to a fixed compact part of the triangle surface, and along every such sequence*

$$\limsup_{n\to\infty} \frac{1}{t_n} \left(\log \|\sigma_1(P)\|_{1,t_n} + \log \|\sigma_2(P)\|_{2,t_n}\right) \geq 1. \tag{132}$$

In particular, this ray cannot be Oseledets regular with outer upper exponents whose sum is less than one.

Proof. A geodesic ray on a finite-area hyperbolic orbifold which eventually leaves every compact set ends at a parabolic fixed point. Since x is not a cusp, compact returns t_n exist. Reduce each return to a fixed compact fundamental set by $g_n \in G_7$. Equivariance of the Hodge metric identifies the norm of $Y(P)$ at z_{t_n} with the norm, at a point of that compact set, of

$$\mathrm{As}(g_n)Y(P) = Y(g_n P).$$

The latter is a nonzero vector in one fixed integral lattice. The shortest nonzero lattice-vector norm has a positive lower bound on a compact set. Consequently

$$\log \|Y(P)\|_{t_n} = \sum_{i=0}^{2} \log \|\sigma_i(P)\|_{i,t_n} \geq -O(1).$$

Insert (131), divide by t_n, and take a limsup. This gives (132). □

The preceding exponent bound hides a stronger, finite-height statement. It is useful to record it because it identifies the precise corridor in which a counterexample would have to escape. For a compact return time t, define the *outer compensation excess*

$$\mathcal{D}_P(t) = \log \|\sigma_1(P)\|_{1,t} + \log \|\sigma_2(P)\|_{2,t} - t. \tag{133}$$

The term t is not a fitted normalization: it is exactly the contraction of the uniformizing factor in (131), up to a bounded endpoint constant.

Proposition 8.2 (Compensation is the height of the reduced orbit)**.** *Fix a compact set C in the triangle surface. Suppose that t_n is a return time and that $g_n \in G_7$ reduces the return to C. Then*

$$\mathcal{D}_P(t_n) = 6h([g_nP]) + O_{P,C}(1), \tag{134}$$

where the error is uniform in n.

Proof. The tensor norm and (131) give

$$\log \|Y(P)\|_{t_n} = \mathcal{D}_P(t_n) + O_P(1).$$

Equivariance identifies the left side with the norm of $Y(g_nP)$ over the reduced point in C. On a compact set, each squared Hodge norm is uniformly comparable with the square of a fixed Euclidean norm on the corresponding real two-space. Therefore

$$\log \|Y(g_nP)\| = 2\sum_{i=0}^{2} \log \max\{|\sigma_i((g_nP)_1)|, |\sigma_i((g_nP)_2)|\} + O_C(1).$$

The vector g_nP is primitive because both g_n and P are integral and unimodular. Formula (3) identifies twice the sum with $6h([g_nP])$, proving (134). □

Proposition 8.3 (Exact escape forced on a counterexample)**.** *Let $x = [P] \in \mathbf{P}^1(K)$ be noncuspidal, and let $t_n \to \infty$ be any sequence along which the geodesic ray to x returns to a fixed compact part of the triangle surface. If*

$$\liminf_{n\to\infty} \mathcal{D}_P(t_n) < \infty, \tag{135}$$

then x is fixed by a projectively nonidentity element of G_7, and hence is a cusp or a special hyperbolic fixed point. Consequently, if x is neither a cusp nor special, then

$$\mathcal{D}_P(t_n) \longrightarrow +\infty \tag{136}$$

along every compact-return sequence.

Proof. Pass to a subsequence on which $\mathcal{D}_P(t_n)$ is bounded above, and choose $g_n \in G_7$ which reduces the n-th return to one fixed compact set. Proposition 8.2 bounds the projective heights $h([g_nP])$. Northcott's theorem leaves only finitely many possibilities for $[g_nP]$. Two are equal. The reducing elements may be taken distinct because $t_n \to \infty$. Hence the projectively nonidentity element $g_n^{-1}g_m$ fixes $[P]$. Proposition 7.1 gives the stated cusp-or-special alternative. □

Proposition 8.4 (The quantitative Northcott barrier)**.** *Let $t_1 < \cdots < t_N$ be compact returns reduced by distinct $g_1, \ldots, g_N \in G_7$ to one fixed compact set. If $[P]$ has trivial projective stabilizer in G_7, then*

$$\max_{1\le n\le N} \mathcal{D}_P(t_n) \ge \log N - O_{K,P,C}(1). \tag{137}$$

Consequently, if for some $\delta > 0$ and arbitrarily large N all the first N reduced states satisfy

$$\mathcal{D}_P(t_n) \le (1-\delta)\log N,$$

then $[P]$ is a cusp or a special hyperbolic fixed point.

Proof. There is a constant A_K such that

$$\#\{y \in \mathbf{P}^1(K) : h(y) \le H\} \le A_K e^{6H}. \tag{138}$$

Indeed, use Lemma 3.2 to choose a primitive representative whose two coordinates have size at most $e^{H+O_K(1)}$ in all three real embeddings. The Minkowski lattice $\mathcal{O}_K^2$ has rank six, so the number of such pairs is $O_K(e^{6H})$.

Triviality of the stabilizer makes the points $[g_nP]$ pairwise distinct. Proposition 8.2 and (138) therefore give

$$N \le A_{K,P,C} \exp\left(\max_{n\le N} \mathcal{D}_P(t_n)\right).$$

Taking logarithms proves (137). The last assertion follows because $N > A_{K,P,C}N^{1-\delta}$ for large N, forcing two reduced states to coincide. □

What is elementary here. This is only quantitative pigeonhole, not the missing rigidity theorem. It shows exactly how much ordinary height counting can see: a nonperiodic orbit must create at least logarithmic compensation in the number of returns. Since $\log N = o(t_N)$ in every ordinary recurrent regime, the bound is compatible with the sublinear equality corridor. The needed Arakelov input must therefore control the cocycle along the stopped orbit, not merely count field points of bounded height.

What is proved here. No ergodic theorem and no estimate for a Rosen generator enters this proposition. It is the full thin-group monodromy, the three-factor integral height, and compact-return reduction. It proves more than the lower exponent bound: a hypothetical nonperiodic field point must make the positive Galois compensation escape through rational points of unbounded height. The least exceptional still possible regime is

$$\mathcal{D}_P(t) \longrightarrow \infty, \qquad \mathcal{D}_P(t) = o(t),$$

so that the outer exponent is exactly 1 although no bounded-height state recurs. Ruling out this drifting equality corridor is the precise boundary case of the new rigidity statement; a complete exceptional-ray theorem must also exclude linearly supercritical outer growth. Ordinary degree zero rules out neither behavior.

8.1 The equality corridor as a Borel-frame problem

Proposition 8.5 (Unit-balanced form of height escape)**.** *There is a constant C_K with the following property. For every primitive $Q = (u,v)^t \in \mathcal{O}_K^2$, one can choose a unit $\varepsilon \in \mathcal{O}_K^\times$ such that, for $i = 0,1,2$,*

$$|\log \max\{|\sigma_i(\varepsilon u)|, |\sigma_i(\varepsilon v)|\} - h([Q])| \le C_K. \tag{139}$$

Consequently, if g_n reduces compact returns at times t_n, then

$$\mathcal{D}_P(t_n) = o(t_n) \tag{140}$$

if and only if there are units ε_n for which all three affine vectors $\sigma_i(\varepsilon_n g_n P)$ have logarithmic size $o(t_n)$.

Proof. Put

$$H_i(Q) = \log \max\{|\sigma_i(u)|, |\sigma_i(v)|\}.$$

Primitivity and (3) give $H_0 + H_1 + H_2 = 3h([Q])$. The logarithmic unit vectors

$$\big(\log|\sigma_0\varepsilon|, \log|\sigma_1\varepsilon|, \log|\sigma_2\varepsilon|\big)$$

form a cocompact lattice in the plane $x_0 + x_1 + x_2 = 0$. Translate $(H_i - h([Q]))_i$ into one fixed fundamental parallelogram of that lattice. This proves (139). The last assertion follows from Proposition 8.2. □

This proposition gives a concrete geometric meaning to the boundary case. A sublinear compensation excess says that, after using the full rank-two unit lattice of the ambient rational Borel, the three Galois components of the reduced endpoint vector are simultaneously subexponential. The desired equality theorem can therefore be stated as a sublinear-tracking rigidity assertion: a thin-group geodesic which tracks a rational Hilbert-modular Borel to sublinear transverse height must actually meet that Borel. The intersection is then parabolic or a unit eigenline by Proposition 7.1. Filip proves this when the transverse height tends strictly downward. The new work is to deduce the intersection when the transverse height drifts upward only sublinearly; the unit balancing above shows that this is an all-place statement rather than a choice of one real embedding.

The word "tracks" in the preceding paragraph can be made algebraic. For $A \in \mathrm{SL}_2(K)$, put

$$\mathfrak{h}_\Delta(A) = \sum_{i=0}^{2} \log \max\{1, \|\sigma_i(A)\|_\infty\}.$$

This is an all-place matrix size; on the integral group it is comparable with displacement in the symmetric space of $\mathrm{Res}_{K/\mathbf{Q}}\,\mathrm{SL}_2$.

Lemma 8.6 (A balanced primitive column has a controlled frame)**.** *There are constants A_K, B_K such that the following holds. Suppose that $Q = (u, v)^t \in \mathcal{O}_K^2$ is primitive and*

$$|\log \max\{|\sigma_i(u)|, |\sigma_i(v)|\} - h([Q])| \le C_K \quad (i = 0, 1, 2). \tag{141}$$

Then Q is the first column of a matrix $H_Q \in \mathrm{SL}_2(\mathcal{O}_K)$ satisfying

$$\mathfrak{h}_\Delta(H_Q) \le A_K h([Q]) + B_K. \tag{142}$$

Proof. Primitivity makes the map

$$\mathcal{O}_K^2 \longrightarrow \mathcal{O}_K, \qquad (a, b) \longmapsto ub - va$$

surjective. Its solution set over 1 is an affine translate of the rank-three Minkowski lattice $\{s(u, v) : s \in \mathcal{O}_K\}$. Reduce one solution modulo a fixed fundamental parallelepiped of this lattice. At the i-th real embedding, the component perpendicular to $\sigma_i(u, v)$ has norm $\|\sigma_i(u, v)\|^{-1}$, because its determinant with that vector is one; the component parallel to it is bounded by a fixed multiple of $\|\sigma_i(u, v)\|$. Hypothesis (141) bounds both by $\exp(h([Q]) + O_K(1))$. Taking the reduced solution as the second column gives determinant one and (142). □

Proposition 8.7 (Exact sublinear Borel factorization)**.** *Fix a primitive endpoint column P and a unimodular frame H_P with first column P. Let $g_n \in G_7$, put $X_n = g_n P$, and choose the balancing units ε_n of Proposition 8.5. There are matrices*

$$C_n \in \mathrm{SL}_2(\mathcal{O}_K), \qquad b_n \in B_P(K) \cap \mathrm{SL}_2(\mathcal{O}_K) \tag{143}$$

such that

$$g_n = C_n b_n, \qquad \mathfrak{h}_\Delta(C_n) \le A_K h([X_n]) + O_{K,P}(1). \tag{144}$$

In particular, along compact returns with $\mathcal{D}_P(t_n) = o(t_n)$,

$$\mathfrak{h}_\Delta(C_n) = o(t_n). \tag{145}$$

Proof. Put $Q_n = \varepsilon_n X_n$. This column is primitive and satisfies (141). Choose its controlled frame H_{Q_n} from Lemma 8.6, and set

$$C_n = H_{Q_n} H_P^{-1}, \qquad b_n = C_n^{-1} g_n.$$

Then

$$b_n P = C_n^{-1} X_n = H_P H_{Q_n}^{-1}(\varepsilon_n^{-1} Q_n) = \varepsilon_n^{-1} P.$$

Thus b_n fixes $[P]$, so it lies in $B_P(K)$; it is integral and unimodular because all the other matrices are. The height bound follows from (142) and the fixed size of H_P. Finally use (134). □

Proposition 8.8 (The quotient height measures the complete frame)**.** *For a right coset in the arithmetic flag orbit, define*

$$\mathfrak{H}_P(g\mathscr{B}_P) = \inf_{b \in \mathscr{B}_P} \mathfrak{h}_\Delta(gb), \qquad \mathscr{B}_P = B_P(K) \cap \mathrm{SL}_2(\mathcal{O}_K). \tag{146}$$

There are constants A, B*, depending only on* K *and the fixed primitive column* P*, such that for every* $g \in \mathrm{SL}_2(\mathcal{O}_K)$*,*

$$3h([gP]) - B \le \mathfrak{H}_P(g\mathscr{B}_P) \le Ah([gP]) + B. \tag{147}$$

In particular, $\mathfrak{H}_P$ *is a proper height on the orbit of the rational flag* $[P]$*. At compact returns,*

$$\tfrac{1}{2}\mathcal{D}_P(t_n) - O_{P,C}(1) \le \mathfrak{H}_P(g_n\mathscr{B}_P) \le \tfrac{A}{6}\mathcal{D}_P(t_n) + O_{P,C}(1). \tag{148}$$

Thus "sublinear tracking of a Borel" controls the complete rational state, not only a diagonal unit parameter.

Proof. For the upper bound, apply the unit balancing and controlled completion used in Proposition 8.7; they produce one representative $C \in g\mathscr{B}_P$ with $\mathfrak{h}_\Delta(C) \le Ah([gP]) + B$.

For the lower bound, let $C = gb$ with $b \in \mathscr{B}_P$. Since $bP = \eta P$ for a unit η, one has $CP = \eta gP$. If

$$M_i = \max\{|\sigma_i((gP)_1)|, |\sigma_i((gP)_2)|\},$$

then the norm of $\sigma_i(C)$ is at least a fixed positive multiple, depending on P, of $|\sigma_i(\eta)|M_i$. Hence

$$\mathfrak{h}_\Delta(C) \ge \sum_{i=0}^{2} \max\{0, \log M_i + \log|\sigma_i(\eta)| - O_P(1)\} \ge \sum_{i=0}^{2} \log M_i - O_P(1).$$

The unit terms sum to zero by the product formula. Since gP is primitive, the last sum is $3h([gP])$. Taking the infimum gives the lower bound. Properness is Northcott's theorem together with the identification of this coset orbit with the orbit of $[P]$. Finally substitute $\mathcal{D}_P(t_n) = 6h([g_nP]) + O_{P,C}(1)$. □

Proposition 8.9 (Exact transition equation for the Borel frames)**.** *Use the factors of Proposition 8.7, put*

$$\mathscr{B}_P = B_P(K) \cap \mathrm{SL}_2(\mathcal{O}_K), \qquad a_n = g_{n+1} g_n^{-1}, \qquad Q_n = C_n P = \varepsilon_n g_n P.$$

Then

$$a_n = C_{n+1} u_n C_n^{-1}, \qquad u_n = b_{n+1} b_n^{-1} \in \mathscr{B}_P. \tag{149}$$

In the fixed frame H_P, there is a uniquely determined $\beta_n \in \mathcal{O}_K$ such that

$$H_P^{-1} u_n H_P = \begin{pmatrix} \varepsilon_n/\varepsilon_{n+1} & \beta_n \\ 0 & \varepsilon_{n+1}/\varepsilon_n \end{pmatrix}. \tag{150}$$

Consequently

$$a_n Q_n = \frac{\varepsilon_n}{\varepsilon_{n+1}} Q_{n+1}. \tag{151}$$

Moreover,

$$C_n \mathscr{B}_P = g_n \mathscr{B}_P, \tag{152}$$

and for $m < n$ the following are equivalent:

1. $C_m \mathscr{B}_P = C_n \mathscr{B}_P$;
2. $[g_m P] = [g_n P]$;
3. $g_n^{-1} g_m \in \mathscr{B}_P$.

Thus repetition of the arithmetic Borel frame is exactly repetition of the complete rational state.

Proof. The identity $g_n = C_n b_n$ gives

$$g_{n+1} g_n^{-1} = C_{n+1} (b_{n+1} b_n^{-1}) C_n^{-1},$$

which proves (149). Since $b_n P = \varepsilon_n^{-1} P$, the element u_n acts on P by $\varepsilon_n/\varepsilon_{n+1}$. Conjugation by the integral unimodular frame H_P identifies $\mathscr{B}_P$ with the integral upper triangular group. This proves (150), and applying (149) to $C_n P$ proves (151).

Since $b_n \in \mathscr{B}_P$, equation (152) is immediate. The coset space $\mathrm{SL}_2(\mathcal{O}_K)/\mathscr{B}_P$ is the orbit of the line $[P]$, so equality of two cosets is equivalent to $[g_m P] = [g_n P]$. Multiplying that equality by g_n^{-1} gives the third condition. □

The entry β_n is not itself an extra state coordinate. It depends on the choice of the second column in each balanced completion. More precisely, for $s_n \in \mathcal{O}_K$ put

$$v_n = H_P \begin{pmatrix} 1 & s_n \\ 0 & 1 \end{pmatrix} H_P^{-1} \in \mathscr{B}_P.$$

Replacing C_n by $C_n v_n$ and b_n by $v_n^{-1} b_n$ leaves g_n, Q_n, and the coset $C_n \mathscr{B}_P$ unchanged. If $d_n = \varepsilon_n/\varepsilon_{n+1}$, it replaces the upper-right entry by

$$\beta_n' = \beta_n + d_n s_n - d_n^{-1} s_{n+1}. \tag{153}$$

Thus a bound for β_n would be coordinate-dependent. Its role in (150) is to prove that the diagonal unit does not determine a lifted transition. The invariant closing variable is the quotient $C_n \mathscr{B}_P$, in which this unipotent gauge has already been removed.

The integral completion used in the preceding Borel calculation is not the geometrically natural past endpoint. We now separate the two normalizations. Assume that $[P]$ is not already a cusp, choose a primitive cusp column $C \in \mathcal{O}_K^2$, and put

$$\delta_P = \det(P, C) \neq 0, \qquad E = \delta_P^{-1} C, \qquad \widehat{H}_P = (P, E) \in \mathrm{SL}_2(K). \tag{154}$$

Thus $[E]$ is a cusp and $\det(P, E) = 1$, but E generally lies in the fixed fractional lattice $\delta_P^{-1}\mathcal{O}_K^2$, not in $\mathcal{O}_K^2$. This is harmless for heights: every change from the integral lattice below is an $O_P(1)$ term. It is logically important, however, because an incidence with the transported line $[E]$ is now genuinely the cusp alternative. From this point on, E denotes this cusp-normalized past column; the earlier integral completion is used only to choose controlled Borel frames.

Proposition 8.10 (What the natural-extension frame costs)**.** *For a reduced return define the dynamically transported cusp frame*

$$Q_n = \varepsilon_n g_n P, \qquad R_n^{\mathrm{dyn}} = \varepsilon_n^{-1} g_n E. \tag{155}$$

Then

$$\det(Q_n, R_n^{\mathrm{dyn}}) = 1 \tag{156}$$

and its three-place transverse proximity is

$$\sum_{i=0}^{2} \log \frac{\|\sigma_i Q_n\|_\infty \|\sigma_i R_n^{\mathrm{dyn}}\|_\infty}{|\det(\sigma_i Q_n, \sigma_i R_n^{\mathrm{dyn}})|} = 3h([g_n P]) + 3h([g_n E]) + O_P(1). \tag{157}$$

If $H_{Q_n} = (Q_n, R_n^{\mathrm{ctl}})$ is any controlled unimodular completion from Lemma 8.6, then there is a unique $s_n \in \delta_P^{-1}\mathcal{O}_K$ such that

$$R_n^{\mathrm{dyn}} = R_n^{\mathrm{ctl}} + s_n Q_n. \tag{158}$$

Thus replacing the controlled completion by the dynamically correct past flag restores exactly one additive Borel gauge. Its height cannot be bounded in terms of $h([g_n P])$ by algebraic number theory alone.

Proof. Both scalings in (155) are units, so

$$\det(Q_n, R_n^{\mathrm{dyn}}) = \det(g_n P, g_n E) = \det(P, E) = 1.$$

Apply the product-formula computation of Proposition 8.32 to these two columns. Their determinant norm is one, and projective height is unchanged by the balancing units. Replacing the primitive integral cusp column C by $E = C/\delta_P$ changes the archimedean sum by the fixed amount $-\log|\mathrm{N}\,\delta_P|$. This gives (157).

Two second columns completing the same primitive first column to determinant one differ by a unique scalar multiple of that first column. Since one column is integral and the other lies in the fixed lattice $\delta_P^{-1}\mathcal{O}_K^2$, primitivity of Q_n puts that scalar in $\delta_P^{-1}\mathcal{O}_K$. This proves (158). Finally, in the fixed rational frame $\widehat{H}_P$, the rational Borel elements

$$\widehat{H}_P \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix} \widehat{H}_P^{-1}, \qquad s \in \delta_P^{-1}\mathcal{O}_K,$$

fix $[P]$ and send E to $E + sP$. The height of the latter line is unbounded while the first line is unchanged. Hence no inequality bounding the second height by the first can hold without using the thin orbit and its Hodge dynamics. □

Why the natural extension does not finish the proof. The pair $(Q_n, R_n^{\mathrm{dyn}})$ has the correct stable and unstable boundary flags, so its geodesic is the actual natural-extension orbit. It pays for that advantage by the uncontrolled second height in (157). The controlled completion has sublinear height in the equality corridor, but (158) changes its past endpoint and loses the exact orbit path. The coefficient s_n is the cumulative version of the gauge in (153). Thus using the past endpoint does not remove the obstruction: it moves it from path comparison into an additive endpoint-height term.

Proposition 8.11 (The natural-extension gauge is an affine unit cocycle)**.** *Fix one deterministic choice of controlled completion at every state and let s_n be defined by (158). With $d_n = \varepsilon_n/\varepsilon_{n+1}$ and β_n as in (150), one has*

$$s_{n+1} = d_n^2 s_n + d_n \beta_n, \qquad d_n \in \mathcal{O}_K^{\times}, \quad s_n \in \delta_P^{-1}\mathcal{O}_K, \quad \beta_n \in \mathcal{O}_K. \tag{159}$$

Thus the diagonal Hodge-neutral transport is multiplicative, but the failure of the controlled past frame to be flat is an additive integral extension of that unit cocycle.

Proof. The dynamically transported frames satisfy

$$a_n Q_n = d_n Q_{n+1}, \qquad a_n R_n^{\mathrm{dyn}} = d_n^{-1} R_{n+1}^{\mathrm{dyn}}.$$

The controlled transition equation gives

$$a_n R_n^{\mathrm{ctl}} = \beta_n Q_{n+1} + d_n^{-1} R_{n+1}^{\mathrm{ctl}}.$$

Apply a_n to $R_n^{\mathrm{dyn}} = R_n^{\mathrm{ctl}} + s_n Q_n$, substitute the three displayed relations, and compare the coefficient of Q_{n+1}. This gives $\beta_n + d_n s_n = d_n^{-1} s_{n+1}$, which is (159). □

Proposition 8.12 (The actual path accumulates the Hodge defect)**.** *Let (P, E) be the determinant-one future–past frame of an oriented geodesic $\gamma(t)$. For each embedding, let $f_{P,i}, f_{E,i}$ be the two period functions and put*

$$R_i(t) = \frac{f_{E,i}(\gamma(t))}{f_{P,i}(\gamma(t))}. \tag{160}$$

Orient arclength so that $d\log|R_0(t)|/dt = 1$, and define

$$\mathfrak{A}_{P,E}(t) = 2\log|R_0(t)| - \log|R_1(t)| - \log|R_2(t)|. \tag{161}$$

Then

$$\frac{d}{dt}\mathfrak{A}_{P,E}(t) \geq 2 - \|DF_1(\gamma(t))\|_{\mathrm{hyp}} - \|DF_2(\gamma(t))\|_{\mathrm{hyp}}. \tag{162}$$

For $t_m < t_n$, write

$$\mathfrak{C}(m,n) = \int_{t_m}^{t_n} \left(2 - \|DF_1(\gamma(s))\|_{\mathrm{hyp}} - \|DF_2(\gamma(s))\|_{\mathrm{hyp}}\right) ds. \tag{163}$$

Consequently, for all $t_m < t_n$,

$$\mathfrak{A}_{P,E}(t_n) - \mathfrak{A}_{P,E}(t_m) \geq \mathfrak{C}(m,n). \tag{164}$$

This is an inequality on the original stopped orbit segment, not on the geodesic joining two rational future states.

Proof. The real Möbius coordinate R_0 sends the central geodesic with endpoints $[P], [E]$ to the imaginary axis. Hyperbolic arclength on that axis is $|d \log |R_0||$, which gives the chosen normalization. By (299), now applied to the pair (P, E),

$$\frac{d \log R_i}{d \log R_0} = \frac{J_i^{P,E}}{J_0^{P,E}}.$$

Proposition 17.7 applies at every point of this very geodesic and gives

$$\left| \frac{d}{dt} \log |R_i(t)| \right| \leq \|DF_i(\gamma(t))\|_{\mathrm{hyp}}.$$

Substitution in the derivative of (161) proves (162); integration proves (164). □

Proposition 8.13 (Exact boundary term between dynamic and controlled Wronskians)**.** *At a reduced return, abbreviate*

$$Q = Q_n, \qquad R = R_n^{\mathrm{ctl}}, \qquad R^{\mathrm{dyn}} = R + s_n Q.$$

For the three period evaluations write $f_{Q,i}, f_{R,i}, f_{\mathrm{dyn},i}$, and let J_i^{ctl}, J_i^{dyn} be the normalized Wronskians of the pairs (Q, R), (Q, R^{dyn}). Then

$$W_i(Q, R^{\mathrm{dyn}}) = W_i(Q, R), \tag{165}$$

$$\frac{J_i^{\mathrm{dyn}}}{J_i^{\mathrm{ctl}}} = \frac{f_{R,i}}{f_{\mathrm{dyn},i}}, \qquad f_{\mathrm{dyn},i} = f_{R,i} + \sigma_i(s_n) f_{Q,i}. \tag{166}$$

In particular, the difference of the cyclic Wronskian metrics is the single endpoint term

$$\mathfrak{B}_n = \log \left| \frac{f_{R,1} f_{R,2} f_{\mathrm{dyn},0}^2}{f_{\mathrm{dyn},1} f_{\mathrm{dyn},2} f_{R,0}^2} \right|, \tag{167}$$

and this is also the difference between the dynamic and controlled potentials (161). Hence every failure to transfer (164) to the controlled integral Wronskian is concentrated in $\mathfrak{B}_n - \mathfrak{B}_m$.

Proof. The two second columns have the same determinant with Q:

$$\det(Q, R + s_n Q) = \det(Q, R) = 1.$$

The factorization (297) therefore gives (165). Dividing by the two period values as in (299) gives (166); flatness of the covectors gives the displayed affine formula for $f_{\mathrm{dyn},i}$. Take the outer product divided by the square of the central term to obtain (167). The same quotient results by subtracting the two definitions of $\mathfrak{A}$. □

Corollary 8.14 (The two-future Wronskian is a rational splitting of the dynamic extension)**.** *Assume $m < n$ and $\Delta_{m,n} \neq 0$. Put*

$$\widetilde{R}_{m,n} = \frac{Q_n}{\Delta_{m,n}}, \qquad A_{m,n} = \det(Q_n, R_m^{\mathrm{dyn}}) \in \delta_P^{-1} \mathcal{O}_K. \tag{168}$$

Then $\det(Q_m, \widetilde{R}_{m,n}) = 1$, and

$$R_m^{\mathrm{dyn}} = \widetilde{R}_{m,n} + s_{m,n} Q_m, \qquad s_{m,n} = -\frac{A_{m,n}}{\Delta_{m,n}} \in (\delta_P \Delta_{m,n})^{-1} \mathcal{O}_K. \tag{169}$$

Moreover, at every embedding,

$$J_i(Q_m, \widetilde{R}_{m,n}) = J_i^{m,n}. \tag{170}$$

Thus the exact boundary formula of Proposition 8.13 applies to the stopped two-future Wronskian after replacing (R, s_n) there by $(\widetilde{R}_{m,n}, s_{m,n})$. Its varying finite denominator is precisely the already present ideal $(\Delta_{m,n})$; the factor (δ_P) belongs to the fixed cusp normalization and contributes only $O_P(1)$.

Proof. The first identity follows immediately from the definition of $\Delta_{m,n}$. Two determinant-one completions of Q_m differ by a multiple of Q_m, so write $R_m^{\mathrm{dyn}} = \widetilde{R}_{m,n} + sQ_m$. Taking determinant with $\widetilde{R}_{m,n}$ gives

$$-s = \det(\widetilde{R}_{m,n}, R_m^{\mathrm{dyn}}) = \frac{A_{m,n}}{\Delta_{m,n}},$$

which proves (169). Finally, scaling the second column by $\sigma_i(\Delta_{m,n})^{-1}$ divides both its Wronskian numerator and its period value by the same scalar. The normalized jet is unchanged, proving (170). □

The two integral coordinates have the elementary expansion

$$Q_n = A_{m,n}Q_m + \Delta_{m,n}R_m^{\mathrm{dyn}}. \tag{171}$$

They also display the two exact closing divisors. The equation $\Delta_{m,n} = 0$ is recurrence of the rational future line and hence already gives a nonidentity Borel stabilizer. If $A_{m,n} = 0$, then $[Q_n] = [R_m^{\mathrm{dyn}}]$; by (154), the right side is in the G_7-orbit of a cusp. Applying g_n^{-1} shows that $[P]$ itself is a cusp. This is why the fixed fractional lattice, rather than an arbitrary integral completion, is used in the active argument. Away from these two divisors, the residual endpoint is the genuine rational slope $-A_{m,n}/\Delta_{m,n}$. Consequently the remaining all-place estimate is a two-coordinate problem with exactly the two desired equality components, not an unspecified comparison of paths.

Proposition 8.15 (The endpoint coordinates are one pointed thin-group orbit)**.** *Let*

$$k_{m,n} = g_m^{-1}g_n \in G_7, \qquad \widehat{H}_P = (P, E) \in \mathrm{SL}_2(K),$$

and write

$$\widehat{H}_P^{-1}k_{m,n}P = \begin{pmatrix} a_{m,n} \\ b_{m,n} \end{pmatrix}. \tag{172}$$

Then

$$a_{m,n} = \det(k_{m,n}P, E), \qquad b_{m,n} = \det(P, k_{m,n}P) = Q_{k_{m,n}}(P), \tag{173}$$

$$A_{m,n} = \frac{\varepsilon_n}{\varepsilon_m}a_{m,n}, \qquad \Delta_{m,n} = \varepsilon_m\varepsilon_n b_{m,n}. \tag{174}$$

The pair $(a_{m,n}, b_{m,n})$ is primitive in the fixed lattice $\widehat{H}_P^{-1}\mathcal{O}_K^2$; explicitly, $a_{m,n} \in \delta_P^{-1}\mathcal{O}_K$ and $b_{m,n} \in \mathcal{O}_K$. In particular, the two-frame Wronskian determinant is, up to a unit, exactly the integral fixed-point form of the full interval quotient, and

$$s_{m,n} = -\varepsilon_m^{-2}\frac{a_{m,n}}{b_{m,n}}. \tag{175}$$

Thus the slopes which occur in the endpoint problem form the pointed orbit

$$\widehat{H}_P^{-1}G_7\widehat{H}_P \cdot [1:0] \subset \mathbf{P}^1(K), \tag{176}$$

with the additional restriction that $k_{m,n}$ is a consecutive quotient along one geodesic prefix sequence.

Proof. If $kP = aP + bE$, taking determinants with E and P gives (173). The second identity there is the definition (496). Since kP is a primitive integral column, its image is primitive in the transported lattice $\widehat{H}_P^{-1}\mathcal{O}_K^2$. Moreover, $a = \delta_P^{-1}\det(kP, C) \in \delta_P^{-1}\mathcal{O}_K$, while the displayed determinant formula puts b in $\mathcal{O}_K$.

Using

$$Q_r = \varepsilon_r g_r P, \qquad R_m^{\mathrm{dyn}} = \varepsilon_m^{-1} g_m E,$$

the equality $kP = aP + bE$, multiplied by $\varepsilon_n g_m$, becomes

$$Q_n = \frac{\varepsilon_n}{\varepsilon_m} a Q_m + \varepsilon_m \varepsilon_n b R_m^{\mathrm{dyn}}.$$

Comparison with (171) proves (174); substitution in (169) gives (175). Projectivizing (172) gives (176). □

> **Where the thin group now enters.** In the ambient arithmetic group, every primitive pair (a, b) occurs, so no height inequality for the universal tautological extension alone can prove the theorem. Proposition 8.15 restricts the pair to one ordered orbit of the Zariski-dense, infinite-index group G_7, and in fact to quotients of nested prefixes of one geodesic. The endpoint absorption theorem must use this restriction. It is therefore neither a calculation with individual Rosen letters nor a general positivity statement on the ambient Hilbert modular threefold; it is a relative domination theorem for the pointed thin orbit.

The endpoint term also has a coordinate-free description. If Z, U, V are nonzero columns in a two-dimensional complex vector space, write

$$[Z; U, V] = \frac{\det(Z, U)}{\det(Z, V)}.$$

For two period points $Z(t_m), Z(t_n)$, the quotient $[Z(t_n); U, V]/[Z(t_m); U, V]$ is the usual cross ratio of the four projective lines. Unlike the separate determinants, it is unchanged by rescaling any of the four columns.

Proposition 8.16 (The endpoint error is one Galois-weighted cross ratio)**.** *Let $Z_i(t) \in \mathbf{P}^1(\mathbf{C})$ be the period point of the i-th rank-two Hodge summand, in a flat trivialization along the stopped segment. With $U = R_m^{\mathrm{dyn}}$ and $V = \widetilde{R}_{m,n}$, put*

$$\mathfrak{B}_{m,n}(t) = 2\log\left|\frac{f_{U,0}(t)}{f_{V,0}(t)}\right| - \sum_{i=1}^{2}\log\left|\frac{f_{U,i}(t)}{f_{V,i}(t)}\right|. \tag{177}$$

Also put

$$\mathrm{CR}_i(m, n) = \frac{[Z_i(t_n); \sigma_i U, \sigma_i V]}{[Z_i(t_m); \sigma_i U, \sigma_i V]}. \tag{178}$$

Whenever none of the displayed incidences vanishes,

$$\mathfrak{B}_{m,n}(t_n) - \mathfrak{B}_{m,n}(t_m) = 2\log|\,\mathrm{CR}_0(m, n)| - \log|\,\mathrm{CR}_1(m, n)| - \log|\,\mathrm{CR}_2(m, n)|. \tag{179}$$

After applying the common projective transformation g_m^{-1}, the two algebraic endpoint lines in every factor are

$$[E], \qquad [k_{m,n}P]. \tag{180}$$

Consequently the boundary problem is a cross-ratio distortion problem for the single pointed orbit $\widehat{H}_P^{-1}G_7\widehat{H}_P[1:0]$, restricted to consecutive geodesic quotients. It is not a statement about independent prime divisors or about arbitrary primitive pairs in K^2.

Proof. Under the polarization, a period evaluation $f_{U,i}(t)$ is, up to a nonzero constant independent of t, $\det(Z_i(t), \sigma_i U)$. The constant cancels from the quotient at the two times, so

$$\frac{f_{U,i}(t_n)/f_{V,i}(t_n)}{f_{U,i}(t_m)/f_{V,i}(t_m)} = \mathrm{CR}_i(m,n).$$

Substitution into (177) proves (179). Projective invariance allows a common application of g_m^{-1}. The definitions of U, V, together with $g_n = g_m k_{m,n}$, then give the two lines in (180); balancing units and the factor $\Delta_{m,n}^{-1}$ disappear after projectivization. □

Proposition 8.17 (The central absorption has an exact visual formula)**.** *Apply a real projective coordinate in which $[E] = 0$, $[P] = \infty$, and the actual central geodesic is $z(t) = ie^t$. If*

$$\widehat{H}_P^{-1} k_{m,n} P = \begin{pmatrix} a \\ b \end{pmatrix}, \qquad b \neq 0,$$

put $v = \sigma_0(a/b)$. For the central logarithmic cross-ratio potential

$$u_v(t) = \log \left| \frac{z(t)}{z(t) - v} \right|$$

one has

$$u_v'(t) = \frac{v^2}{e^{2t} + v^2}, \tag{181}$$

$$\int_{t_m}^{t_n} (1 - u_v'(t))\, dt = \frac{1}{2} \log \frac{e^{2t_n} + v^2}{e^{2t_m} + v^2}. \tag{182}$$

In particular, for a sequence of intervals with $t_n - t_m \to \infty$, the central cross ratio can absorb all but $O(1)$ of its maximal growth only if

$$|\sigma_0(a/b)| \geq e^{t_n - O_P(1)}. \tag{183}$$

Equivalently, the pointed orbit line $[k_{m,n}P] = [a : b]$ is exponentially close to $[P] = [1 : 0]$ in the central visual metric.

Proof. Since $|ie^t| = e^t$ and $|ie^t - v| = (e^{2t} + v^2)^{1/2}$,

$$u_v(t) = t - \frac{1}{2} \log(e^{2t} + v^2).$$

Differentiation gives (181), and integration of its complement gives (182). If the latter is bounded while the interval length tends to infinity, its numerator and denominator must remain comparable; this forces $v^2 \gg e^{2t_n}$ up to a bounded factor, which is (183). □

Lemma 8.18 (What Schwarz–Pick alone says about the endpoint)**.** *For distinct $u, v \in \mathbf{P}^1(\mathbf{R})$, let*

$$\psi_{u,v}(z) = \log \left| \frac{z - u}{z - v} \right|$$

in an affine chart. Then

$$\|d\psi_{u,v}(z)\|_{\mathrm{hyp}} = \frac{\mathrm{Im}(z)|u - v|}{|z - u|\,|z - v|} \leq 1, \tag{184}$$

with equality exactly on the geodesic joining u and v. Hence, along the i-th period path,

$$\left|\frac{d}{dt}\log\left|\frac{f_{U,i}(t)}{f_{V,i}(t)}\right|\right| \le \|DF_i(\gamma(t))\|_{\mathrm{hyp}}. \tag{185}$$

For $i = 0$, the right side is 1.

Proof. The Euclidean gradient of $\psi_{u,v}$ has norm $|u-v|/(|z-u||z-v|)$. Multiplication by $\mathrm{Im}(z)$ converts it to the hyperbolic cotangent norm. The elementary inequality

$$\mathrm{Im}(z)|u-v| \le |z-u|\,|z-v|$$

is an equality exactly on the semicircle from u to v. Apply the chain rule to $Z_i(t) = F_i(\gamma(t))$. □

The lemma is deliberately recorded with its signs visible. It gives only

$$\frac{d}{dt}\mathfrak{B}_{m,n}(t) \le 2 + \|DF_1\|_{\mathrm{hyp}} + \|DF_2\|_{\mathrm{hyp}},$$

whereas the useful curvature is $2 - \|DF_1\|_{\mathrm{hyp}} - \|DF_2\|_{\mathrm{hyp}}$. Thus another application of Schwarz–Pick cannot prove (334): it loses exactly the two outer signs. Recovering those signs from the simultaneous Galois position of $[a : b]$ is the nontrivial thin-group input.

There is a useful exact way to isolate those signs. Put

$$L_0 = t_n - t_m, \qquad L_i = \int_{t_m}^{t_n} \|DF_i(\gamma(t))\|_{\mathrm{hyp}}\,dt \quad (i = 1, 2), \tag{186}$$

and set

$$X_i = \log|\,\mathrm{CR}_i(m,n)|, \qquad \alpha_i = L_i - X_i \quad (i = 0, 1, 2). \tag{187}$$

Here the definition of L_0 uses $F_0 = \mathrm{id}$. Lemma 8.18 gives $X_i \le |X_i| \le L_i$, hence $\alpha_i \ge 0$. The number α_i measures how much length of the i-th period path fails to appear as oriented displacement in the cross-ratio coordinate determined by the two algebraic endpoint lines.

Proposition 8.19 (Endpoint absorption is a competition of alignment defects)**.** *With the notation above,*

$$\mathfrak{B}_{m,n}(t_n) - \mathfrak{B}_{m,n}(t_m) = \mathfrak{C}(m,n) - 2\alpha_0 + \alpha_1 + \alpha_2. \tag{188}$$

Consequently Problem 17.16 is equivalent to the existence of $\eta_P > 0$ and C_P such that

$$2\alpha_0 - \alpha_1 - \alpha_2 \ge \eta_P\mathfrak{C}(m,n) - C_P \tag{189}$$

before cusp incidence or frame recurrence.

Proof. By (179), the left side of (188) is $2X_0 - X_1 - X_2$. On the other hand,

$$\mathfrak{C}(m,n) = 2L_0 - L_1 - L_2.$$

Substitute $X_i = L_i - \alpha_i$. Rearrangement proves both statements. □

The signed scalar inequality (189) is a sufficient route, but it is not the only way to keep the Hodge defect. There is a second exact calculation which avoids comparing the two outer alignment defects with the central one. It keeps the whole rational two-future frame instead of projecting that frame to the single boundary scalar $\mathfrak{B}_{m,n}$.

Put $Q = Q_m$, $U = R_m^{\mathrm{dyn}}$, and $V = \widetilde{R}_{m,n}$, and define the *controlled two-future potential*

$$\mathfrak{A}^{\mathrm{ctl}}_{m,n}(t) = 2\log\left|\frac{f_{V,0}(t)}{f_{Q,0}(t)}\right| - \sum_{i=1}^{2}\log\left|\frac{f_{V,i}(t)}{f_{Q,i}(t)}\right|. \tag{190}$$

It is the cyclic Hodge potential of the rational determinant-one pair $(Q_m, Q_n/\Delta_{m,n})$. In particular, its only varying finite denominator is the same ideal $(\Delta_{m,n})$ which already occurs in the two-frame Wronskian.

Proposition 8.20 (The Hodge defect has two exact channels)**.** *Assume $\Delta_{m,n} \neq 0$, and use the notation* (186)– (187)*. Then*

$$\mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_n) - \mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_m) \geq \mathfrak{C}(m,n) - 2X_0. \tag{191}$$

Consequently

$$\max\left\{X_0, \mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_n) - \mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_m)\right\} \geq \frac{1}{3}\mathfrak{C}(m,n). \tag{192}$$

There is also an exact visual interpretation of the first channel. Put

$$G_{m,n}(t) = \Lambda\big(Z_0(t); [Q_m], [Q_n]\big), \tag{193}$$

where Λ is the projectively invariant two-boundary Green function (213)*. Then*

$$X_0 = G_{m,n}(t_m) - G_{m,n}(t_n) \leq G_{m,n}(t_m). \tag{194}$$

Thus every stopped segment puts at least one third of its Hodge defect either into the decrease of the central Green proximity of the two rational future lines or into the cyclic Hodge growth of their controlled Wronskian frame.

Proof. Apply g_m^{-1} and use the central coordinate of Proposition 8.17. The three relevant lines are then

$$[Q] = [P] = \infty, \qquad [U] = [E] = 0, \qquad [V] = v,$$

and the actual geodesic is $z(t) = ie^t$, after translating the origin of t. Up to constants independent of t,

$$\log\left|\frac{f_{U,0}}{f_{Q,0}}\right| = t, \qquad \log\left|\frac{f_{U,0}}{f_{V,0}}\right| = u_v(t).$$

Subtracting gives

$$\left[\log\left|\frac{f_{V,0}}{f_{Q,0}}\right|\right]_{t_m}^{t_n} = L_0 - X_0 = \alpha_0.$$

For $i = 1, 2$, Lemma 8.18, now applied to the endpoint pair $([Q], [V])$, gives

$$\left|\left[\log\left|\frac{f_{V,i}}{f_{Q,i}}\right|\right]_{t_m}^{t_n}\right| \leq L_i.$$

Substitution in (190) gives

$$\mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_n) - \mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_m) \geq 2\alpha_0 - L_1 - L_2 = \mathfrak{C}(m,n) - 2X_0,$$

which proves (191). If $X_0 \geq \mathfrak{C}(m,n)/3$, the first entry in the maximum proves (192); otherwise (191) makes the second entry larger than $\mathfrak{C}(m,n)/3$.

Finally, in the same coordinate,

$$\Lambda(ie^t;\infty,v)=\frac{1}{2}\log\left(1+v^2e^{-2t}\right)=-u_v(t).$$

The invariance of Λ under real projective transformations and the definition of X_0 therefore give (194). Nonnegativity of Λ gives the final inequality. □

Lemma 8.21 (Green size and affine shape are the sum and difference)**.** *For $w\in\mathbb{H}$ and a real column $q=(q_1,q_2)^t$, use the determinant-one Hodge norm*

$$\|q\|_w=\frac{|q_1-q_2w|}{\sqrt{\operatorname{Im} w}}.$$

Write $f_q(w)=q_1-q_2w$. If $\det(q,v)=1$, then

$$\Lambda(w;[q],[v])=\log\|q\|_w+\log\|v\|_w, \tag{195}$$

$$\log\left|\frac{f_v(w)}{f_q(w)}\right|=\log\|v\|_w-\log\|q\|_w. \tag{196}$$

Consequently, if the two quantities on the right sides of (195) and (196) are denoted by G and r, respectively, then

$$\log\|q\|_w=\frac{G-r}{2},\qquad \log\|v\|_w=\frac{G+r}{2},\qquad \max\{\log\|q\|_w,\log\|v\|_w\}=\frac{G+|r|}{2}. \tag{197}$$

Proof. In an affine chart write the endpoints of q,v as x_q,x_v. Cancelling their second coordinates against $\det(q,v)$ gives

$$\frac{\|q\|_w\|v\|_w}{|\det(q,v)|}=\frac{|w-x_q|\,|w-x_v|}{\operatorname{Im}(w)|x_q-x_v|}.$$

This is $\exp\Lambda(w;x_q,x_v)$. The period evaluation is the same determinant pairing, so its quotient is the quotient of the two Hodge norms. The last three identities are elementary linear algebra. □

The lemma explains exactly, without terminology, what "rank-two fusion" means. At one real place, the Green coordinate and the affine Hodge coordinate already determine the sizes of both basis vectors. Over K, however, an Arakelov subline is one line $[q+sv]$ with $s\in K$. It cannot choose the larger of q and v independently at the three embeddings without paying the finite height of s. Ordinary adelic Roth estimates this simultaneous choice at the critical exponent; the desired strict gain must additionally use that $(q,v)=(Q_m,Q_n/\Delta_{m,n})$ comes from a consecutive quotient in the pointed thin orbit. Thus the local maximum in (197) is elementary; converting it to one global maximal-slope line is the deep step.

> **What has just been proved.** This proposition removes the apparent need to prove that the two outer paths cannot backtrack. Outer backtracking can destroy the scalar alignment inequality, but it cannot destroy both channels. If the line $[k_{m,n}P]$ follows $[P]$ closely enough to absorb the central Hodge growth, the central Green function records that visual recurrence. If it does not, the determinant-one two-future frame retains a definite part of the curvature. This is a geometric statement valid for the whole stopped segment; no Rosen generator calculation or arithmetic approximation theorem enters its proof.

What is not yet proved. The maximum in (192) is not yet an Arakelov height. The Green term and the controlled Hodge term are two coordinates of the same rank-two affine extension, and both use the same finite divisor $(\Delta_{m,n})$. Adding two independently normalized scalar heights would charge that divisor twice and repeat the critical multiplicity failure of Proposition 8.34. The remaining arithmetic task is therefore a stopped rank-two transport statement: assemble the two coordinates in the Galois completion of (203), charge $(\Delta_{m,n})$ once, and retain whichever of the two proved channels is large. Proposition 17.20 proves that the static slope of the extension is symmetric and cannot supply this assertion; Proposition 17.23 constructs the two cyclic tensor lines on which the stopped metric must act. Rank drop is $\Delta_{m,n} = 0$; vanishing of the complementary affine coordinate is $A_{m,n} = 0$. These are already the recurrence and cusp alternatives.

Formula (189) remains the shortest *scalar* form of the problem. The central defect α_0 is the explicit quantity in (182), while α_1, α_2 measure outer backtracking. A proof of that scalar inequality would still suffice. Proposition 8.20 gives a weaker and more structural alternative: retain the full rank-two extension and prove one all-place fusion theorem instead of controlling the signs of the two outer cross-ratios separately.

The hard statement, without jargon. The preceding proposition is elementary; it is not the missing theorem. It says exactly what a bad interval would have to do. The interval quotient must send P exponentially close back to P in the central visual metric. At the same time, the two outer cross ratios in (179) must align with their Hodge paths with the signs needed to absorb the outer Schwarz defect. The two coordinates a, b are conjugate algebraic numbers in one fixed fractional lattice, and they come from a consecutive quotient in G_7, not from an arbitrary matrix in $\mathrm{SL}_2(\mathcal{O}_K)$.

Thus the residual assertion is a simultaneous Galois visual-approximation theorem for one pointed thin orbit. Ordinary rational approximation sees only the central inequality (183); an ambient Arakelov height sees all three conjugates but allows every primitive pair. The desired proof must combine both pieces and use the consecutive thin-orbit restriction. The two exact failures are $a = 0$, which is the cusp incidence, and $b = 0$, which is Borel recurrence.

Proposition 8.22 (Affine equality is exactly the special alternative)**.** *Suppose an interval quotient fixes $[P]$. In a determinant-one frame (P, E), write*

$$k = \begin{pmatrix} d & \beta \\ 0 & d^{-1} \end{pmatrix}.$$

If $d^2 \neq 1$, the affine transport $s \mapsto d^2 s + d\beta$ has the unique fixed splitting

$$s_* = -\frac{d\beta}{d^2 - 1}, \tag{198}$$

and the line $[E + s_* P]$ is the other eigenline of k. Thus this neutral affine holonomy is precisely the special hyperbolic alternative. If $d^2 = 1$ and the projective element is nonidentity, then $\beta \neq 0$, there is no finite fixed splitting, and k is parabolic; this is the cusp alternative.

Proof. The fixed-point equation for the affine action (204) is $s = d^2 s + d\beta$, which gives (198). Directly,

$$k(E + sP) = d^{-1}\big(E + (d\beta + d^2 s)P\big),$$

so a fixed affine splitting is exactly a second projective eigenline. For $d^2 \neq 1$ the two eigenvalues are distinct. For $d^2 = 1$, a nonzero upper-right entry gives a nontrivial unipotent projective element. □

What the metric calculation now proves. There is no longer an unquantified path-comparison error. The strict Schwarz defect accumulates on the actual natural-extension path by (164). Passing to the height-reduced arithmetic frame changes the cyclic metric by the explicit quantity $\mathfrak{B}_n - \mathfrak{B}_m$, and (166) shows that this quantity is the Hodge norm of the affine shear s_n. Proposition 8.24 identifies its all-place height with the past-endpoint height, modulo the controlled future height.

The remaining endpoint theorem. The signed real number $\mathfrak{B}_n$ is not a proper height: a unit can make it large at one embedding and small at the other two. The proportional scalar gap (334) would suffice, but Proposition 8.20 proves that it is not necessary to control that signed scalar. The remaining argument may instead take the maximal Arakelov slope of the full Galois orbit of the tautological extension and select between its Green and controlled-Wronskian coordinates as in Problem 17.17. The path inequality and the two-channel alternative are proved; only this one-divisor all-place conversion remains.

Corollary 8.23 (Balancing removes the multiplicative part)**.** *Put*

$$z_n = \varepsilon_n^2 s_n, \qquad \xi_n = \varepsilon_n \varepsilon_{n+1} \beta_n. \tag{199}$$

Then $z_n \in \delta_P^{-1}\mathcal{O}_K$, $\xi_n \in \mathcal{O}_K$, *and*

$$z_{n+1} - z_n = \xi_n, \qquad z_n = z_m + \sum_{j=m}^{n-1} \xi_j. \tag{200}$$

Thus the linear unit character in the affine recurrence is a coboundary along the balanced ray. The remaining quantity is the height of an accumulated extension class in one fixed fractional lattice.

Proof. Multiply (159) by ε_{n+1}^2 and use $d_n = \varepsilon_n/\varepsilon_{n+1}$. This gives

$$\varepsilon_{n+1}^2 s_{n+1} = \varepsilon_n^2 s_n + \varepsilon_n \varepsilon_{n+1} \beta_n.$$

All three factors in ξ_n are integral and the first two are units; the fixed lattice $\delta_P^{-1}\mathcal{O}_K$ contains $\mathcal{O}_K$. Iteration proves the second formula. □

Equation (159) explains both the relevance and the limitation of the S-unit direction. After the additive terms have been confined to a fixed finite-rank multiplicative group, division by one nonzero term gives an S-unit equation and the Evertse–Schlickewei–Schmidt theorem supplies finiteness. Before that confinement, only d_n is a unit; s_n and β_n range through the fixed additive lattice $\delta_P^{-1}\mathcal{O}_K$. The affine equation itself therefore does not satisfy the hypotheses of an S-unit theorem. The stopped Hodge or holonomy estimate has a sharply stated job: bound the height of this affine extension, or force it onto the periodic unit eigendirection. Corollary 8.23 makes the first alternative still more precise: after balancing, one must bound the partial sums of the integral increments ξ_n, globally across all three embeddings. Estimating the increments separately cannot do this, because cancellation is part of the extension class.

Proposition 8.24 (The affine gauge is the height of the past endpoint)**.** *Let* $M_n = (Q_n, R_n^{\mathrm{ctl}}) \in \mathrm{SL}_2(\mathcal{O}_K)$ *be the controlled frame and let* s_n *be defined by* (158)*. Then*

$$h([s_n : 1]) = h([g_n E]) + O\big(h([g_n P]) + 1\big), \tag{201}$$

where the implied constant depends only on the fixed choice of controlled reduction. In particular, boundedness of

$$h([g_n P]) + h([s_n : 1]) \tag{202}$$

along infinitely many compact returns forces a projectively nonidentity element of $G_7 \cap B_P(K)$.

Proof. Equation (158) is the projective identity

$$[R_n^{\mathrm{dyn}}] = M_n[s_n : 1].$$

For any $M \in \mathrm{GL}_2(K)$, the elementary functorial height estimate, applied once to M and once to M^{-1}, gives

$$|h(Mx) - h(x)| \le h_{\mathrm{mat}}(M) + h_{\mathrm{mat}}(M^{-1}) + O_K(1).$$

Here $\det M_n = 1$, so the entries of M_n^{-1} are, up to signs, the entries of M_n. The controlled-completion lemma bounds their matrix height by $O(h([Q_n]) + 1)$. Since multiplication by the balancing unit does not change projective height,

$$h([Q_n]) = h([g_nP]), \qquad h([R_n^{\mathrm{dyn}}]) = h([g_nE]).$$

This proves (201).

If (202) is bounded, the same is true of the two projective endpoints $([g_nP], [g_nE])$. Northcott's theorem leaves only finitely many such pairs. Two distinct returns therefore give the same pair. Their quotient fixes $[P]$ and $[E]$; choosing two returns whose quotient is not projectively the identity gives a nonidentity element of $G_7 \cap B_P(K)$. □

Proposition 8.24 makes the remaining endpoint cost intrinsic. The active stopped estimate need not control a chosen matrix entry. It must control the proper height of the ordered future–past pair, or show that the affine recurrence lies on a unit eigendirection. The first alternative closes by Northcott; the second is the special hyperbolic alternative when the diagonal unit is nontrivial, with its unipotent degeneration giving the cusp alternative.

Proposition 8.25 (The affine cocycle is the universal tautological extension)**.** *Let V be the standard rank-two $\mathcal{O}_K$-local system and let $\mathscr{L} \subset V$ be the tautological line on the projective flat bundle $\mathbf{P}(V)$; its fibre over a line $[L]$ is simply L itself. There is a canonical exact sequence*

$$0 \longrightarrow \mathscr{L}^{\otimes 2} \longrightarrow \mathscr{L} \otimes V \longrightarrow \det V \simeq \mathcal{O} \longrightarrow 0. \tag{203}$$

On the fibre over $[P]$, a determinant-one completion of P is a splitting of the last arrow. If

$$b(d, \beta) = \begin{pmatrix} d & \beta \\ 0 & d^{-1} \end{pmatrix} \in \mathscr{B}_P,$$

then its action on the affine splitting coordinate is the linear action

$$\rho_{\mathrm{aff}}(b) = d\, b = \begin{pmatrix} d^2 & d\beta \\ 0 & 1 \end{pmatrix}. \tag{204}$$

In particular, (159) is precisely parallel transport in (203), written in the controlled splittings.

Proof. Tensor the tautological sequence

$$0 \longrightarrow \mathscr{L} \longrightarrow V \longrightarrow V/\mathscr{L} \longrightarrow 0$$

by $\mathscr{L}$. Since V has determinant one, $\mathscr{L} \otimes (V/\mathscr{L}) \simeq \det V \simeq \mathcal{O}$, giving (203).

In a frame (P, E), the diagonal character of the tautological line is d. Therefore the action on $\mathscr{L} \otimes V$ is the standard matrix b multiplied by d, which is (204). It is a representation, not just a formula along one sequence: if $b_j = b(d_j, \beta_j)$, direct multiplication gives

$$\rho_{\mathrm{aff}}(b_1 b_2) = \rho_{\mathrm{aff}}(b_1)\rho_{\mathrm{aff}}(b_2).$$

Finally,

$$\rho_{\mathrm{aff}}(b(d_n,\beta_n))\begin{pmatrix} s_n \\ 1 \end{pmatrix} = \begin{pmatrix} d_n^2 s_n + d_n\beta_n \\ 1 \end{pmatrix},$$

which is the recurrence already proved. □

Corollary 8.26 (The transverse Wronskian differentiates the affine coordinate)**.** *The character d^2 on the subline $\mathscr{L}^{\otimes 2}$ in (203) is the linear part of the affine transport (204). If two flat covectors give local period functions f_m, f_n, then the logarithmic derivative of their projective coordinate is*

$$d\log(f_n/f_m) = \frac{f_m f_n' - f_n f_m'}{f_m f_n}. \tag{205}$$

Thus the normalized transverse jet (299) is the logarithmic cotangent of the same tautological affine extension; it is not a separate auxiliary construction.

Proof. Two splittings of (203) differ by a section of $\mathscr{L}^{\otimes 2}$, and (204) shows that its transition character is d^2. Formula (205) is the quotient rule. Its numerator is the Gauss–Manin Wronskian and its denominator removes the two incidence divisors, exactly as in (299). □

This proposition removes a construction problem from the list. The affine extension is neither an invented normal function nor a collection of local shear parameters; it is the canonical bundle $\mathscr{L} \otimes V$ on the projective flat bundle, with its integral lattice. Corollary 8.26 also identifies its logarithmic cotangent with the already constructed Wronskian. The remaining Hodge-theoretic task is now precise: equip the Galois completion of this extension with the fundamental curved metric coming from $\mathcal{N}_2$, and prove that its stopped Arakelov degree is (333). The algebraic extension and its integral affine coordinate no longer have to be guessed, and neither does its transverse first jet.

> **What is elementary, and what is not.** The balanced completion, the comparison (147), and the transition identity (149) are reduction theory and matrix algebra. Their role is diagnostic: they identify the variable which has to recur and prove that its size is the same height already measured by the Asai norm. The substantive theorem would be the dynamical assertion that this particular thin-group return cocycle cannot visit infinitely many such cosets with $\mathfrak{H}_P = o(t)$ while remaining in the double-horoball-mismatch channel. None of the displayed algebraic identities proves that assertion.

Problem 8.27 (Sublinear rational-Borel intersection)**.** *Let $P \in \mathcal{O}_K^2$ be primitive and let g_n be reductions of the geodesic ray to $[P]$ at compact return times $t_n \to \infty$. Suppose that*

$$\mathfrak{H}_P(g_n \mathscr{B}_P) = o(t_n).$$

Prove that $G_7 \cap B_P(K)$ contains a projectively nonidentity element.

Propositions 8.7 and 8.8 prove that this problem is exactly the group-theoretic form of the exponent-one drifting equality case. It is not an instance of ambient split-torus rigidity. The factors b_n all lie in one rational Borel, but the errors C_n vary, and the lifted Teichmüller curve is not a torus orbit. A proof must use that the left factors come from one geodesic prefix sequence in the thin group, together with the two modular embeddings. For example, a suitable relative-Morse theorem would say that sublinear tracking of an ambient rational Borel by this Hodge curve forces an actual intersection. No such conclusion follows from sublinear tracking for an arbitrary subgroup of the Hilbert modular group.

Proposition 8.9 also prevents a more subtle overclaim. The logarithmic unit vector records only the diagonal, or Levi, entry $\varepsilon_n/\varepsilon_{n+1}$ in (150). The lifted equation also contains the gauge-dependent additive entry β_n, so its diagonal part cannot determine the invariant coset $C_n\mathscr{B}_P$. A limiting direction for the units therefore does not make those cosets repeat. The actual equality theorem must lock the complete frame coset. Unit-direction locking describes the allowed neutral holonomy after that frame has recurred; it is not by itself the recurrence theorem.

Proposition 8.28 (The Morse audit for the frame problem)**.** *Neither of the two natural higher-rank representations puts Problem 8.27 under a standard Morse lemma.*

1. *The Galois-diagonal representation in* $\mathrm{SL}_2(\mathbb{R})^3$ *is not uniformly regular for the full product Borel.*
2. *The block representation*
$$\rho_\star = \mathrm{Sym}^2\rho_0 \oplus (\rho_1 \otimes \rho_2)$$
is not projectively divergent on the peripheral sequence T^m*.*

Proof. By Proposition 7.3, there are distinct $g_k \in G_7$ whose two outer images remain in one compact set. The distinguished image must escape every compact set because it is discrete. The two outer simple-root gaps therefore remain bounded along an escaping sequence, proving the first assertion.

At every real embedding, the largest singular value of the shear $\rho_i(T^m)$ is $\asymp |m|$. Hence the largest singular values of both $\mathrm{Sym}^2\rho_0(T^m)$ and $(\rho_1\otimes\rho_2)(T^m)$ are $\asymp m^2$. The ratio between the largest and second-largest singular values of their direct sum stays bounded. Thus $\rho_\star(T^m)$ is not projectively divergent. □

The failure can be stated more sharply in the language used by the relative-Anosov literature. Write $a_i(g) = \log s_1(\sigma_i(g))$, where s_1 is the largest singular value in the Hodge-adapted norm.

Proposition 8.29 (Exact root switching in the block representation)**.** *For*
$$\rho_\star = \mathrm{Sym}^2\rho_0 \oplus (\rho_1 \otimes \rho_2),$$
the three largest logarithmic singular values, up to the bounded change between adapted norms at the ends of a compact return, are
$$2a_0, \qquad a_1 + a_2, \qquad |a_1 - a_2|.$$
Consequently its first two simple singular gaps are
$$\alpha_1(\rho_\star(g)) = 2a_0 - a_1 - a_2, \qquad \alpha_2(\rho_\star(g)) = 2\min(a_1, a_2). \tag{206}$$
Neither gap is divergent on every escaping sequence in G_7*:*

1. *for* $g = T^N$*,* $\alpha_1 = O(1)$ *and* $\alpha_2 = 2\log|N| + O(1)$*;*
2. *there are escaping* g_n *for which both outer images remain in a compact set, and then* $\alpha_1 \to \infty$ *while* $\alpha_2 = O(1)$*.*

Thus $\rho_\star$ *is neither relatively* P_1*-Anosov nor relatively* P_2*-Anosov. The quantity available uniformly is the adaptive one*
$$\max\{\alpha_1, \alpha_2\},$$
or equivalently the top-to-third gap $\alpha_1 + \alpha_2 = 2a_0 - |a_1 - a_2|$*, not a fixed simple-root gap.*

Proof. The singular logarithms of $\operatorname{Sym}^2 \rho_0(g)$ are $2a_0, 0, -2a_0$. Those of $\rho_1(g) \otimes \rho_2(g)$ are

$$a_1 + a_2, \quad |a_1 - a_2|, \quad -|a_1 - a_2|, \quad -a_1 - a_2.$$

The doubled-central domination places $2a_0$ first along the relevant returns, giving (206). Every conjugate of T^N is a shear of size $\asymp |N|$, proving the first asymptotic. The second is Proposition 7.3. Relative P_k-Anosov representations are P_k-divergent on every escaping sequence, including peripheral ones; hence the two countersequences rule out $k = 1, 2$ [10, 11]. □

This proposition explains exactly what relative-Anosov theory contributes. It provides the correct language for the cusp-adapted norms and for each individual contraction. It does not choose between two different flag types as the orbit alternates between the two walls. Problem 17.17 is precisely the arithmetic assertion that the one-divisor affine extension makes that adaptive maximum available without paying for its determinant twice. Replacing it by the phrase "relative Morse lemma" would therefore conceal the same missing selection.

The first representation remembers all three rational flags but lacks uniform regularity. The second has the correct doubled-central character on thick segments but loses its singular gap on a cusp excursion. A relative-Anosov theorem can organize the already proved contraction; it cannot infer recurrence of the arithmetic coset $C_n \mathscr{B}_P$ from these hypotheses. This is why the required input is a stopped arithmetic frame theorem rather than an unqualified Morse statement.

The number-theoretic black boxes fail at the same, now visible, place.

Proposition 8.30 (The S-unit and moving-target audit)**.** *The exact frame equation does not satisfy the hypotheses needed for either of the two natural applications of the Subspace Theorem.*

1. *The Evertse–Schlickewei–Schmidt theorem for a linear equation in a fixed finite-rank multiplicative group does not apply to (149). Only the diagonal entry of u_n lies in the unit group; β_n and the entries of C_n vary additively in $\mathcal{O}_K$.*
2. *The relation $g_n \in C_n B_P$ is one exact slowly moving hyperplane incidence, but it makes the sequence degenerate over the field generated by that moving target. The nondegenerate conclusion of the Ru–Vojta moving-target theorem therefore cannot make $C_n B_P$ independent of n.*

Proof. The first assertion follows directly from (150): there is no closed linear equation whose variables all belong to one fixed multiplicative group. Such an equation appears only after the remainder support has already been bounded, which is why the S-unit theorem is a valid closing argument but not the source of boundedness [18].

For the second assertion, represent a matrix X by its four entries. The condition that XP be proportional to $C_n P$ is

$$L_n(X) = \det(XP, C_n P) = 0. \tag{207}$$

This is linear in the entries of X, and its coefficient height is $O(\mathfrak{h}_\Delta(C_n) + 1) = o(t_n)$. Thus it is indeed a slowly moving target. But $L_n(g_n) = 0$ identically, so the points are already linearly degenerate over the target coefficient field. A degenerate alternative returns the known containment in a moving hyperplane; it does not prove that two of those hyperplanes agree [23]. □

Taking all the coordinates of the Asai state does not repair this degeneracy.

Proposition 8.31 (All state-generated Asai incidences give the same moving Borel)**.** *For $C \in \mathrm{SL}_2(K)$, let*

$$I_C = \{\ell \in W^\vee : \ell(\mathrm{As}(C) Y(P)) = 0\}.$$

Then for every $A \in \mathrm{SL}_2(K)$,

$$\ell(\mathrm{As}(A)Y(P)) = 0 \quad (\ell \in I_C) \quad \Longleftrightarrow \quad A \in CB_P(K). \tag{208}$$

A basis of I_C has coefficient height $O(\mathfrak{h}_\Delta(C) + 1)$. Consequently, along the sublinear frame factorization, even the full family of Asai hyperplanes is slowly moving but is identically degenerate on the same moving Borel coset.

Proof. The simultaneous vanishing on the left says that $\mathrm{As}(A)Y(P)$ lies in the one-dimensional annihilator of I_C, namely the line spanned by $\mathrm{As}(C)Y(P)$. Equality of these decomposable tensor lines implies equality of their distinguished factor lines. Thus $[AP] = [CP]$, which is equivalent to $C^{-1}A \in B_P(K)$. The converse is immediate. Finally, the Plücker height of the annihilator of a line is its projective height up to a fixed error. Since the Asai representation has fixed degree, this height is $O(\mathfrak{h}_\Delta(C) + 1)$. □

There is, however, a genuinely transverse arithmetic incidence once two different frames are compared.

Proposition 8.32 (The two-frame transverse minor)**.** *For two reduced returns put*

$$\Delta_{m,n} = \det(Q_m, Q_n) \in \mathcal{O}_K, \qquad m < n. \tag{209}$$

Its vanishing and the absolute norm of its principal ideal are independent of all balancing and completion choices. Moreover,

$$\Delta_{m,n} = 0 \quad \Longleftrightarrow \quad C_m\mathscr{B}_P = C_n\mathscr{B}_P, \tag{210}$$

and, if $\Delta_{m,n} \neq 0$,

$$0 \leq \log|\,\mathrm{N}\,\Delta_{m,n}| \leq 3h([g_mP]) + 3h([g_nP]) + 3\log 2. \tag{211}$$

If $M_{k,i}$ is the max norm of $\sigma_i(Q_k)$, define

$$\mathcal{P}_{m,n} = \sum_{i=0}^{2} \log \frac{M_{m,i}M_{n,i}}{|\sigma_i(\Delta_{m,n})|}.$$

Then, whenever the minor is nonzero,

$$\mathcal{P}_{m,n} = 3h([g_mP]) + 3h([g_nP]) - \log|\,\mathrm{N}\,\Delta_{m,n}|. \tag{212}$$

Thus any strict lower bound exceeding the two state heights forces frame recurrence.

Proof. Changing a balancing unit multiplies Q_m or Q_n by a unit, and changing a second column leaves both first columns fixed. Hence vanishing and absolute ideal norm are invariant. Equation (210) follows from Proposition 8.9. For each embedding,

$$|\sigma_i(\Delta_{m,n})| \leq 2M_{m,i}M_{n,i}.$$

Multiply over the three embeddings and use $\sum_i \log M_{k,i} = 3h([g_kP])$. A nonzero algebraic integer has positive integral absolute norm, which proves (211). The same product formula, without the factor-two estimate, gives (212). □

Proposition 8.33 (The transverse proximity is a geodesic Green function)**.** *For $w \in \mathbb{H}$ and distinct $a, b \in \mathbb{R}$, put*

$$\Lambda(w; a, b) = \log \frac{|w-a|\,|w-b|}{\operatorname{Im}(w)\,|a-b|}. \tag{213}$$

Then $\Lambda \geq 0$, with equality exactly when w lies on the geodesic with endpoints a, b. Fix a compact set in each of the three period coordinates and choose w_i in it. If $x_{r,i} = [\sigma_i(Q_r)]$, then, in compatible affine charts,

$$\log \frac{M_{m,i} M_{n,i}}{|\sigma_i(\Delta_{m,n})|} = \Lambda(w_i; x_{m,i}, x_{n,i}) + O(1), \tag{214}$$

where the error is uniform in the two rational frames. Consequently $\mathcal{P}_{m,n}$ is, up to $O(1)$, the sum of the three geodesic Green functions.

Proof. The inequality

$$|w-a|\,|w-b| \geq \operatorname{Im}(w)|a-b|$$

is the elementary equation of the semicircle through a, b; equality holds on that semicircle. Write $Q_r = (p_r, q_r)^t$ and $x_r = p_r/q_r$. Then

$$M_r = |q_r| \max\{1, |x_r|\}, \qquad |\det(Q_m, Q_n)| = |q_m q_n|\,|x_m - x_n|.$$

For w in a fixed compact subset of $\mathbb{H}$, the quotient $|w-x|/\max\{1,|x|\}$ is bounded above and below uniformly for $x \in \mathbb{R}$. Substitution proves (214). The chart containing infinity gives the same statement by projective invariance. □

This formulation exposes the double-horoball mismatch without shorthand. Schwarz–Pick directly controls the image under F_i of a central geodesic. The outer terms in (214), however, use the geodesic whose endpoints are the rational conjugates $x_{m,i}, x_{n,i}$. These are not known to be the boundary images of the two central endpoints; indeed Proposition 17.72 shows that exact rational boundary matching already forces the cusp alternative. The stopped Wronskian estimate must compare these two geodesics through the integral first-jet lattice. Merely applying Schwarz–Pick to the Hodge image silently substitutes the wrong outer endpoints.

This minor is the first algebraic quantity in the frame analysis which is not an equation of a single moving Borel. It also shows exactly what the Hodge jet must accomplish: create an overcritical lower bound for the two-frame proximity $\mathcal{P}_{m,n}$. Ordinary height theory gives only the opposite, critical inequality in (212). The modular embeddings can potentially supply the excess by comparing two returns of the Hodge flag, but the double-horoball mismatch is precisely the obstruction to replacing their outer limiting flags by the rational conjugate lines of Q_m and Q_n. Proposition 17.4 below constructs the relevant first jet on the generic fibre. Thus the remaining problem is not to guess a transverse covariant, but to prove the stopped integral norm estimate for that explicit covariant.

Nor does assigning extra weight to the central incidence create that surplus.

Proposition 8.34 (Weighted moving targets charge repeated support)**.** *Let $H = (L = 0)$ be a hyperplane and let $D = mH$ be the degree-m hypersurface defined by L^m. At every place and away from H,*

$$\frac{1}{m}\lambda_D(x) = \lambda_H(x) + O(1). \tag{215}$$

In the weighted moving-target Subspace Theorem, if hyperplanes or hypersurfaces of total weight c have the same support H, then the concentration term $\max_W \alpha(W)/\operatorname{codim} W$ is at least c, by taking $W = H$. Thus repeating the central incidence, or replacing it by a power, multiplies the proximity and the critical coefficient by the same amount. It gives no strict inequality.

Proof. The local Weil function of L^m is m times that of L, up to the bounded normalization of coefficient norms, proving (215). For targets with common support, the definition of the weighted concentration function gives $\alpha(H) = c$ and $\operatorname{codim} H = 1$. This is precisely the factor which appears on the right side of the weighted theorem [22]. □

These propositions narrow the viable synthesis. The extra incidence cannot be another coordinate, tensor power, or weighted copy of the current rational state: all such equations remain in the ideal of the moving Borel, and central multiplicity pays for itself. What is needed is a *Hodge-transverse arithmetic jet*: a covariant coupling two frames, transverse to the moving Borel, whose stopped proximity receives the positive divisor P_7. The Wronskian identity in Proposition 17.4 constructs precisely this covariant and identifies its zero set. The unresolved checks are its word-uniform integral lattice and the stopped proximity inequality. Only after those metric and integral checks have broken the state-generated degeneracy can a moving-target or arithmetic-holonomy theorem serve as the final number-theoretic engine.

Proposition 8.35 (Asai repetition and neutral unit holonomy)**.** *Let $P \in \mathcal{O}_K^2$ be primitive.*

1. *If distinct $g_m, g_n \in G_7$ satisfy*

$$[Y(g_m P)] = [Y(g_n P)] \quad \text{in } \mathbf{P}(W),$$

then $g_n^{-1} g_m$ fixes $[P]$. If the quotient is projectively nonidentity, $[P]$ is a cusp or a special hyperbolic fixed point.

2. *Conversely, if a projectively nonidentity $h \in G_7$ fixes $[P]$, write $hP = \varepsilon P$. Then $\varepsilon \in \mathcal{O}_K^\times$ and*

$$\operatorname{As}(h) Y(P) = \mathrm{N}_{K/\mathbf{Q}}(\varepsilon) Y(P) = \pm Y(P). \tag{216}$$

Thus the special alternative is exactly neutral unit holonomy on an integral twisted-Segre line.

Proof. Equality of two decomposable tensors implies equality of their three factor lines. In the distinguished factor this says $g_m P = c g_n P$ for some $c \in K^\times$. Hence $g_n^{-1} g_m P = cP$, and Proposition 7.1 gives the first assertion.

For the converse, ε is a root of the monic polynomial $X^2 - \operatorname{Tr}(h) X + 1 \in \mathcal{O}_K[X]$, and the same is true of its inverse. Therefore it is a unit. Tensoring the three conjugate eigenvalue equations gives (216); a unit has rational norm ± 1. □

For the elementary fixed point $\alpha = \lambda^2 - 1$, the matrix in the introduction has eigenvalue

$$u_7 = \lambda\alpha + 1 = \lambda^2 + \lambda.$$

Its other two conjugates lie between zero and one and $\mathrm{N}(u_7) = 1$. Consequently, along the corresponding closed ray, central contraction is exactly balanced by the two outer expansions. In return orientation the identity is

$$-\log u_{7,0} - \log u_{7,1} - \log u_{7,2} = 0, \tag{217}$$

where the last two terms are positive because $0 < u_{7,i} < 1$. This is an exact equality case of the integral Asai obstruction, not a failure of its estimate.

For Liouville-almost every ray, the two outer top exponents are, with the embedding labels fixed above,

$$\lambda_1^{\mathrm{H}} = \frac{1}{5}, \qquad \lambda_2^{\mathrm{H}} = \frac{3}{5}.$$

Every fixed outer vector has upper growth at most the corresponding top exponent at an Oseledets-regular point. Hence a regular rational endpoint would give

$$\limsup \frac{1}{t} \log \|Y(P)\|_t \leq -1 + \frac{1}{5} + \frac{3}{5} = -\frac{1}{5},$$

contradicting the proposition. The numerical value $-1/5$ is the ordinary integral Asai gap; unlike the pairwise determinant, this argument does not discard the integral lattice.

> **Where the mathematics comes from.** The lower bound 1 is algebraic number theory: it is the product formula and discreteness of one integral tensor lattice. The value 4/5 is Hodge theory: it is the sum of the two Lyapunov exponents of the conjugate variations. The assertion that an algebraic ray cannot remain in the resulting 1/5-gap unless it closes is neither ingredient alone. It is the precise hybrid rigidity theorem still required.

This yields a clean conditional completion which is genuinely about the full group.

Conjecture 8.36 (Trace-field Hodge-gap rigidity)**.** *If $x \in \mathbf{P}^1(K)$ is neither a cusp nor fixed by a hyperbolic element of G_7, then, along its returns to some fixed compact part of the triangle surface,*

$$\limsup_{t\to\infty} \frac{\log \|\sigma_1(P)\|_{1,t} + \log \|\sigma_2(P)\|_{2,t}}{t} < 1. \tag{218}$$

Theorem 8.37 (Conditional global classification)**.** *Conjecture 8.36 implies Conjecture 1.1.*

Proof. Let $x \in \mathbf{P}^1(K)$. If it is a cusp there is nothing to prove. Otherwise Proposition 8.1 contradicts (218). Thus a hyperbolic element fixes x. Its eigenvalue is in K, by Proposition 7.1, so it is special. □

Conjecture 8.36 asks only for a strict inequality; it does not require the typical values 1/5 and 3/5, digit frequencies, or convergence of every Lyapunov average. There is a familiar, stronger conjecture from continued-fraction dynamics which implies it. Let μ_{BM} be the Bowen–Margulis probability measure for the geodesic flow of the $(2, 7, \infty)$ orbifold, represented on the natural extension of the accelerated Rosen map.

Conjecture 8.38 (Algebraic Rosen genericity)**.** *If $x \in \mathbf{P}^1(K)$ is neither terminating nor periodic under the Rosen algorithm, then its natural-extension orbit is generic for μ_{BM}.*

Corollary 8.39 (A dynamical conditional completion)**.** *Conjecture 8.38 implies the heptagonal cusp classification.*

Proof. The two outer Hodge cocycles are integrable for the accelerated Rosen natural extension. Genericity and the multiplicative ergodic theorem, with cocycle sums divided by the accumulated hyperbolic roof time, therefore give their top exponents 1/5 and 3/5. A fixed outer vector has upper exponent no larger than the corresponding top exponent, so the Asai tensor attached to x has upper exponent at most $-1 + 1/5 + 3/5 = -1/5$. At compact returns this contradicts Proposition 8.1. Hence the orbit terminates or is periodic. Termination gives a cusp. A period gives a projectively nonidentity stabilizer, and Proposition 7.1 identifies it as the special hyperbolic alternative. □

Algebraic Rosen genericity is the heptagonal analogue of the usual expectation that a non-quadratic algebraic number is generic for the Gauss continued-fraction map. The full classification needs much less: it is enough to prove that the sum of the two outer upper exponents is strictly less than one on every nonperiodic trace-field ray. The difference is important. The conditional

proof uses the invariant distribution and the two Hodge cocycles of the whole thin group; it does not require digit frequencies one by one, and it does not assume full statistical normality when the weaker exponent inequality is available.

This theorem also locates the limitation of Roth's theorem. Roth controls the depth of each exceptionally good cusp approximation relative to the height accumulated before that excursion. Oseledets regularity controls a time average of the two outer Hodge cocycles. A sequence can have every individual cusp excursion sublinear in the elapsed time while the total proportion of time spent in successively longer cusps tends to one. Therefore even an optimal fixed-ε Roth theorem does not imply Conjecture 8.36. The required extra statement is an arithmetic rigidity or equidistribution theorem for the full thin-group orbit.

The finite-alphabet presentation of the group shows why the known "stammering" refinements of Roth do not automatically supply that extra statement. A Rosen digit a_j is a run of $|a_j|$ parabolic letters, but its hyperbolic cusp time is only $\log|a_j| + O(1)$. For the abstract model

$$|a_j| = \left\lceil e^{\sqrt{j}} \right\rceil$$

one has

$$\sum_{j\le N} \log|a_j| \asymp N^{3/2}, \qquad \frac{\log|a_N|}{\sum_{j<N}\log|a_j|} \longrightarrow 0.$$

Thus every single excursion is compatible with Roth, although cusp time has density one. On the other hand

$$\frac{|a_N|}{\sum_{j<N}|a_j|} \asymp N^{-1/2} \longrightarrow 0,$$

so no parabolic run repeats a block occupying a fixed proportion of the preceding finite-alphabet word. The stammering criteria of Adamczewski–Bugeaud therefore do not apply, even if one grants their natural Rosen analogue. This example is not proposed as an algebraic endpoint; it proves the logical point that Roth plus existing repetition criteria still leave exactly the cusp-density regime detected by the integral Asai obstruction [4].

There is nevertheless a useful stronger conclusion for the part of the problem which really does stammer. Theorem 18.6 below combines the same periodic surrogate with whole-group trace dominance and proves an exact termination-or-periodicity criterion for trace-field endpoints. At $q = 7$ its threshold is $3/2$, rather than the $9/2$ furnished by the general algebraic-approximation theorem. The model above still evades that sharper criterion, so this is a genuine partial theorem and not a completion disguised as a word-complexity statement.

> **What this reduction accomplishes.** The integral Asai tensor, its compact-return lattice lower bound, and the three Hodge exponents belong to the full monodromy representation. The only conditional sentence is the classification of its exceptional trace-field rays. Generator computations play no role in the proof of Theorem 8.37.

There is a nearby theorem which must not be over-applied. Filip's integral-vector theorem for a log-Anosov representation says that an integral vector in its strictly decreasing boundary subbundle forces a parabolic endpoint [20, Theorem 4.5.3]. Its proof is short and important: at returns to a compact set the vector would converge to zero, while the nonzero vectors in one fixed integral lattice have a uniform positive lower bound.

The Galois product representation is log-Anosov for the distinguished root, since its first projection is uniformizing. In the Asai tensor, however, the weights above the central contracting line are

$$(-1,+1,+1), \quad (-1,+1,-1), \quad (-1,-1,+1), \quad (-1,-1,-1).$$

They are incomparable with zero in the distinguished-root partial order: the two outer coordinates have not vanished. Thus the integral endpoint vector does not belong to the theorem's decreasing subbundle. This is a statement about the Levi factor of the full Galois representation, not a defect in the estimate used in Filip's proof.

Proposition 8.40 (Exact scope of the integral-vector theorem)**.** *Let $H = \mathrm{Res}_{K/\mathbf{Q}}\,\mathrm{SL}_2$, and give the Galois-diagonal representation of G_7 the log-Anosov structure coming from the distinguished simple root.*

1. *Filip's decreasing subbundle detects strict contraction in the distinguished factor only after the weights in the two outer Levi factors have been fixed.*
2. *Neither the integral Asai endpoint vector $Y(P)$, nor either pairwise vector $w_i(P)$, is a rational vector contained in that decreasing subbundle.*
3. *No extension which declares the normalized endpoint state strictly decreasing can prove the desired classification: at every special endpoint the normalized state has neutral unit holonomy.*

Proof. For the chosen root, two weights can be compared only if their two outer Levi weights agree. Each displayed Asai weight has a nonzero outer weight, so it is incomparable with the zero weight. The same calculation for $V_0 \otimes V_i$ gives the two weights $(-1,+1)$ and $(-1,-1)$, again with a nonzero outer Levi coordinate. After restriction of scalars a rational vector contains the full cyclic Galois orbit of these components, so it cannot be placed in just the distinguished decreasing component.

For the last assertion, suppose $hP = \varepsilon P$ at a special endpoint. Proposition 8.35 gives

$$\mathrm{As}(h)Y(P) = \mathrm{N}(\varepsilon)Y(P) = \pm Y(P).$$

The normalized state tensor introduced in (401) has the same property. Its norm can therefore have zero period around this closed geodesic. A theorem which made it strictly decreasing would incorrectly classify this permitted special point as a cusp. □

Proposition 8.3 separates two statements which are easy to conflate. Once the norm of the Galois-completed integral vector is bounded at infinitely many reduced returns, repetition is an elementary lattice argument; no extension of Filip's discreteness argument is needed. What Filip's theorem supplies in its range is the prior, much deeper assertion that an adapted Hodge metric makes the vector strictly decrease. What is missing here is the corresponding prior assertion at degree zero: the first-record condition must keep the compensating Levi coordinate in a bounded set, or put it on one periodic unit line. Calling the desired result an "equality-case extension of Filip" refers to this boundedness theorem, not to the final pigeonhole step.

Ordinary semistability cannot give that boundedness. In the fundamental Galois orbit the three degrees are

$$-\frac{1}{14}, \qquad \frac{1}{7}, \qquad -\frac{1}{14},$$

so their sum is zero. Semistability records the sum; a lattice height records the positive part. The positive middle component may therefore drift to infinity while the total degree remains zero. The thin-group input has to rule out precisely this drift for the sequence of rational Borels selected by the Rosen stopping rule.

The same obstruction is visible numerically on a parabolic power T^n. The three top singular-value logarithms satisfy

$$a_i(n) = \log n + \log|\lambda_i| + o(1).$$

The most dangerous ordinary integral Asai endpoint component has scalar weight

$$-a_0 + a_1 + a_2 = \log n + O_7(1) > 0,$$

so it is not in that decreasing subbundle. The pairwise weight has

$$-2a_0 + a_1 + a_2 = -3\log \lambda_0 + o(1) < 0,$$

but this is only a constant while the cusp distance is $2\log n + O(1)$, and its selected Hodge line is not a flat integral subrepresentation. Thus the existing log-Anosov theorem proves the compact-return step once the correct integral decreasing object exists; it does not construct that object here.

Filip's refined poset formalism does recover one additional piece which is useful here. For a uniformizing rank-two system together with one Galois-conjugate Hodge system, strict Schwarz–Pick contraction refines the root order to

$$E_0^+ \succ E_i \succ E_0^-.$$

It follows that the distinguished real component of $w_i(P) = P \otimes \tau^i(P)$ has nonpositive weight $-a_0 + a_i$, with strict decrease across compact returns. This is the representation-theoretic source of Proposition 17.46; it is a theorem about the full monodromy representation, not about its generators.

The rationality hypothesis still fails after descent. The restriction of scalars of the pairwise representation has three cyclic real components. A rational vector has all three conjugates, while the decreasing filtration contains only the distinguished component. The degree table (419) records the result: every negative component is accompanied by positive conjugate degree and the Galois norm has degree zero. Thus no nonzero rational vector lies wholly in the decreasing subspace to which the integral-vector theorem applies.

This is also the archimedean meaning of the closed-Bruhat channel in (642). Its finite vector is Borel-fixed; its distinguished real component decreases; and its two conjugates can compensate through the reciprocal incidence ideal. The required theorem is therefore an equality-case, adelic extension of the integral-vector theorem: at first norm-record returns, the compensation must either become strictly decreasing, giving a parabolic endpoint, or remain on one recurrent zero-weight projective line, giving the special hyperbolic stabilizer. A stronger pointwise domination estimate would lose the second alternative. The normalized state tensor in Proposition 17.47 makes this target completely explicit; Section 17.3 states the result which has to be proved.

9 The Rosen reduction

9.1 The exact algorithm

Let

$$I = [-\lambda_0/2, \lambda_0/2).$$

For $x \in I \setminus \{0\}$, choose the integer $n(x) \neq 0$ such that

$$\mathcal{R}(x) = -\frac{1}{x} - n(x)\lambda_0 \in I. \tag{219}$$

At a half-open endpoint we use the displayed convention. This is the nearest-λ-multiple continued fraction. The branch with digit n is represented over K by

$$B(n) = \begin{pmatrix} -n\lambda & -1 \\ 1 & 0 \end{pmatrix} \in G_7. \tag{220}$$

The determinant is one; a digit can be arbitrarily large because it records the duration of one excursion around the cusp.

Choose a primitive integral representative $P = (r_0, r_{-1})^t \in \mathcal{O}_K^2$ of the initial point. Exact iteration of (219) gives

$$\begin{pmatrix} r_{j+1} \\ r_j \end{pmatrix} = B(n_j) \begin{pmatrix} r_j \\ r_{j-1} \end{pmatrix}, \qquad r_{j+1} = -r_{j-1} - n_j \lambda r_j. \tag{221}$$

Every branch is unimodular, so the ideal (r_j, r_{j-1}) remains equal to $\mathcal{O}_K$. In particular, adjacent remainders have no common prime-ideal divisor.

Proposition 9.1 (The exact consecutive ray transport)**.** *At every nonterminal transition define*

$$\rho_j = [\overline{r_j}] \in \mathscr{R}_{(r_{j-1})} = (\mathcal{O}_K/(r_{j-1}))^\times / \operatorname{im}(\mathcal{O}_K^\times).$$

Then

$$\rho_{j+1} = [\overline{r_{j+1}}] = [-\overline{r_{j-1}}] = [\overline{r_{j-1}}] \quad \text{in } \mathscr{R}_{(r_j)}. \tag{222}$$

Moreover the single CRT class

$$\Theta_j = [\overline{r_j - r_{j-1}}] \in (\mathcal{O}_K/(r_j r_{j-1}))^\times / \operatorname{im}(\mathcal{O}_K^\times) \tag{223}$$

maps to ρ_j modulo (r_{j-1}) and to ρ_{j+1} modulo (r_j).

Proof. Reducing the recurrence modulo (r_j) removes the digit term and gives $r_{j+1} \equiv -r_{j-1} \pmod{r_j}$. The sign disappears in the quotient because -1 is a global unit. Since adjacent remainders are coprime, the Chinese remainder theorem applies to the relatively prime ideals (r_j) and (r_{j-1}). The element $r_j - r_{j-1}$ reduces to r_j modulo the second ideal and to $-r_{j-1}$ modulo the first, which proves the last assertion. □

The CRT class contains slightly more information than the ordered pair of its two displayed images: quotienting the two factors separately permits two independent unit changes, whereas (223) permits only one diagonal global unit. Thus consecutive fusion does not merely list two unrelated ray classes; it retains their relative unit normalization. The calculation itself is elementary. Its structural content is that the shear $T(-n_j\lambda)$ does not change the residue state at all: inversion exchanges numerator and modulus, and the digit chooses the next modulus. This is the precise finite-place information carried by the word "consecutive."

There is also a classical ambient reduction against which this restricted quotient can be measured. Recall that the Euclidean minimum of a number field is

$$M(K) = \sup_{\xi \in K} \inf_{q \in \mathcal{O}_K} |\operatorname{N}(\xi - q)|.$$

Proposition 9.2 (The exact ambient Euclidean correction)**.** *For the real cubic field $K = \mathbb{Q}(2\cos(\pi/7))$, one has $M(K) = 1/7$. Consequently, at every nonterminal Rosen step there are $q_j \in \mathcal{O}_K$ and*

$$\kappa_j = q_j + n_j \in \mathcal{O}_K$$

such that

$$\left|\operatorname{N}\big(r_{j+1} + \kappa_j \lambda r_j\big)\right| \le \frac{1}{7} |\operatorname{N}(r_j)|. \tag{224}$$

If r_{j+1} is a new absolute-norm record, then every such κ_j is nonzero.

Proof. The equality $M(K) = 1/7$ is the discriminant- 49 entry in the exact tables of Cavallar–Lemmermeyer [16]. Apply its defining inequality to

$$\xi_j = \frac{r_{j-1}}{\lambda r_j}.$$

Since λ is a unit, some $q_j \in \mathcal{O}_K$ satisfies

$$|\,\mathrm{N}(r_{j-1} - q_j \lambda r_j)| \le \frac{1}{7}|\,\mathrm{N}(r_j)|.$$

The recurrence and $\kappa_j = q_j + n_j$ give

$$r_{j-1} - q_j \lambda r_j = -r_{j+1} - \kappa_j \lambda r_j,$$

which proves (224). If $\kappa_j = 0$, that inequality reads $|\,\mathrm{N}\, r_{j+1}| \le |\,\mathrm{N}\, r_j|/7$, which is incompatible with a new record. □

The comparison is also an exact factorization in the matrix group. Put

$$A(q) = \begin{pmatrix} q\lambda & -1 \\ 1 & 0 \end{pmatrix}.$$

Then

$$A(q_j) = T(\kappa_j \lambda) B(n_j), \qquad A(q_j)\begin{pmatrix} r_j \\ r_{j-1} \end{pmatrix} = \begin{pmatrix} q_j \lambda r_j - r_{j-1} \\ r_j \end{pmatrix}. \tag{225}$$

Thus the full norm-reducing step is literally the thin Rosen step followed by one labelled correction translation. Products of the matrices $A(q_j)$ do not arise by inserting arbitrary ambient elements: the translation at time j is determined by the same state which determined the Rosen digit n_j. This is the exact relative cocycle on which a consecutive inequality can act.

Proposition 9.3 (Euclidean reduction stays in the same ray fibre)**.** *Define the ambient reduced remainder*

$$r^{\mathrm{E}}_{j+1} = q_j \lambda r_j - r_{j-1} = r_{j+1} + \kappa_j \lambda r_j.$$

Then

$$(r^{\mathrm{E}}_{j+1}, r_j) = \mathcal{O}_K, \tag{226}$$

$$r^{\mathrm{E}}_{j+1} \equiv r_{j+1} \pmod{r_j}, \tag{227}$$

$$|\,\mathrm{N}\, r^{\mathrm{E}}_{j+1}| \le \frac{1}{7}|\,\mathrm{N}\, r_j|. \tag{228}$$

Thus r^{E}_{j+1} and r_{j+1} give the same element of $\mathscr{R}_{(r_j)}$. If r_{j+1} is a new norm record and $r^{\mathrm{E}}_{j+1} \neq 0$, then

$$\frac{|\,\mathrm{N}\, r_{j+1}|}{|\,\mathrm{N}\, r^{\mathrm{E}}_{j+1}|} > 7. \tag{229}$$

At the next modulus one has the exact gcd and residue relations

$$(r^{\mathrm{E}}_{j+1}, r_{j+1}) = (\kappa_j, r_{j+1}), \tag{230}$$

$$r^{\mathrm{E}}_{j+1} \equiv \kappa_j \lambda r_j \pmod{r_{j+1}}. \tag{231}$$

In particular, when κ_j is a unit, the ambient reduced remainder represents the consecutive reciprocal ray class $[\overline{r_j}] \in \mathscr{R}_{(r_{j+1})}$ as well.

Proof. Equation (227) is the definition of r^{E}_{j+1}. Since $(r_{j+1}, r_j) = \mathcal{O}_K$, it implies (226); it also proves equality of the ray classes. Inequality (228) is (224). Combining it with $|\,\mathrm{N}\, r_{j+1}| > |\,\mathrm{N}\, r_j|$ proves (229).

For the last two assertions, reduce the same identity modulo (r_{j+1}). Since λ is a unit and $(r_j, r_{j+1}) = \mathcal{O}_K$, localization at each prime gives

$$\min\{v_{\mathfrak{p}}(r^{\mathrm{E}}_{j+1}), v_{\mathfrak{p}}(r_{j+1})\} = \min\{v_{\mathfrak{p}}(\kappa_j), v_{\mathfrak{p}}(r_{j+1})\}.$$

This is (230); the congruence is (231). If κ_j is a unit, its factor, and the factor λ, disappear in the ray quotient. □

In the splitting notation, put

$$s_j = -\frac{r_{j+1}}{r_j}, \qquad s^{\mathrm{E}}_j = -\frac{r^{\mathrm{E}}_{j+1}}{r_j}.$$

Then

$$s^{\mathrm{E}}_j = s_j - \kappa_j\lambda, \qquad (s_j, r_j^{-1}) = (s^{\mathrm{E}}_j, r_j^{-1}) = r_j^{-1}\mathcal{O}_K. \tag{232}$$

Thus the two splitting lines have exactly the same finite content. The factor seven is a comparison between two integral translates in one Arakelov ray fibre, not between unrelated moduli. Any remaining loss is therefore wholly in their three archimedean Hodge norms.

There is a telescoping version which uses exactly the amount needed for the classification. Put

$$R_j = \max_{-1\le k\le j} |\,\mathrm{N}\, r_k|$$

and, at a record-producing transition, define the regularized Euclidean excess

$$\varepsilon^{\mathrm{E}}_j = \log \frac{|\,\mathrm{N}\, r_{j+1}|}{7\max\{1, |\,\mathrm{N}\, r^{\mathrm{E}}_{j+1}|\}}. \tag{233}$$

Proposition 9.4 (Record growth is bounded by Euclidean excess)**.** *If r_{j+1} is a new record and $R_j \ge 7$, then*

$$0 < \log \frac{R_{j+1}}{R_j} \le \varepsilon^{\mathrm{E}}_j. \tag{234}$$

Consequently, if $\mathcal{R}$ denotes all record-producing transitions, then

$$\log \sup_j |\,\mathrm{N}\, r_j| \le \log\max\{7, R_0\} + \sum_{j\in\mathcal{R}} \max\{0, \varepsilon^{\mathrm{E}}_j\}. \tag{235}$$

In particular, the cumulative estimate

$$\sum_{j\in\mathcal{R}} \max\{0, \varepsilon^{\mathrm{E}}_j\} \le 6h(P) + O_7(1) \tag{236}$$

implies the norm-record bound and the full classification.

Proof. By (228),

$$\max\{1, |\,\mathrm{N}\, r^{\mathrm{E}}_{j+1}|\} \le R_j/7$$

when $R_j \ge 7$. Since $R_{j+1} = |\,\mathrm{N}\, r_{j+1}| > R_j$, division proves (234). Summing the logarithmic record increments telescopes and proves (235). Under (236), all norms are bounded by $\exp(6h(P)+O_7(1))$; Proposition 9.10 then gives termination or repetition. □

When $r_{j+1}^{\mathrm{E}} \neq 0$, the excess is the boundary Arakelov-degree difference of the two same-content splitting lines, less $\log 7$:

$$\varepsilon_j^{\mathrm{E}} = \sum_{i=0}^{2} \log \left| \frac{\sigma_i(s_j)}{\sigma_i(s_j^{\mathrm{E}})} \right| - \log 7. \tag{237}$$

Here the boundary splitting target has been normalized to zero. If the Euclidean remainder vanishes, $s_j^{\mathrm{E}} = 0$ lies on the ambient cusp divisor; the regularization in (233) is the finite-height version of that logarithmic boundary value. Thus (236) asks the stopped Hodge transport to bound a sum of positive degree differences in one fixed finite content, after removing the strict norm-Euclidean contribution $\log 7$ at every record.

The preceding Euclidean representative can be unnecessarily small. If $|\mathrm{N}\, r_j| \ll R_j$, charging its entire gap to the next record counts again the descent which has already occurred since the preceding record. A global geometry-of-numbers construction produces a representative at the correct record scale and simultaneously removes its unit imbalance.

Lemma 9.5 (Uniform Arakelov covering radius)**.** *There is a constant $C_{\mathrm{cov}} = C(K)$ such that, for every nonzero integral ideal $\mathfrak{a}$, every $x = (x_0, x_1, x_2) \in K \otimes_{\mathbb{Q}} \mathbb{R}$, and every coset $c + \mathfrak{a}$, there is $e \in c + \mathfrak{a}$ satisfying*

$$\max_i |\sigma_i(e) - x_i| \leq C_{\mathrm{cov}} (\mathrm{N}\, \mathfrak{a})^{1/3}. \tag{238}$$

Proof. The field K has class number one, so $\mathfrak{a} = (r)$. Dirichlet's unit theorem, applied to the compact quotient of the trace-zero logarithmic plane by the unit lattice, gives a unit u for which

$$|\sigma_i(ur)| \asymp_K |\mathrm{N}\, r|^{1/3} \qquad (0 \leq i \leq 2).$$

The ideal does not change: $(ur) = (r)$. Multiply one fixed integral basis of $\mathcal{O}_K$ by ur. Every edge of the resulting fundamental parallelotope has sup norm $O_K(|\mathrm{N}\, r|^{1/3})$. Reducing $x - c$ into that parallelotope proves (238). □

Proposition 9.6 (A balanced representative at the old-record shell)**.** *Let $a, r \in \mathcal{O}_K$ satisfy $(a, r) = \mathcal{O}_K$, put*

$$R \geq |\mathrm{N}\, r|, \qquad T = (R/7)^{1/3}, \qquad \rho = C_{\mathrm{cov}} |\mathrm{N}\, r|^{1/3},$$

and suppose $T > 2\rho$. Then the residue class of a modulo (r) contains an element e for which

$$T - 2\rho \leq \sigma_i(e) \leq T \qquad (0 \leq i \leq 2), \tag{239}$$

$$\frac{R}{7}(1 - \delta)^3 \leq \mathrm{N}\, e \leq \frac{R}{7}, \qquad \delta = 2C_{\mathrm{cov}} \left(\frac{7|\mathrm{N}\, r|}{R} \right)^{1/3} < 1. \tag{240}$$

Moreover $(e, r) = \mathcal{O}_K$, and there is a $\kappa \in \mathcal{O}_K$ such that

$$e = a + \kappa \lambda r.$$

Proof. Apply Lemma 9.5 to the coset $a + (r)$ and the Minkowski target $(T - \rho, T - \rho, T - \rho)$. This gives (239); positivity and multiplication of the three inequalities give (240). Congruence to a preserves primitivity with r. Finally $e - a \in (r) = (\lambda r)$, because λ is a unit. □

Apply the proposition with $a = r_{j+1}$, $r = r_j$, and $R = R_j$. For every sufficiently deep inter-record descent, namely

$$R_j > 56 C_{\mathrm{cov}}^3 |\mathrm{N}\, r_j|,$$

it supplies a same-ray splitting $s_j^{\mathrm{sh}} = -e/r_j$ with the same finite content $r_j^{-1}\mathcal{O}_K$, no uncontrolled unit shape, and

$$\log \frac{R_{j+1}}{R_j} \le \left(\log \frac{|\,\mathrm{N}\, r_{j+1}|}{|\,\mathrm{N}\, e|} - \log 7 \right) \le \log \frac{R_{j+1}}{R_j} - 3\log(1-\delta_j), \tag{241}$$

where

$$\delta_j = 2C_{\mathrm{cov}} \left(\frac{7|\,\mathrm{N}\, r_j|}{R_j} \right)^{1/3}.$$

Thus the excessive gaps seen when the norm-Euclidean remainder is much smaller than the old record are not intrinsic. They are removed by one all-place Arakelov reduction in the same affine ray fibre.

> **The obstruction after global balancing.** The record problem has now split into two precise regimes. On a deep descent, Proposition 9.6 supplies a balanced second splitting line whose boundary degree gap differs from the actual record increment only by the explicit covering error $-3\log(1-\delta_j)$. On a shallow descent,
>
> $$|\,\mathrm{N}\, r_j| \ge (56C_{\mathrm{cov}}^3)^{-1} R_j,$$
>
> the finite coordinate of the current pair is already comparable with the record; no missing ideal size has to be recovered. The remaining Hodge statement is therefore not an unrestricted ray-class lifting theorem. It is a stopped two-line inequality for a balanced shell representative in the deep regime, together with pairwise-energy control when the current denominator is comparable with the record. This is the first reduction which removes the large inter-record recharge by a global Arakelov argument rather than attempting to account for it one prime or one generator at a time.

The two transverse additive directions also have a canonical structure. Choose the generator τ of $\mathrm{Gal}(K/\mathbb{Q})$ for which

$$\tau(\lambda) = \lambda^2 - \lambda - 1, \qquad \tau(\lambda^2) = 2 - \lambda.$$

Proposition 9.7 (The missing quotient is an Eisenstein lattice)**.** *Modulo $\mathbb{Z}$, with basis $(\bar{\lambda}, \overline{\lambda^2})$, the action of τ is*

$$\begin{pmatrix} b \\ c \end{pmatrix} \longmapsto \begin{pmatrix} -1 & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} b \\ c \end{pmatrix}. \tag{242}$$

Its invariant positive quadratic form is

$$Q(b,c) = b^2 + bc + c^2.$$

More intrinsically, for $\kappa = a + b\lambda + c\lambda^2$,

$$\sum_{i=0}^{2} \left(\sigma_i(\kappa) - \frac{1}{3}\,\mathrm{Tr}(\kappa) \right)^2 = \frac{14}{3} Q(b,c). \tag{243}$$

Consequently every τ-stable additive module

$$\mathbb{Z} \subsetneq M \subseteq \mathcal{O}_K$$

has rank three and finite index in $\mathcal{O}_K$. The group generated by G_7 and $T(\lambda M)$ therefore has finite index in $\mathrm{SL}_2(\mathcal{O}_K)$.

Proof. The two displayed formulas for τ follow from the minimal polynomial $\lambda^3 - \lambda^2 - 2\lambda + 1 = 0$. Reduction modulo $\mathbb{Z}$ gives (242); direct substitution shows that Q is invariant. Newton sums give

$$\operatorname{Tr}(\lambda) = 1, \quad \operatorname{Tr}(\lambda^2) = 5, \quad \operatorname{Tr}(\lambda^3) = 4, \quad \operatorname{Tr}(\lambda^4) = 13.$$

Expanding the centered square gives (243).

The matrix in (242) has irreducible characteristic polynomial $X^2 + X + 1$ over $\mathbb{Q}$. Hence the τ-orbit of every nonzero element of $(\mathcal{O}_K/\mathbb{Z}) \otimes \mathbb{Q}$ spans that two-dimensional space. The image of M therefore has rank two, so M has rank three and finite index in $\mathcal{O}_K$. The translations $T(\lambda M)$ form a lattice in the full three-dimensional Minkowski horosphere. Benoist–Miquel, as in Proposition 7.6, proves the final assertion. □

> **Why a natural intermediate group breaks.** The quotient $\mathcal{O}_K/\mathbb{Z}$ is one rank-one module over the Eisenstein integers $\mathbb{Z}[\zeta_3]$; its two ordinary integer directions are exchanged irreducibly by Galois. A rank-two quotient algorithm must choose one non-Galois-invariant integer line, as H_2 does. Closing that line under the full Galois/Hodge symmetry supplies both missing directions, a full cusp lattice, and hence arithmeticity. The Galois-compatible object which remains thin is therefore the Eisenstein-valued correction $\bar{\kappa}_j$ attached to each labelled step, not a group generated by all of its possible values.

The same Eisenstein norm has an exact Hodge-geometric interpretation. Let $F_0(z) = z, F_1(z), F_2(z)$ be the three modular embeddings and let $\mathcal{K} \Subset \mathbb{H}$ be a compact return set. For $\bar{\kappa} \in \mathcal{O}_K/\mathbb{Z}$, define

$$\mathcal{D}_z(\bar{\kappa}) = \inf_{m \in \mathbb{Z}} \sum_{i=0}^{2} d_{\mathbb{H}}\left(F_i(z), T\big(\sigma_i(\lambda(\kappa - m))\big) F_i(z)\right). \tag{244}$$

Changing the lift κ only changes the indexing integer m.

Proposition 9.8 (The correction norm is transverse Hodge distance)**.** *There are constants $0 < c_{\mathcal{K}} \le C_{\mathcal{K}}$ such that, for every $z \in \mathcal{K}$ and $\kappa = a + b\lambda + c\lambda^2$,*

$$c_{\mathcal{K}} \log\big(1 + Q(b,c)\big) \le \mathcal{D}_z(\bar{\kappa}) \le C_{\mathcal{K}} \log\big(1 + Q(b,c)\big). \tag{245}$$

In particular, $\mathcal{D}_z(\bar{\kappa}) = 0$ exactly when $\kappa \in \mathbb{Z}$.

The left side of (244) is the distance, in the product of the three real upper half-planes, from the ambient translate $T(\lambda\kappa) \cdot (F_0(z), F_1(z), F_2(z))$ to the thin cusp orbit

$$\{T(m\lambda) \cdot (F_0(z), F_1(z), F_2(z)) : m \in \mathbb{Z}\}.$$

Proof. For $w \in \mathbb{H}$ and $t \in \mathbb{R}$,

$$d_{\mathbb{H}}(w, T(t)w) = 2 \operatorname{arsinh} \frac{|t|}{2 \operatorname{Im} w}.$$

On the compact set $\mathcal{K}$, all three positive functions $\operatorname{Im} F_i$ are bounded above and below. Multiplication by λ is an invertible linear map on the three-dimensional Minkowski space. Moreover (243) says that the squared distance of κ from the rational line $\mathbb{R} \cdot 1$ is a fixed positive multiple of $Q(b,c)$. Finite-dimensional norm equivalence therefore gives, uniformly in m,

$$\max_i |\sigma_i(\lambda(\kappa - m))| \gg \sqrt{Q(b,c)}.$$

Taking $m = a$ gives the reverse upper bound for one term in the infimum. Finally $\operatorname{arsinh} x \asymp \log(1 + x)$ away from zero, and the nonzero values of Q are positive integers. This proves (245). Equivariance $F_i(z + m\lambda_0) = T(m\lambda_i)F_i(z)$ identifies the minimizing set with the thin cusp orbit. □

A weaker cumulative target already closes. Let $\mathcal{R}$ be the set of new norm-record transitions, and choose one Euclidean correction at each of them. Put

$$\chi(\kappa) = \begin{cases} \log(1 + Q(b,c)), & Q(b,c) > 0, \\ \log(1 + |a|), & \kappa = a \in \mathbb{Z} \setminus \{0\}. \end{cases}$$

Every record correction is nonzero, so $\chi(\kappa) \geq \log 2$. Consequently the cumulative estimate

$$\sum_{j \in \mathcal{R}} \chi(\kappa_j) \leq 6h(P) + O_7(1) \tag{246}$$

would already imply the classification: it permits only finitely many new records, hence all remainder norms are bounded, and Proposition 9.10 closes the orbit.

For $Q > 0$, Proposition 9.8 identifies this charge with normal displacement from the thin cusp orbit. For $Q = 0$, the alternative Euclidean matrix is another triangle-group branch and $\log(1 + |a|)$ is ordinary cusp separation. Thus (246) is a two-channel formulation of the remaining task: normal Hodge displacement for a transverse correction, and central Green displacement for an integral correction. It is formally weaker than the pointwise record bound (249), but it retains the full consecutive constraint and has a strictly positive cost at every record.

What the factor 1/7 changes. The deep input here is the exact Euclidean minimum; the displayed deduction from it is elementary. The Rosen algorithm permits only the quotient $-n_j \in \mathbb{Z}$, whereas ambient norm reduction permits $q_j \in \mathcal{O}_K$. Hence

$$\overline{\kappa}_j = \overline{q_j} \in \mathcal{O}_K/\mathbb{Z}$$

is a rank-two correction. A nonzero class is a genuinely unavailable quotient direction. If the class is zero, the ambient correction is a different integer quotient, and the remaining discrepancy is the distinguished nearest-integer rule. Thus norm-Euclideanity does not by itself prove the record bound. It replaces the vague Roth obstruction by a sharper target: the consecutive Hodge estimate must prevent these ambient corrections from carrying the compensation at successive record times. Equivalently, it must force the factor-1/7 descent to occur inside the restricted Rosen line often enough to close the orbit.

Proposition 9.9 (The exact closing criterion)**.** *If a remainder vanishes, the initial point is a cusp. If a nonzero projective state $[r_j : r_{j-1}]$ repeats, the initial point belongs to the G_7-orbit of a fixed point of a hyperbolic element whose eigenvalues lie in K.*

Proof. If $r_j = 0$, the inverse prefix word carries zero to the initial point, and $0 = S\infty$. Thus the initial point belongs to $G_7\infty$.

If the states at times $a < b$ agree, the nonempty branch word between them fixes the state at time a. On the interior of I, each branch of $\mathcal{R}$ has derivative $1/x^2 > 4/\lambda_0^2 > 1$. An interior periodic word is therefore hyperbolic; the parabolic endpoint codes terminate under the half-open convention. Conjugating by the prefix gives an element fixing the initial point. If $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ fixes $x \in K$, then

$$g\begin{pmatrix} x \\ 1 \end{pmatrix} = (cx + d)\begin{pmatrix} x \\ 1 \end{pmatrix}.$$

Thus one eigenvalue is $cx + d \in K$, and the other is its inverse. □

Elementary dynamical skeleton. The recurrence, preservation of primitivity, and closing criterion use only the explicit branch matrices and expansion of the real Rosen map. They explain what termination and repetition mean, but do not show that one of them must occur.

9.2 Why bounded norms force closing

The remaining arithmetic implication is an S-unit argument, but it is important to state it correctly. A single recurrence relation contains the unrestricted integer digit n_j, so one cannot simply call its three terms S-units. The Galois difference removes the digit.

At a nonterminal transition put

$$z_j = \frac{r_{j-1}}{\lambda r_j}, \qquad w_j = \frac{r_{j+1}}{\lambda r_j}.$$

Then

$$z_j + w_j = -n_j, \qquad z_j + w_j - \tau(z_j) - \tau(w_j) = 0, \tag{247}$$

where τ generates the cyclic Galois group of $K/\mathbf{Q}$.

Proposition 9.10 (Fixed finite support closes)**.** *Suppose that every nonzero remainder is supported on one finite set of prime ideals. Then the Rosen orbit terminates or its projective state repeats.*

Proof. Enlarge the finite set so that it is stable under Galois. The four coordinates in the second relation of (247) then lie in one finitely generated multiplicative group. The projective S-unit theorem leaves only finitely many projective classes for relations having no vanishing proper subsum [18].

There is still a projective scalar to remove. Galois compatibility forces that scalar to lie in $\mathbf{Q}^\times$. Two consecutive transitions satisfy

$$z_{j+1} = \frac{1}{\lambda^2 w_j}. \tag{248}$$

Once the two projective four-term relations are fixed, (248) fixes the product of their rational scalars and hence the nonzero integer $n_j n_{j+1}$. That integer has only finitely many factorizations, so only finitely many exact two-transition blocks can occur.

For completeness, a degenerate four-term relation forces $z_j, w_j \in \mathbf{Q}^\times$: its nonzero terms can cancel only in two Galois-compatible pairs. The following transition is nondegenerate because $z_{j+1} = 1/(\lambda^2 w_j) \notin \mathbf{Q}$. If the next one were also degenerate, the recurrence would force a nonzero rational multiple of λ^2 to be rational. Hence an infinite orbit contains infinitely many consecutive nondegenerate pairs.

There are only finitely many exact pairs, so an infinite deterministic orbit repeats a state. Proposition 9.9 identifies the resulting periodic alternative. □

> **Deep number-theoretic input.** The finiteness theorem for projective S-unit equations is deep. The Galois subtraction and two-step glue are elementary but essential: without them the variable Rosen digit and the remaining projective scalar would leave a logical gap.

If all absolute norms $|\mathrm{N}\, r_j|$ are bounded, only finitely many prime ideals can divide the remainders, because a prime divisor of r_j has norm at most $|\mathrm{N}\, r_j|$. Proposition 9.10 therefore reduces the classification to one estimate:

$$\log|\mathrm{N}\, r_j| \le 6h(P) + B_7 \qquad \text{for every } j. \tag{249}$$

It is enough to prove this only when r_j first exceeds all earlier remainder norms. The coefficient six is the natural scale rather than a convenient overestimate: the initial projective height has three real components, while the dangerous cusp cancellation leaves a square. Exact unit families attain this scale up to a bounded term.

10 The stopped scalar

Immediately before a transition, write the primitive state as $X = (u, v)$ and put

$$H_i(X) = \log \max\{|\sigma_i(u)|, |\sigma_i(v)|\}, \qquad \eta_i = H_i(X) - \log|\sigma_i(v)|. \tag{250}$$

Each η_i is nonnegative. Primitivity gives

$$H_0(X) + H_1(X) + H_2(X) = 3h(X). \tag{251}$$

Define

$$\Theta(X) = \exp(\eta_0 - \eta_1 - \eta_2). \tag{252}$$

η_i is a projective proximity to the coordinate $v = 0$ at the i-th embedding. For example,

$$e^{\eta_0} = \max\{1, |\sigma_0(u/v)|\}.$$

Thus η_0 records the central cusp depth, while $\eta_1 + \eta_2$ records how much of that depth is already visible in the two conjugate directions. The ratio Θ is large only when the distinguished state is much closer to the cusp than both conjugate states.

The following estimate is the elementary bridge from this real geometry to the norm of the next algebraic remainder.

Proposition 10.1 (Global cusp transfer)**.** *Let $X = (u, v)$ be a primitive integral state with $v \neq 0$, and suppose*

$$r = -u - n\lambda v, \qquad |\sigma_0(r)| \le C_c|\sigma_0(v)|.$$

Then

$$\log|\mathrm{N}\, r| \le 3h(X) + \log(1 + \Theta(X)) + C(C_c, K). \tag{253}$$

Proof. Write $P_i = e^{H_i(X)}$, $v_i = \sigma_i(v)$, and $\lambda_i = \sigma_i(\lambda)$. At the distinguished embedding,

$$|n|\,|\lambda_0|\,|v_0| \le |u_0| + |r_0| \le (1 + C_c)P_0.$$

The same integer n occurs at both outer embeddings, so

$$|r_i| \le P_i + C\frac{P_0}{|v_0|}|v_i|, \qquad i = 1, 2.$$

Multiplying these two inequalities and using $|r_0| \le C_c|v_0|$, we find that $|\mathrm{N}\, r|$ is at most a field constant times the sum of

$$|v_0|P_1P_2, \quad P_0P_1|v_2|, \quad P_0P_2|v_1|, \quad \frac{P_0^2|v_1v_2|}{|v_0|}.$$

The first three terms are bounded by a constant times $P_0P_1P_2$. The last is exactly

$$P_0P_1P_2\,\frac{P_0|v_1v_2|}{|v_0|P_1P_2} = P_0P_1P_2\,\Theta(X).$$

Finally, $P_0P_1P_2 = e^{3h(X)}$, which proves the result. □

For an exact Rosen transition, the next ratio r/v lies in I, so one may take $C_c = \lambda_0/2$. The estimate therefore treats an arbitrarily large digit in one step. The digit and all of its prime factors have disappeared; the only exceptional term is the explicit archimedean scalar Θ.

There are two useful ways to read the result. Directly,

$$|\,\mathrm{N}\, r| \ll_7 e^{3h(X)}(1+\Theta(X)).$$

Using the product formula, one also has the exact identity

$$3h(X) + \log\Theta(X) = \log|\,\mathrm{N}\, v| + 2\eta_0. \tag{254}$$

Thus the dangerous part of the next norm is the old remainder norm multiplied by the square of the central cusp depth. This is the arithmetic form of the missing-central-copy phenomenon in Section 5.

It follows that the classification would follow from

$$3h(X) + \log(1+\Theta(X)) \le 6h(P) + O_7(1) \tag{255}$$

whenever the transition creates a new norm record. This is the precise place at which the finite stopping rule must meet Hodge contraction.

> **Elementary but structurally important.** Proposition 10.1 is a triangle-inequality and product-formula calculation; it is not a Hodge theorem. Its significance is that it replaces every place-by-place divisibility question by one scalar. Hodge theory is asked only to control that scalar at the correct stopping times.

11 Replacement ideas and what they actually establish

The pairwise Arakelov construction below was reached only after several more immediate replacements for McMullen's height had been tested. This section separates the useful mathematical content of those attempts from the claims they do not prove.

11.1 The ordinary three-factor height

The most natural projective height uses all three Hodge factors with equal weight. It passes the unit test (20), but along the distinguished geodesic its derivative is bounded only by

$$-1 + \|DF_1\|_{\mathrm{hyp}} + \|DF_2\|_{\mathrm{hyp}}.$$

Near the cusp, both derivative norms tend to one, so this upper bound tends to 1, not to a negative number. This is a sign obstruction, not a failure to sharpen a numerical constant. The ordinary height has one central contracting copy whereas the cubic norm estimate needs two.

11.2 The two pairwise Hodge products

Each product involving the central Hodge factor and one outer factor has strictly negative derivative. Their product has logarithmic derivative

$$-2 + \|DF_1\|_{\mathrm{hyp}} + \|DF_2\|_{\mathrm{hyp}} < 0,$$

which is precisely the right analytic sign. Taken in isolation, however, it has weight $(2,1,1)$ and changes under multiplication by a unit. It is therefore not a function on $\mathbf{P}^1(K)$. The pairwise

products are not a failed insight; they are the two archimedean weights which the new Arakelov line will retain. Two separate issues were previously conflated: they need a finite completion to become projectively well defined, and they need a flat transport law before their fixed-covector contraction can be applied along an orbit. The first issue is solved below; the second is not.

11.3 The negative Hodge line

The degree calculation

$$\deg(\mathcal{L}_1\mathcal{L}_2\mathcal{L}_0^{-1}) = -\frac{1}{28}$$

shows that the variation contains a canonical correction of the required sign. The corresponding ratio of Higgs fields identifies a logarithmic cotangent factor at the cusp. In an adapted nilpotent-orbit metric it transfers one copy of the central cusp depth with bounded endpoint error. This is genuine Hodge-theoretic structure.

It still does not specify the integral lattice of an evaluated remainder. A line of negative geometric degree can be twisted by many inequivalent finite metrics. Unless the finite metric is fixed and its divisor is shown to be the remainder norm, the construction is not an arithmetic height. The incidence tensor in Section 13 supplies this missing integral information.

11.4 Real word contraction and finite-state fusion

For a return word g, let

$$a_i(g) = \log \rho(\sigma_i(g)), \qquad \delta(g) = a_0(g) - \max\{a_1(g), a_2(g)\}.$$

Strict Schwarz–Pick contraction makes $\delta(g) > 0$ on a genuine return through the compact part. By inducing past maximal parabolic powers, one obtains a coarse inequality

$$\delta(g) \geq c_7\, \nu(g) - C_7,$$

where $\nu(g)$ counts compact returns rather than raw letters. Exact finite-state calculations with the six matrices A_j give much sharper versions on large symbolic classes.

These estimates are useful: they measure how much real Hodge contraction is available. They do not contain the norm divisor of a remainder. Even the optimal real fusion constant would therefore be only the archimedean half of the proof. Its proper role is to provide curvature or strict decrease for the same metrized line whose finite part records $|\,\mathrm{N}\, r|$.

11.5 Roth-type approximation

Roth's theorem says that, for a fixed algebraic irrational ξ and a fixed $\varepsilon > 0$, there are only finitely many rational approximations satisfying

$$|\xi - p/q| < q^{-2-\varepsilon}.$$

The continued-fraction constant C/q^2 is qualitatively weaker. When $C < 1/2$ it essentially says that p/q is a convergent; with merely $C = O(1)$ it does not force periodicity or termination.

The heptagonal near-parabolic family produces a more delicate threshold: the exponent has the form

$$2 + \frac{2\delta_7}{\log N} + o\!\left(\frac{1}{\log N}\right).$$

For every fixed N it exceeds two, but the excess approaches zero. A standard application of Roth fixes ε before the denominator grows and therefore does not see this regime. Quantitative forms of Roth also deteriorate as $\varepsilon \to 0$.

Moving-target versions of the Subspace Theorem permit targets of small height relative to the approximating points. Here the natural target comes from the Rosen prefix, and its height is comparable with the parabolic parameter itself. Thus the usual small-target hypothesis is not satisfied. A conjectural theorem strong enough to finish by this route would have to exploit the common triangle-group origin of target and approximant; a generic optimal constant in C/q^2 would not suffice.

The order-seven factorization makes the obstruction stronger. With the notation of Proposition 17.32, put

$$\beta_g = \frac{E_g}{2D_g} \in K.$$

Then

$$x - \beta_g = \frac{A_g(P)}{2qD_g}. \tag{256}$$

Thus both the target x and every natural elliptic-center approximant β_g already lie in K. Equivalently, after the Cayley transformation the local target is the root of unity -1. A fixed-target irrationality theorem, no matter how sharp its exponent or effective constant, therefore addresses the wrong datum. The required statement is an orbit-height theorem which forces equality in (256); it is not an irrationality measure for a target outside the approximation field.

11.6 Vojta and global height conjectures

Vojta's framework is closer in spirit because it places all finite valuations into one truncated boundary term. Applied only to the original projective line, however, the log-canonical degree is critical: the apparent boundary gain is canceled by the same unit normalization which caused the degree-three defect. A meaningful conditional statement must first lift the orbit to the Hodge–Arakelov incidence space, where the two outer contractions contribute positive curvature while the primitive finite metric remains unchanged. Vojta on that explicitly constructed lift could imply the stopped inequality. Invoking it before constructing the lift would merely give a new name to the missing estimate.

11.7 Finite-divisor bookkeeping (not a proof route)

One can record when a prime ideal first divides a remainder, how long it takes to reappear, and which continuant it then divides. The return divisibility statements are exact and useful diagnostics. They do not control a record whose divisor is spread among many primes, and a single later record can replace most of its support while increasing its norm by only a small factor. The product formula suggests the correct repair: sum the entire finite divisor first, then compare it with one global archimedean degree. This is the Arakelov viewpoint adopted below.

11.8 Operator norms versus evaluated tensors

The Hodge cocycle has expanding directions. Passing to an operator norm removes dependence on the selected vector, but it also discards the incidence angle which makes a particular remainder small. Conversely, tracking only the evaluated scalar retains the angle but can obscure the ambient Hodge contraction. The construction below keeps an evaluated one-dimensional tensor inside

the two components having the desired Hodge weights. This is why it is more specific than a norm bound for the full 27-dimensional Asai representation. It still changes with the prefix, so membership in those components is not itself a contraction statement.

route	what survives in the present argument
ordinary Hodge height	the correct projective normalization, but not the needed cusp sign;
pairwise products	the two strict real contractions for a fixed flat covector, and the unresolved transport problem for prefix covectors;
negative Hodge line	the missing central cotangent correction;
finite-state fusion	a quantitative compact-return contraction budget;
Roth and moving targets	the critical exponent and the need for correlated targets;
Vojta	the possibility of a conditional global inequality after an explicit Hodge–Arakelov lift;
prime divisors	exact return identities, but not the global norm estimate;
evaluated incidence	retention of the line selected by the norm-record stopping rule.

The new direction. The surviving pieces point to one object: a K-line whose two Hodge metrics have the pairwise weights and whose finite fractional ideal is the reciprocal remainder ideal. The next section constructs that line first at the elementary metric level; Section 13 then realizes it pointwise inside the Hodge representation. Neither construction supplies the missing transport law.

12 The pairwise Arakelov energy

The new observation is that the pointwise sum of the two Hodge quantities has a direct arithmetic completion. Here "pointwise" is essential: the completion constructed below is a genuine metrized line at one Rosen state, but this fact alone says nothing about how lines belonging to two successive states are transported.

The construction is forced by the preceding calculations. The real Hodge descent wants the weight $2\eta_0 + \eta_1 + \eta_2$. To make this homogeneous, one must add the finite divisor of the coordinate v. The product formula then turns the apparent weight $(2, 1, 1)$ into the degree of a single K-line. No prime ideal is chosen, and no unit is chosen during the orbit.

Definition 12.1 (Pairwise Arakelov energy)**.** For a primitive state $X = (u, v)$ with $v \neq 0$, put

$$\mathscr{E}(X) = \log|\mathrm{N}\, v| + 2\eta_0 + \eta_1 + \eta_2. \tag{257}$$

This is not the ordinary height of X. The coefficient two on η_0 records the fact that the distinguished Hodge factor occurs in both pairwise contractions (18). The finite term makes the expression invariant under multiplication of (u, v) by a unit.

Three elementary regimes help interpret the definition.

1. If v is not small relative to u at any embedding, all η_i vanish and $\mathscr{E}(X) = \log|\mathrm{N}\, v|$.
2. If only the distinguished coordinate v_0 is small, then $\mathscr{E}$ charges twice the central depth. This is exactly the square in (254).
3. Multiplying both coordinates by a unit changes none of the η_i and changes $\log|\mathrm{N}\, v|$ by zero. Thus a large real realignment by a unit is recorded without changing the arithmetic line.

Lemma 12.2 (Two exact identities)**.** *For a primitive integral state,*

$$\mathscr{E}(X) = 3h(X) + \eta_0, \tag{258}$$

$$3h(X) + \log(1 + \Theta(X)) \le \mathscr{E}(X) + \log 2. \tag{259}$$

Proof. Put $A = \log|\mathrm{N}\, v|$. The product formula gives

$$3h(X) = H_0 + H_1 + H_2 = A + \eta_0 + \eta_1 + \eta_2,$$

which proves (258). Moreover

$$3h(X) + \log(1 + \Theta) = A + \log\left(e^{\eta_0+\eta_1+\eta_2} + e^{2\eta_0}\right).$$

Both exponents inside the last logarithm are at most $2\eta_0 + \eta_1 + \eta_2$. This proves (259). □

Lemma 12.3 (The initial budget)**.** *If $P = (u, v)$ is primitive and $v \neq 0$, then*

$$\mathscr{E}(P) \le 6h(P). \tag{260}$$

Proof. It is enough, by (258), to prove $\eta_0 \le 3h(P)$. Since v is a nonzero algebraic integer, $|\mathrm{N}\, v| \ge 1$, and hence

$$-\log|\sigma_0(v)| \le \log|\sigma_1(v)| + \log|\sigma_2(v)|.$$

The last two terms are at most $H_1(P) + H_2(P)$. Adding $H_0(P)$ gives the claim. □

Proposition 12.4 (The energy is one global metrized line)**.** *On the affine chart $v \neq 0$ of $\mathbf{P}^1_K$, give the tautological line the finite-place norm*

$$\|(u, v)\|_{\mathfrak{p}} = \max\{|u|_{\mathfrak{p}}, |v|_{\mathfrak{p}}\}$$

and the following norms at the three real places:

$$\|(u, v)\|_{\mathrm{pair},0} = \frac{\max\{|\sigma_0(u)|, |\sigma_0(v)|\}^2}{|\sigma_0(v)|}, \tag{261}$$

$$\|(u, v)\|_{\mathrm{pair},i} = \max\{|\sigma_i(u)|, |\sigma_i(v)|\}, \qquad i = 1, 2. \tag{262}$$

These are metrics on one K-line: multiplication by $k \in K^\times$ multiplies the i-th norm by $|\sigma_i(k)|$. For a primitive integral state the finite norms are one, and their summed real logarithmic size is exactly

$$\sum_{i=0}^{2} \log\|(u, v)\|_{\mathrm{pair},i} = \mathscr{E}(X). \tag{263}$$

The central metric has a logarithmic singularity at $v = 0$, the cusp divisor; it is smooth away from that divisor.

Proof. The scaling assertion follows directly from the homogeneity of (261): both powers in the numerator and the one power in the denominator leave one power of $|\sigma_0(k)|$. The outer formulas are plainly homogeneous of degree one. Finally,

$$\sum_i \log\|(u, v)\|_{\mathrm{pair},i} = 2H_0(X) + H_1(X) + H_2(X) - \log|\sigma_0(v)|.$$

The product formula for v rewrites the right side as $\log|\mathrm{N}\, v| + 2\eta_0 + \eta_1 + \eta_2$, which is (257). □

Elementary but decisive. Lemmas 12.2, 12.3, and Proposition 12.4 use only the product formula and the inequality $\log(a+b) \le \max(\log a, \log b) + \log 2$. Their importance is structural: they show that the sum of the two Hodge descents has exactly the right coefficient-six initial budget, dominates the stopped scalar actually needed, and is a genuine all-place metric rather than a formal sum of selected embeddings.

The fact that this is one global line is the advance. A formal linear combination of real logarithms can always be written down, but it need not be invariant under rescaling and it need not have an integral model. Here each local norm has degree-one homogeneity, the finite norms are fixed, and the cusp singularity is the divisor of the same coordinate v. The product formula, rather than an informal balance of embeddings, proves the descent.

There is also an exact formula for the change of this energy. It is more informative than a numerical monotonicity test because it isolates, place by place, the term which a proof would have to control. Write

$$X_j = (r_{j-1}, r_j), \qquad x_i = \frac{\sigma_i(r_j)}{\sigma_i(r_{j-1})}, \qquad y_i = \frac{\sigma_i(r_{j+1})}{\sigma_i(r_j)},$$

and put $\log^+ t = \max\{0, \log t\}$.

Lemma 12.5 (Exact energy cocycle)**.** *For every nonterminal transition,*

$$\begin{aligned}\mathscr{E}(X_{j+1}) - \mathscr{E}(X_j) = {} & 2\log|x_0| + 2\log^+|y_0| - \log|y_0| - 2\log^+|x_0| \\ & + \sum_{i=1}^{2}\bigl(\log|x_i| + \log^+|y_i| - \log^+|x_i|\bigr). \end{aligned} \tag{264}$$

After the first Rosen step, $|x_0|, |y_0| < 1$, so the central summand is simply $2\log|x_0| - \log|y_0|$. Equivalently,

$$\mathscr{E}(X_j) = \log|\mathrm{N}\, r_{j-1}| + D(x_j), \qquad D(x) = |\log|x_0|| + \log^+|x_1| + \log^+|x_2|. \tag{265}$$

Proof. The finite part changes by $\log|\mathrm{N}\, r_{j+1}| - \log|\mathrm{N}\, r_j| = \sum_i \log|y_i|$. At the i-th real place the new proximity is $\log^+|y_i|^{-1}$, while the old one is $\log^+|x_i|^{-1}$. Substitution, with weight two at $i = 0$, and the identity $\log^+(t^{-1}) = \log^+ t - \log t$ give (264). Applying the same identity directly to (257) gives (265). □

The formula exhibits the obstruction without any experiment. A large positive outer term can be paid for by a central term of either sign, and the first-record condition controls only the sum $\sum_i \log|y_i|$. Thus neither the recurrence nor the record inequality forces $\mathscr{E}$ to decrease. Any successful argument must use an additional global relation among the three embeddings; this is precisely what the proposed incidence line is meant to encode.

The remaining assertion now has a particularly clean form.

Conjecture 12.6 (Pairwise Arakelov stopping inequality)**.** *There is a constant B_{pair}, depending only on the heptagonal system, such that every state immediately preceding a new norm record satisfies*

$$\mathscr{E}(X) \le 6h(P) + B_{\mathrm{pair}}. \tag{266}$$

Theorem 12.7 (Logical consequence)**.** *Conjecture 12.6 implies Conjecture 1.1.*

Proof. At a record-producing transition, combine (253), (259), and (266). This gives (249). Every norm larger than the initial two first occurs at such a transition, so all remainder norms are bounded. Proposition 9.10 then gives termination or repetition, which are precisely the two alternatives in Conjecture 1.1. □

13 The Hodge tensor behind the distinguished metric

Proposition 12.4 settles the elementary descent of the scalar to one metrized tautological line. We next identify the Hodge tensor which produces its nonstandard central metric.

The representation used here is easier to understand from its purpose. The vector $P = (p, q)$ has three real conjugates. Their tensor product remembers all three at once:

$$Y(P) = \sigma_0(P) \otimes \sigma_1(P) \otimes \sigma_2(P).$$

If a row $w = (a, b)$ evaluates P to the remainder $r = ap + bq$, the rotated column $w^\sharp = (-b, a)^t$ satisfies $\det(P, w^\sharp) = r$. Taking the exterior product of $Y(P)$ with $Y(w^\sharp)$ therefore records three conjugate determinant calculations in one rational tensor.

Let $V_i \simeq \mathbf{R}^2$ be the three real factors of K^2, and put $W = V_0 \otimes V_1 \otimes V_2$. For $P \in K^2$, write

$$Y(P) = \sigma_0(P) \otimes \sigma_1(P) \otimes \sigma_2(P).$$

The exterior square has a canonical decomposition

$$\begin{aligned} \bigwedge^2 W = (&\bigwedge^2 V_0 \otimes \operatorname{Sym}^2 V_1 \otimes \operatorname{Sym}^2 V_2) \\ &\oplus (\operatorname{Sym}^2 V_0 \otimes \bigwedge^2 V_1 \otimes \operatorname{Sym}^2 V_2) \\ &\oplus (\operatorname{Sym}^2 V_0 \otimes \operatorname{Sym}^2 V_1 \otimes \bigwedge^2 V_2) \oplus \bigotimes_i \bigwedge^2 V_i. \end{aligned} \tag{267}$$

Denote the first three components by $\mathcal{Z}_0, \mathcal{Z}_1, \mathcal{Z}_2$. The last component is one-dimensional.

In $\mathcal{Z}_i$, the i-th two-dimensional factor is alternating and the other two are symmetric. Alternation extracts the determinant r_i; the symmetric factors retain the two projective directions needed for a Hodge norm. The last one-dimensional summand alternates in every factor and is therefore the product $r_0 r_1 r_2 = \mathrm{N}\, r$. This explains both the norm scalar and the three pairwise proximity terms below.

Let $w = (a, b)$ be a primitive prefix row and set

$$c = w^\sharp = (-b, a), \qquad r = wP = ap + bq.$$

The scalar component of $Y(P) \wedge Y(c)$ is

$$s = \mathrm{N}(r). \tag{268}$$

If $\lambda_i(w, P)$ denotes the logarithmic projective proximity at the i-th real embedding, direct expansion of (267) gives

$$\log \frac{\|\mathcal{Z}_0\|}{|s|} = \lambda_1 + \lambda_2 + O(1), \tag{269}$$

$$\log \frac{\|\mathcal{Z}_1\|}{|s|} = \lambda_0 + \lambda_2 + O(1), \tag{270}$$

$$\log \frac{\|\mathcal{Z}_2\|}{|s|} = \lambda_0 + \lambda_1 + O(1). \tag{271}$$

The last two are exactly the two pairwise Hodge quantities.

Let τ generate $\mathrm{Gal}(K/\mathbf{Q})$, numbered so that it cycles the three real factors. The coordinate ideals of the three components satisfy

$$I(\mathcal{Z}_0) = (r), \qquad I(\mathcal{Z}_1) = (\tau r), \qquad I(\mathcal{Z}_2) = (\tau^2 r). \tag{272}$$

Proof of (272). The first component has the visible determinant factor r. After it is removed, its remaining coordinates are the coefficients of the symmetric products of two primitive pairs. Over a discrete valuation ring those coefficients cannot all lie in the maximal ideal: their reductions form the product of two nonzero linear forms over a field. Thus their coefficient ideal is the unit ideal. This proves $I(\mathcal{Z}_0) = (r)$. The other two identities are its Galois conjugates. □

Definition 13.1 (Pairwise incidence tensor)**.** Define the evaluated K-line

$$\mathscr{P}_{w,P} = K \cdot \frac{\mathcal{Z}_1(w,P) \otimes \mathcal{Z}_2(w,P)}{s(w,P)}. \tag{273}$$

It carries the tensor-product Hodge metric at the distinguished real place and the lattice metric induced by (272) at finite places. The remaining real metrics of the global tautological line are those in (262).

Proposition 13.2 (Exact integral and distinguished-real realization)**.** *The finite fractional ideal of the displayed generator in* (273) *is exactly* $(r)^{-1}$*. In the Rosen state frame* $X = (u, v) = gP$*, with* $r = v$*, the sum of its two relative Hodge logarithms is*

$$2\eta_0 + \eta_1 + \eta_2 + O_7(1). \tag{274}$$

Its finite Arakelov contribution is $\log|\mathrm{N}\, v|$*; consequently the total pairwise degree is*

$$\mathscr{E}(X) + O_7(1). \tag{275}$$

Proof. Since $s = r\tau(r)\tau^2(r)$, equation (272) gives

$$I\left(\frac{\mathcal{Z}_1 \otimes \mathcal{Z}_2}{s}\right) = \frac{(\tau r)(\tau^2 r)}{(r)(\tau r)(\tau^2 r)} = (r)^{-1}.$$

This is the reciprocal remainder lattice. Its degree is $\log|\mathrm{N}\, r|$.

For the real contribution, write the prefix matrix as g. Its selected row is w, and the determinant-one identity sends $w^\sharp$ to a coordinate vector. Functoriality gives

$$\mathrm{As}(g)(Y(P) \wedge Y(w^\sharp)) = -Y(X) \wedge Y(e_1).$$

The two components in which V_1, respectively V_2, is alternating then have relative norms

$$\eta_0 + \eta_2 + O(1), \qquad \eta_0 + \eta_1 + O(1).$$

Their tensor product gives $2\eta_0 + \eta_1 + \eta_2 + O(1)$, and the finite fractional ideal supplies $\log|\mathrm{N}\, v|$. The comparison constant is uniform on the finite Rosen cross-section. □

The $O_7(1)$ in this proposition has a geometric origin. The explicit max norms in Proposition 12.4 and the smooth Hodge norms are two metrics on the same lines. Their ratio is bounded above and below on the compact Rosen cross-section. Deep in a cusp, the adapted nilpotent-orbit metric supplies the same comparison after the explicit logarithmic singularity has been removed. Thus the error does not grow with a Rosen digit.

Pointwise identification only. Proposition 12.4 gives the all-place descent, while Proposition 13.2 supplies its Hodge and integral realization at a fixed state. The two displayed components have the same weights as the pairwise McMullen functions, and their product has exactly the reciprocal remainder ideal. This does not make them a single contracting section along the orbit: the prefix row w changes with the state. Confusing equality of pointwise weights with transport of a fixed flat vector was an error in the shorter version of this manuscript. Proposition 17.46 does provide fixed flat tensors with exact transport, but their pairing has the positive divisor in (399), rather than the reciprocal divisor required here. The two statements therefore identify opposite halves of the same missing construction.

13.1 What "global" means here

There are two different descent questions which should not be confused. A metrized line over the totally real field K may have different metrics at its three real places. Such a line is a legitimate Arakelov object; permutation-invariance of those metrics is not required. It would be required if one insisted on descending the entire metrized line to $\mathbf{Q}$. Equations (261) and (262) specify all three metrics and verify their homogeneity, so no metric is left implicit.

Nevertheless, arbitrary independent metrics would prove nothing. The metrics on $\mathscr{P}_{w,P}$ are fixed before P and w are chosen: they are the Hodge metrics on the uniformizing factor and the two modular embeddings. The distinguished label is supplied by the triangle-group uniformization itself. Thus the construction is asymmetric for a geometric reason, not because the desired inequality was inserted into the norm.

14 The flat-line/record-line obstruction

The exact obstruction can be stated in two dimensions, before taking any Asai tensor. Let

$$G_j \in \mathrm{SL}_2(\mathcal{O}_K), \qquad X_j = G_j P,$$

be the Rosen prefix and the transported endpoint. Write $\ell_P = (-q, p)$ for the primitive row annihilating $P = (p, q)^t$.

Proposition 14.1 (Flat line versus record line)**.** *The transported row*

$$\ell_j = \ell_P G_j^{-1}$$

annihilates X_j, is flat along the orbit, and has the same primitive content ideal as ℓ_P. Thus it is the line to which McMullen's Busemann argument applies, but it carries no growing remainder divisor.

If $w_j = e_1^t G_j$ is the prefix row, then

$$r_j = w_j P, \qquad c_j = w_j^\sharp = (-b_j, a_j)^t, \qquad \det(P, c_j) = r_j.$$

Hence c_j carries the desired remainder divisor, but its line depends on j. The comparison between Kc_j and the endpoint line has determinant r_j at every embedding.

Proof. The identity

$$\ell_j X_j = \ell_P G_j^{-1} G_j P = 0$$

proves flatness and annihilation. Right multiplication by a unimodular matrix preserves the ideal generated by the two coordinates, so the content ideal of ℓ_j is constant. The remaining statements follow from the definitions and

$$\det \begin{pmatrix} p & -b_j \\ q & a_j \end{pmatrix} = a_j p + b_j q = w_j P.$$

□

The archimedean price of changing the flat covector can be computed exactly. This elementary calculation is the local form of the obstruction and is worth separating from the Hodge theory.

Lemma 14.2 (Angular defect of a moving Busemann center)**.** *Normalize a geodesic in* $\mathbb{H}$ *so that its forward endpoint is* 0 *and write it as* $\gamma(t) = ie^{-t}$*. For* $d \in \mathbf{R}$*, let*

$$\beta_d(z) = \log \frac{|z-d|^2}{\Im z};$$

this is the Busemann function belonging to a covector whose zero is d*. Then*

$$\frac{d}{dt}\beta_d(\gamma(t)) = \frac{d^2e^{2t}-1}{d^2e^{2t}+1}, \tag{276}$$

$$\beta_d(\gamma(t)) - \beta_d(\gamma(0)) + t = \log \frac{1+d^2e^{2t}}{1+d^2}. \tag{277}$$

In particular, the derivative is identically -1 *only for* $d = 0$*. For* $d \neq 0$*, it changes sign when* $|d|e^t = 1$*.*

Proof. Along the normalized geodesic,

$$\beta_d(\gamma(t)) = \log(d^2e^t + e^{-t}).$$

Differentiation gives (276), and subtracting the value at $t = 0$ gives (277). □

For McMullen's flat covector the center follows the actual endpoint, so $d = 0$ and there is no defect. For the prefix covector the center is the cusp approximant. After applying the same Möbius normalization, $|d|$ is the projective separation between that approximant and the endpoint; before normalization its numerator is $|\det(P, c_j)| = |r_j|$. Formula (277) therefore says more than that a uniform contraction estimate is unavailable: the exact correction needed to compare the two Busemann functions is the Green function of the incidence divisor $\det(P, c_j) = 0$. Treating it as a bounded change of metric would discard the quantity that the proof is supposed to estimate.

The same incidence has a simple global height identity. It makes precise which part of the problem is Diophantine and which part is Hodge theoretic. For primitive integral columns $P = (p, q)^t$ and $c = (-b, a)^t$, put

$$\lambda_i(P, c) = \log \frac{\max\{|\sigma_i(p)|, |\sigma_i(q)|\} \max\{|\sigma_i(a)|, |\sigma_i(b)|\}}{|\sigma_i(\det(P, c))|}.$$

This is the logarithmic chordal Green function, up to the harmless choice of max norm rather than Euclidean norm.

Proposition 14.3 (Global incidence identity)**.** *For the primitive prefix cusp* $c_j = w_j^\sharp$ *and* $r_j = \det(P, c_j) \neq 0$*,*

$$\log|\operatorname{N} r_j| = 3h(P) + 3h(c_j) - \sum_{i=0}^{2} \lambda_i(P, c_j). \tag{278}$$

Consequently the desired estimate $\log|\operatorname{N} r_j| \leq 6h(P) + B_7$ *is equivalent to*

$$\sum_{i=0}^{2} \lambda_i(P, c_j) \geq 3h(c_j) - 3h(P) - B_7. \tag{279}$$

Proof. Sum the definition of λ_i over the three real embeddings. The two max-norm sums are $3h(P)$ and $3h(c_j)$, because both columns are primitive. The remaining sum is $\log|\operatorname{N}\det(P,c_j)| = \log|\operatorname{N} r_j|$. Rearrangement proves (278); the second assertion is the same identity with terms moved to opposite sides. □

Equation (279) clarifies the role of Roth-type theorems. Such theorems ordinarily give an *upper* bound for the proximity of a high-height rational point to a fixed algebraic target. Here the required statement is a *lower* bound: at the special prefix cusps selected by a norm record, the sum of the three real incidences must absorb essentially all of the growing height of c_j. The identity is algebraic number theory. Any reason that the lower bound holds must use the correlated triangle-group orbit, and the modular embeddings are the available source of that correlation. This is the exact sense in which the remaining problem is a number-theory/Hodge-theory hybrid.

This proposition explains why the difficulty repeatedly reappears. The flat line contracts but has no norm record. The record line has the norm record but is not fixed. Transferring from one to the other costs the incidence angle

$$\frac{|r_{j,i}|}{\|P_i\|\,\|c_{j,i}\|},$$

and the product of its three numerators is $|\operatorname{N} r_j|$, the quantity to be bounded. Thus a comparison theorem which simply replaces the record line by the flat line is circular.

There is also an index mismatch at a record-producing transition. The energy immediately before the transition uses the finite ideal of r_j, whereas the maximal-slope coordinate selected after the transition is r_{j+1}. No formal property of a direct sum identifies those two lines. A valid stopping argument must either transport the r_{j+1}-line backward with a controlled incidence cost or formulate a post-record Hodge degree which directly dominates $\log|\operatorname{N} r_{j+1}|$.

> **The precise gap.** The missing theorem is not merely a bound for an already contracting Arakelov line. It is a theorem which joins two different lines: the flat endpoint line selected by Hodge contraction and the changing prefix line selected by the finite norm record.

15 What the pairwise tensor fixes, and what it does not

The component $\mathcal{Z}_0$ contains the desired outer proximity $\lambda_1+\lambda_2$. It was therefore natural to try to control $\mathcal{Z}_0/s$ directly. That route meets two problems.

First, an ordinary rational norm sees $\mathcal{Z}_0$ together with both of its Galois conjugates, and hence counts every proximity twice. Selecting only $\mathcal{Z}_0$ requires a nonparallel metric whose compatibility with the integral lattice must be proved.

Second, $\mathcal{Z}_0$ is the Hodge component with positive limiting weight. Controlling its full vector norm introduces expanding components which are not measured by the scalar incidence. Passing to an operator norm of a cotangent removes those components but also removes the incidence angle which makes the evaluated scalar small.

The tensor $\mathscr{P}_{w,P}$ changes the question. It uses $\mathcal{Z}_1$ and $\mathcal{Z}_2$, the two components whose pointwise weights match those in (18). Its degree is slightly larger than the stopped scalar, but Lemma 12.2 shows that the loss is at most $\log 2$. It therefore fixes the homogeneity and integral-lattice problems without changing the coefficient-six budget. It does not fix the transport problem in Proposition 14.1: a new prefix produces a new evaluated tensor.

16 How the Hodge metric crosses a cusp

The continuous Hodge estimate and the discrete Rosen algorithm use different clocks. This section explains the comparison needed to pass between them.

Let q be a local coordinate vanishing at the cusp. The monodromy around the cusp is unipotent; write it as e^N, where N is nilpotent. The standard adapted metric for the corresponding limiting mixed Hodge structure has the form

$$h_{\mathrm{ad}}(x+iy) = e^{xN} e^{-\frac{1}{2}(\log y)Y} h_{\mathrm{ad}}(i), \tag{280}$$

where Y completes N to an $\mathfrak{sl}_2$-triple. It is equivariant under integral translation:

$$h_{\mathrm{ad}}(\tau+n) = e^{nN} h_{\mathrm{ad}}(\tau).$$

The SL_2-orbit theorem compares the actual Hodge metric uniformly with this model high in the cusp; a convenient formulation for the present setting is given by Filip [20]. Therefore a parabolic power is accounted for by automorphy and an endpoint error bounded independently of the power. It does not create an extra $O(\log|n|)$ loss.

There is also a line-bundle version of the same comparison. The Higgs maps have the form

$$\theta_i : \mathcal{L}_i \longrightarrow \mathcal{L}_i^{-1} \otimes \Omega^1_{\log}.$$

Their ratio

$$\Xi = \frac{\theta_1 \theta_2}{\theta_0} \in H^0\big(\Omega^1_{\log} \otimes \mathcal{M}^{-2}\big) \tag{281}$$

relates the logarithmic cotangent line to the missing-central-copy line. If $L = -\log|q|$, the nilpotent-orbit asymptotics give

$$\|\mathcal{M}^{-2}\| \asymp L^{-1}, \qquad \left\|\frac{dq}{q}\right\| \asymp L, \qquad \|\Xi\| \asymp 1. \tag{282}$$

Thus Ξ does not manufacture an unbounded gain. It canonically moves one $\log L$ between the cotangent metric and the Hodge metric. This is exactly the central copy missing from the ordinary three-factor height.

For a Rosen branch with state coordinate x, the hyperbolic duration of the excursion is

$$\ell(x) = -2\log|x| = \log|\mathcal{R}'(x)|. \tag{283}$$

The same quantity, up to a bounded cross-section error, is the central depth η_0. Equations (282) and (283) therefore explain why the logarithmic singularity in (261) is the correct one. They also explain why the error in Proposition 13.2 is uniform in a large Rosen digit.

> **What is deep here.** The nilpotent-orbit comparison and the Higgs-field interpretation are Hodge theory. The identities of line bundles and the cusp powers in (282) are then direct calculations. Their role is metric transfer: they do not by themselves select a norm record.

17 The remaining degree inequality

The construction does not by itself prove Conjecture 12.6. The pairwise energy is not monotone at the discrete Rosen return times. Hyperbolic cycles explain why one should not expect monotonicity:

the Hodge decrease around a closed geodesic is compensated by unit holonomy. Such a cycle creates no new norm record.

The stopping rule has an exact Arakelov formulation. For a nonzero remainder r, let $\overline{\mathcal{N}}(r)$ be the K-line whose finite unit ball at $\mathfrak{p}$ is $r^{-1}\mathcal{O}_{K,\mathfrak{p}}$, with standard real norms. Then

$$\widehat{\deg}\overline{\mathcal{N}}(r) = \log|\,\mathrm{N}\, r|.$$

For a finite collection of such lines, give the direct sum the maximum norm at every place. The product formula proves

$$\widehat{\mu}_{\max}\left(\left(\bigoplus_{k\le j}\overline{\mathcal{N}}(r_k)\right)_{\ell^\infty}\right) = \max_{k\le j}\log|\,\mathrm{N}\, r_k|. \tag{284}$$

Indeed, a global line in the direct sum uses one global vector; it cannot choose a different coordinate at each place.

If r_{j+1} is a first record, then $\overline{\mathcal{N}}(r_{j+1})$ is a maximal-slope coordinate line. The pairwise tensor used in the pre-record cusp-transfer estimate, however, has finite ideal $(r_j)^{-1}$. Proposition 13.2 does not give a morphism from the selected r_{j+1}-line into that pre-record tensor. The shorter draft silently shifted this index.

There are two logically valid targets. One can prove the pre-record inequality (266) directly, without claiming that its finite line is selected by maximal slope. Alternatively, one can align the indices and prove the stronger post-record statement

$$\mathscr{E}(r_j, r_{j+1}) \le 6h(P) + B_{\mathrm{post}} \quad \text{whenever } r_{j+1} \text{ is a first record.} \tag{285}$$

The order of the two arguments matters: by convention a state is $X_{j+1} = (u, v) = (r_j, r_{j+1})$, and the finite term in $\mathscr{E}(u, v)$ is $\log|\,\mathrm{N}\, v|$. Because all η_i are nonnegative, (285) immediately bounds $\log|\,\mathrm{N}\, r_{j+1}|$, and its finite line is exactly the coordinate selected by (284).

Problem 17.1 (Stopped incidence transport)**.** *Prove either* (266) *or* (285)*. A Hodge-theoretic proof of the second must construct a controlled comparison from the flat endpoint line in Proposition 14.1 to the maximal-slope record line. The comparison must retain the incidence determinant rather than replace it by an operator norm.*

The local Hodge information remains relevant. Along a geodesic aimed at the zero of a fixed covector, the two pairwise functions have derivative bounded by (19); across a cusp, (281) transfers the logarithmic singularity with bounded metric error. For the prefix row w_j, however, the central Busemann function is centered at the cusp approximant $c_j = w_j^\sharp$, not at the endpoint P. Its derivative along the geodesic to P is an angle cosine, not identically -1. The accumulated angular error is the incidence between P and c_j, whose determinant is r_j.

This suggests the correct geometric object: an incidence or biextension line joining the flat endpoint section to the parabolic vanishing-cycle section on the Rosen return correspondence. Its finite divisor should be (r_j), while its archimedean Green function is the missing angular term. An arithmetic Hodge-index or Vojta inequality for that explicitly constructed line would be a meaningful conditional route. Without the incidence line, applying Schwarz–Pick to the pointwise tensor is unjustified.

There is a canonical relative cohomology class from which such a line may be constructed. Let Γ_∞ be the stabilizer of the standard cusp, and choose a primitive column $e \in \mathcal{O}_K^2$ fixed by Γ_∞. Put

$$\kappa(g) = g^{-1}e - e. \tag{286}$$

Then

$$\kappa(gh) = h^{-1}\kappa(g) + \kappa(h), \qquad \kappa|_{\Gamma_\infty} = 0.$$

It is the relative class obtained from the connecting map

$$H^0(\Gamma_\infty, K^2) \longrightarrow H^1(\Gamma, \Gamma_\infty; K^2).$$

For a prefix G_j, the moving cusp is $c_j = G_j^{-1}e = e + \kappa(G_j)$. Pairing with the fixed endpoint covector gives the exact period identity

$$r_j = \ell_P(c_j) = \ell_P(e) + \ell_P(\kappa(G_j)). \tag{287}$$

Thus all motion of the record line is carried by one cusp-framed extension class; it is not an arbitrary sequence of targets.

This suggests a precise Hodge construction. Regard the relative class in (286) as a cusp-framed extension of the trivial variation by the weight-minus-one homology variation, and pair it with the flat dual line $K\ell_P$. The resulting relative biextension should have finite evaluation (287) and archimedean metric (277). Positivity theorems for admissible biextension metrics make their boundary height-jump divisors effective [9, 66]. An application here would still require three proofs: admissibility of the cusp-framed extension, equality of its metric with the pairwise incidence metric, and a passage from height-jump effectiveness to the stopped inequality (279). These are concrete construction tasks; the cited positivity theorem alone does not supply the record estimate.

There is a more basic applicability boundary. The published height-jump theorem concerns the pullback of a semipositive biextension metric along a holomorphic map of an algebraic test curve into a normal-crossings compactification. The endpoint $[P]$ here is rational in the Betti local system, but it does not define a single-valued algebraic section of the triangle curve unless its monodromy stabilizer has already been found. Its natural domain is the horizontal real leaf in the projective flat bundle, and the record times are a discrete subset of one real geodesic. They are not a holomorphic test curve approaching a boundary divisor. Passing to the universal projective flat bundle makes the tautological line algebraic, but it does not make that horizontal leaf algebraic. Consequently the known effectiveness theorem cannot be inserted as the fourth step of the proof. What is required is a stopped real-leaf analogue whose input includes the first-record maximal slope. Establishing such an analogue for the explicit normalized Asai variation would be new mathematics, while the existing theorem remains the correct local model for the sign of its boundary term.

There is a useful limitation on this proposal. In ordinary group cohomology the cocycle $\kappa(g) = g^{-1}e - e$ is the coboundary of e. Its absolute extension class is therefore split; only the splitting framed at the cusp is nontrivial. Analytically, the same fact appears in

$$\beta_d(z) - \beta_0(z) = \log|z - d|^2 - \log|z|^2,$$

whose curvature vanishes away from the boundary incidences. Thus the relative biextension supplies the correct divisor and angular correction, but no interior curvature. It must be tensored with the nontrivial Hodge line $\mathcal{M} = \mathcal{L}_1\mathcal{L}_2\mathcal{L}_0^{-1}$, or with its Asai log-conormal realization, before height-jump positivity could have the required force. This rules out a bare biextension proof.

The Hodge factor can be identified without guesswork. Put

$$\mathcal{N} = \mathcal{L}_1 \otimes \mathcal{L}_2 \otimes \mathcal{L}_0^{-2} = \mathcal{M} \otimes \mathcal{L}_0^{-1}. \tag{288}$$

This is the line which contains two copies of the uniformizing, or central, contraction and one copy of each outer expansion. Its degree is

$$\deg \mathcal{N} = \frac{3}{28} + \frac{1}{28} - 2\frac{5}{28} = -\frac{3}{14}. \tag{289}$$

More is true than negativity of the total degree. Along a unit-speed central geodesic its logarithmic norm has derivative bounded by

$$-2 + \|DF_1\|_{\mathrm{hyp}} + \|DF_2\|_{\mathrm{hyp}} < 0$$

in the interior. The loss tends to zero only during motion arbitrarily deep in the cusp. Thus $\mathcal{N}$, rather than $\mathcal{M}$ alone, is the exact Hodge line behind the two pairwise contractions.

17.1 The rank-two pairwise bundle

The determinant line $\mathcal{N}$ still forgets one essential piece of unit geometry. Put

$$\mathcal{N}_1 = \mathcal{L}_1 \otimes \mathcal{L}_0^{-1}, \qquad \mathcal{N}_2 = \mathcal{L}_2 \otimes \mathcal{L}_0^{-1}, \qquad \mathcal{E}_{\mathrm{pair}} = \mathcal{N}_1 \oplus \mathcal{N}_2. \tag{290}$$

Then

$$\det \mathcal{E}_{\mathrm{pair}} = \mathcal{N}, \qquad \deg \mathcal{N}_1 = -\frac{1}{7}, \qquad \deg \mathcal{N}_2 = -\frac{1}{14}.$$

Each summand has the strict pointwise contraction

$$-1 + \|DF_i\|_{\mathrm{hyp}} < 0.$$

This splitting is forced by the elliptic monodromy of the whole triangle group, rather than by a choice of continued-fraction generators. Let $R = ST$, so that projectively

$$G_7 = \langle S, R \mid S^2 = R^7 = 1 \rangle \simeq C_2 * C_7.$$

The two ratios of Higgs fields are

$$\Xi_i = \frac{\theta_i}{\theta_0} \in H^0\big(C, \mathcal{N}_i^{-2}\big), \qquad i = 1, 2. \tag{291}$$

Proposition 17.2 (The two elliptic defects)**.** *With the embedding labels fixed by $\lambda_1 = 2\cos(3\pi/7)$ and $\lambda_2 = -2\cos(2\pi/7)$, one has*

$$\mathrm{div}(\Xi_1) = 2P_7, \qquad \mathrm{div}(\Xi_2) = P_7. \tag{292}$$

In particular, under the maximal-Higgs trivialization supplied by the nowhere-vanishing section θ_0, one has $\Xi = \Xi_1\Xi_2$, and the divisor $3P_7$ is the determinant of two genuinely different defects, not one scalar defect repeated three times.

Proof. Take a local coordinate at the order-seven point in which the uniformizing monodromy R acts by a primitive seventh root. In the σ_1-representation its projective rotation is the third power; the corresponding Schwarz map has local degree three. In the signed σ_2-representation the projective class of the trace is $2\cos(2\pi/7)$, so the local degree is two. The derivatives therefore vanish to orders two and one. These already exhaust the degrees

$$\deg \mathcal{N}_1^{-2} = \frac{2}{7}, \qquad \deg \mathcal{N}_2^{-2} = \frac{1}{7},$$

so there are no further zeros. Since θ_0 is maximal Higgs, it has no zero and introduces no additional divisor. □

The three cyclic Galois components have an even sharper global description. Read indices modulo three and put

$$\mathcal{Q}_j = \mathcal{L}_{j+1} \otimes \mathcal{L}_{j+2} \otimes \mathcal{L}_j^{-2}.$$

Thus $\mathcal{Q}_0 = \mathcal{N}$, while the degrees of $\mathcal{Q}_0, \mathcal{Q}_1, \mathcal{Q}_2$ are respectively $-3/14, 3/14, 0$. Define the meromorphic cyclic Higgs ratios

$$\Upsilon_j = \frac{\theta_{j+1}\theta_{j+2}}{\theta_j^2} \in H^0_{\mathrm{mer}}(C, \mathcal{Q}_j^{-2}). \tag{293}$$

Proposition 17.3 (The cyclic Higgs divisor ledger)**.** *The three divisors and their product are*

$$\operatorname{div}(\Upsilon_0) = 3P_7, \qquad \operatorname{div}(\Upsilon_1) = -3P_7, \qquad \operatorname{div}(\Upsilon_2) = 0, \qquad \Upsilon_0\Upsilon_1\Upsilon_2 = 1. \tag{294}$$

Consequently

$$\mathcal{Q}_0^2 \simeq \mathcal{O}(-3P_7), \qquad \mathcal{Q}_1^2 \simeq \mathcal{O}(3P_7), \qquad \mathcal{Q}_2^2 \simeq \mathcal{O}. \tag{295}$$

Proof. The three Higgs divisors are $0, 2P_7, P_7$. Taking $D_{j+1} + D_{j+2} - 2D_j$ gives respectively $3P_7, -3P_7, 0$. Multiplication of the three displayed ratios cancels every Higgs field. The line-bundle identities follow from the divisors of the meromorphic sections. □

Proposition 17.4 (The transverse minor is a Hodge Wronskian)**.** *Let Q_m, Q_n be two balanced integral return vectors and put $\Delta_{m,n} = \det(Q_m, Q_n)$. On a simply connected chart of the triangle curve, choose a nonzero local section ω_i of $\mathcal{L}_i$ and a logarithmic tangent vector ∂. Use the polarization to regard $\sigma_i(Q_m)$ and $\sigma_i(Q_n)$ as flat covectors $\ell_{m,i}, \ell_{n,i}$, and define the first-jet determinant*

$$W_i^{m,n} = \det \begin{pmatrix} \ell_{m,i}(\omega_i) & \ell_{n,i}(\omega_i) \\ \ell_{m,i}(\nabla_\partial \omega_i) & \ell_{n,i}(\nabla_\partial \omega_i) \end{pmatrix}. \tag{296}$$

After fixing the determinant-one polarization consistently, one has the exact factorization

$$W_i^{m,n} = \sigma_i(\Delta_{m,n})\, \vartheta_i, \qquad \vartheta_i = \omega_i \wedge \nabla_\partial \omega_i. \tag{297}$$

Here ϑ_i is the scalar form of the Higgs field θ_i in the chosen frame. In particular, when $\Delta_{m,n} \neq 0$,

$$\frac{W_1^{m,n} W_2^{m,n}}{(W_0^{m,n})^2} = \frac{\sigma_1(\Delta_{m,n})\sigma_2(\Delta_{m,n})}{\sigma_0(\Delta_{m,n})^2}\, \Upsilon_0 = \frac{\mathrm{N}(\Delta_{m,n})}{\sigma_0(\Delta_{m,n})^3}\, \Upsilon_0. \tag{298}$$

Writing $f_{r,i} = \ell_{r,i}(\omega_i)$, the normalized transverse jet is the logarithmic divided difference

$$J_i^{m,n} = d \log \frac{f_{n,i}}{f_{m,i}} = \frac{W_i^{m,n}}{f_{m,i} f_{n,i}}. \tag{299}$$

Its numerator is $\sigma_i(\Delta_{m,n})\vartheta_i$, while its two poles are the rational incidence divisors $f_{m,i} = 0$ and $f_{n,i} = 0$. Taking the cyclic quotient gives the two-frame Green–Wronskian factorization

$$\frac{J_1^{m,n} J_2^{m,n}}{(J_0^{m,n})^2} = \frac{\mathrm{N}(\Delta_{m,n})}{\sigma_0(\Delta_{m,n})^3}\, \Upsilon_0 \prod_{r \in \{m,n\}} \frac{f_{r,0}^2}{f_{r,1} f_{r,2}}. \tag{300}$$

The central Wronskian vanishes exactly when the two frame cosets agree. The two outer Wronskians have, in addition, vanishing orders 2 and 1 at P_7. Thus their cyclic product has the effective Hodge defect $3P_7$.

Proof. For a two-dimensional vector space, two covectors u, v, and two vectors a, b, elementary alternating algebra gives

$$\det \begin{pmatrix} u(a) & v(a) \\ u(b) & v(b) \end{pmatrix} = (u \wedge v)(a \wedge b).$$

Apply this with $a = \omega_i$ and $b = \nabla_\partial \omega_i$. Flat transport and the chosen polarization give $(\ell_{m,i} \wedge \ell_{n,i}) = \sigma_i(\Delta_{m,n})$, proving (297). Rescaling ω_i or ∂ changes both sides by the same factor, so the assertion is independent of the displayed trivialization. Taking the cyclic ratio and using the definition of Υ_0 proves the first equality in (298); the product formula identity $\sigma_0(\Delta)\sigma_1(\Delta)\sigma_2(\Delta) = \mathrm{N}(\Delta)$ proves the second. The uniformizing Higgs field θ_0 is nowhere zero, while Proposition 17.2 gives the stated outer vanishing orders. Finally, Proposition 8.32 identifies $\Delta_{m,n} = 0$ with equality of the frame cosets. Finally, (299) is the elementary identity $d\log(f_n/f_m) = (f_m\, df_n - f_n\, df_m)/(f_m f_n)$. Dividing (298) by the six period values appearing in (299) proves (300). □

Proposition 17.5 (The Wronskian has fixed denominator type)**.** *For $k = 1, 2, 3$, put*

$$a_k = \frac{3}{4} - \frac{k}{14}, \qquad b_k = \frac{1}{4} - \frac{k}{14}.$$

The corresponding period equation is

$$t(1-t)y'' + \left[1 - \left(2 - \frac{k}{7}\right) t\right] y' - a_k b_k y = 0. \tag{301}$$

If $\mathcal{W}_k$ is the Wronskian of a Frobenius basis, then

$$\mathcal{W}_k(t) = c_k\, t^{-1}(1-t)^{k/7-1}, \qquad c_k \neq 0. \tag{302}$$

The Galois order used in this paper is $k = 1, 3, 2$ for $i = 0, 1, 2$. Consequently

$$\frac{\mathcal{W}_3 \mathcal{W}_2}{\mathcal{W}_1^2} = c(1-t)^{3/7} = c\, u^3, \qquad u^7 = 1 - t, \tag{303}$$

for a fixed nonzero normalization constant c. In particular, on the order-seven orbifold chart the determinant line of the first-jet evaluation has fixed algebraic denominator type. For the two integral frames its only return-dependent finite factor is the principal ideal $(\Delta_{m,n})$, up to a fixed ideal coming from the choice of de Rham lattice.

Proof. After division by $t(1-t)$, the coefficient of y' in (301) is

$$p_k(t) = \frac{1}{t} + \frac{-1 + k/7}{1-t}.$$

Abel's identity $\mathcal{W}_k' = -p_k \mathcal{W}_k$ gives (302). Substitution of $k = 3, 2, 1$ gives

$$(3/7 - 1) + (2/7 - 1) - 2(1/7 - 1) = 3/7,$$

which proves (303). A Frobenius basis can be normalized over the coefficient field of the differential equation; changing that basis multiplies c_k by a fixed determinant. Combining this fixed Wronskian with (297) shows that the varying determinant of the integral evaluation lattice is precisely $\Delta_{m,n}$. Any primes introduced by the chosen integral de Rham model form one fixed finite set and contribute a bounded term independent of the return word. □

Proposition 17.6 (Exact coefficient denominators of the three periods)**.** *Let*

$$F_k(t) = {}_2F_1(a_k, b_k; 1; t) = \sum_{n\geq 0} c_{k,n} t^n \qquad (k = 1, 2, 3),$$

with a_k, b_k as in Proposition 17.5. The denominator of $c_{k,n}$ is supported only at 2 and 7, and

$$\begin{aligned} v_2(\operatorname{den} c_{k,n}) &= 4n + 2v_2(n!), \\ v_7(\operatorname{den} c_{k,n}) &= 2n + 2v_7(n!). \end{aligned} \tag{304}$$

Consequently, for

$$C_{\text{hyp}} = 2^6\, 7^3 = 21952, \tag{305}$$

one has

$$F_k(C_{\text{hyp}} x) \in \mathbb{Z}[[x]] \quad \textit{for all } k.$$

Among positive integral changes $t = Cx$, the two exponents in (305) are minimal.

Proof. For $k = 1, 2, 3$, the two numerator parameters, with denominator 28, are respectively

$$(19, 5), \qquad (17, 3), \qquad (15, 1).$$

Thus

$$c_{k,n} = \frac{\prod_{j=0}^{n-1} (r_k + 28j)(s_k + 28j)}{28^{2n} (n!)^2}.$$

Every numerator factor is a unit at 2 and at 7, which gives (304). If $p \nmid 28$, then $a_k, b_k \in \mathbb{Z}_p$ and $(a_k)_n/n!, (b_k)_n/n! \in \mathbb{Z}_p$, so no other prime occurs. Legendre's formula gives

$$\frac{v_2(n!)}{n} \longrightarrow 1, \qquad \frac{v_7(n!)}{n} \longrightarrow \frac{1}{6}.$$

Hence an integral scaling requires $v_2(C) \geq 6$ and $v_7(C) \geq 3$; those exponents also plainly suffice in (304). □

> **What the denominator computation does and does not do.** The Picard–Fuchs solutions are globally bounded G-functions after one fixed change of coordinate. Therefore the denominator term in an arithmetic-holonomy argument is explicit and independent of the return word. This does not create a free positive margin. Under $t = C_{\text{hyp}} x$, the singular point $t = 1$ moves to $x = C_{\text{hyp}}^{-1}$; the gain in coefficient integrality is paid for by the smaller analytic radius. The holonomy slope is designed to combine precisely these two quantities. Any useful strict surplus must therefore come from relative overconvergence of a stopped combination, not from a better normalization of the three ordinary Frobenius series.

Proposition 17.7 (Two-frame Schwarz contraction)**.** *Let $\mathcal{G}_{m,n} \subset \mathbb{H}$ be the geodesic with endpoints $x_{m,0}$ and $x_{n,0}$. At a point $z \in \mathcal{G}_{m,n}$, the normalized Wronskians satisfy*

$$\left| \frac{J_i^{m,n}(z)}{J_0^{m,n}(z)} \right| \leq \|DF_i(z)\|_{\text{hyp}} < 1, \qquad i = 1, 2. \tag{306}$$

Consequently

$$\left| \frac{J_1^{m,n}(z) J_2^{m,n}(z)}{(J_0^{m,n}(z))^2} \right| \leq \|DF_1(z)\|_{\text{hyp}}\, \|DF_2(z)\|_{\text{hyp}} < 1. \tag{307}$$

Proof. The quotient $R_s = f_{n,s}/f_{m,s}$ is a real Möbius transformation of the s-th period coordinate. After changing an irrelevant real sign, R_0 and R_i map their respective half-planes to $\mathbb{H}$. Put

$$H_i = R_i \circ F_i \circ R_0^{-1}.$$

If $z \in \mathcal{G}_{m,n}$, then $\zeta = R_0(z)$ lies on the imaginary axis. By (299),

$$\frac{J_i^{m,n}}{J_0^{m,n}}(z) = \frac{\zeta H_i'(\zeta)}{H_i(\zeta)}.$$

Since $|\zeta| = \operatorname{Im}\zeta$ and $|H_i(\zeta)| \geq \operatorname{Im} H_i(\zeta)$, Schwarz–Pick gives

$$\left|\frac{\zeta H_i'(\zeta)}{H_i(\zeta)}\right| \leq \frac{\operatorname{Im}\zeta\,|H_i'(\zeta)|}{\operatorname{Im} H_i(\zeta)} = \|DH_i(\zeta)\|_{\mathrm{hyp}} = \|DF_i(z)\|_{\mathrm{hyp}}.$$

The last equality uses invariance of the hyperbolic derivative under real Möbius transformations. The outer modular embeddings are not automorphisms, so Schwarz–Pick is strict. Multiplication proves (307). □

There is a second way to use the same two flat covectors. The Wronskian above compares their values and first derivatives at one time. The next construction instead compares their values at two different times. Its advantage is that no choice of a moving affine splitting appears.

Proposition 17.8 (The two-time period determinant)**.** *Let $s < t$ lie on a lifted central geodesic, and use Gauss–Manin parallel transport along that segment to identify the flat fibres. At the i-th real place choose polarization-normalized vectors $\widehat{\omega}_i(s)$ and $\widehat{\omega}_i(t)$ spanning the two Hodge lines. For $\Delta_{m,n} \neq 0$, put*

$$\mathcal{E}_i^{m,n}(s,t) = \det\begin{pmatrix} \ell_{m,i}(\widehat{\omega}_i(s)) & \ell_{n,i}(\widehat{\omega}_i(s)) \\ \ell_{m,i}(\widehat{\omega}_i(t)) & \ell_{n,i}(\widehat{\omega}_i(t)) \end{pmatrix}. \tag{308}$$

If

$$d_i(s,t) = d_{\mathbb{H}}\big(Z_i(s), Z_i(t)\big)$$

is the distance between the two Hodge lines in the i-th period coordinate, then

$$\mathcal{E}_i^{m,n}(s,t) = \sigma_i(\Delta_{m,n})\big(\widehat{\omega}_i(s) \wedge \widehat{\omega}_i(t)\big), \tag{309}$$

$$\frac{|\mathcal{E}_i^{m,n}(s,t)|}{|\sigma_i(\Delta_{m,n})|} = 2\sinh\frac{d_i(s,t)}{2}. \tag{310}$$

Define the two-time cyclic Hodge separation by

$$\mathfrak{T}(s,t) = 2\log\frac{|\mathcal{E}_0^{m,n}(s,t)|}{|\sigma_0(\Delta_{m,n})|} - \sum_{i=1}^{2}\log\frac{|\mathcal{E}_i^{m,n}(s,t)|}{|\sigma_i(\Delta_{m,n})|}. \tag{311}$$

If the central geodesic is parametrized by arclength, then, in the extended real sense when an outer determinant vanishes,

$$\mathfrak{T}(s,t) \geq \frac{1}{2}\mathfrak{C}(s,t). \tag{312}$$

Thus this two-time determinant retains a fixed positive proportion of the complete integrated Schwarz–Pick defect without replacing the actual orbit segment by the rational geodesic joining $[Q_m]$ and $[Q_n]$.

Proof. For covectors u, v and vectors a, b, alternating algebra gives

$$\det \begin{pmatrix} u(a) & v(a) \\ u(b) & v(b) \end{pmatrix} = (u \wedge v)(a \wedge b).$$

The first factor is $\sigma_i(\Delta_{m,n})$, which proves (309). In a determinant-one period frame, a normalized lift of $z \in \mathbb{H}$ is $(z,1)/\sqrt{\operatorname{Im} z}$. Hence

$$\left|\det\left(\frac{(z,1)}{\sqrt{\operatorname{Im} z}}, \frac{(w,1)}{\sqrt{\operatorname{Im} w}}\right)\right| = \frac{|z-w|}{\sqrt{\operatorname{Im} z \operatorname{Im} w}} = 2\sinh\frac{d_{\mathbb{H}}(z,w)}{2},$$

proving (310).

Put $L = t - s$, and let

$$L_i = \int_s^t \|DF_i(\gamma(u))\|_{\mathrm{hyp}}\, du.$$

Then $d_0 = L$ and $d_i \le L_i$. For $\phi(x) = \log(2\sinh(x/2))$, one has $\phi'(x) = \frac{1}{2}\coth(x/2) \ge \frac{1}{2}$. Therefore

$$\begin{aligned} \mathfrak{T}(s,t) &= \sum_{i=1}^{2} \big(\phi(L) - \phi(d_i)\big) \\ &\ge \frac{1}{2}\sum_{i=1}^{2}(L - d_i) \ge \frac{1}{2}(2L - L_1 - L_2) = \frac{1}{2}\mathfrak{C}(s,t). \end{aligned}$$

If an outer distance is zero, the corresponding term on the first line is $+\infty$, and the assertion follows by continuity. □

The proposition solves the real two-path comparison for this particular kernel. Its finite-place audit explains why it does not yet solve the classification.

Proposition 17.9 (The cyclic two-time determinant has no radial degree)**.** *Let τ generate $\mathrm{Gal}(K/\mathbb{Q})$, with the central factor indexed by zero. The algebraic coefficient factors of the three determinants in (309) are*

$$\Delta_{m,n}, \qquad \tau(\Delta_{m,n}), \qquad \tau^2(\Delta_{m,n}).$$

Consequently the coefficient of the cyclic quotient

$$\frac{(\mathcal{E}_0^{m,n})^2}{\mathcal{E}_1^{m,n}\mathcal{E}_2^{m,n}} \tag{313}$$

is

$$c(\Delta) = \frac{\Delta^2}{\tau(\Delta)\tau^2(\Delta)} = \frac{\Delta^3}{\mathrm{N}(\Delta)}, \qquad |\mathrm{N}(c(\Delta))| = 1. \tag{314}$$

The cyclic quotient therefore records the unit shape of Δ, but not $\log|\mathrm{N}\,\Delta|$. Conversely, the full product of the three coefficient determinants is $\mathrm{N}(\Delta)$, so it records the radial divisor once, but its Hodge factor is

$$\prod_{i=0}^{2} 2\sinh\frac{d_i(s,t)}{2}, \tag{315}$$

which is expanding rather than the cyclic defect in (312).

Proof. The three coefficient factors are immediate from (309). Their cyclic quotient is the first expression in (314); the second uses the product formula. Taking its field norm gives one. Multiplying the three factorizations instead gives $\mathrm{N}(\Delta)$ and (315). □

The coefficient identity $c(\Delta) = \Delta^3/\mathrm{N}(\Delta)$ suggests multiplying the cyclic quotient by $\mathrm{N}(\Delta)$ and taking a cube root. The fundamental Hodge relation $\mathcal{N} \simeq \mathcal{N}_2^{\otimes 3}$ below shows that the underlying Hodge *line* does have such a cube root. The two-time *section* does not. This distinction is another concrete form of the fusion obstruction.

Proposition 17.10 (The two-time cyclic section is not a cube)**.** *On a local product of two copies of the triangle curve, let*

$$K_i(u,v) = \omega_i(u) \wedge \omega_i(v)$$

denote the two-point determinant of local holomorphic Hodge lifts. The cyclic section

$$K_{\mathrm{cyc}}(u,v) = \frac{K_0(u,v)^2}{K_1(u,v)K_2(u,v)} \tag{316}$$

lives in

$$p_1^*(\mathcal{N}^{-1}) \otimes p_2^*(\mathcal{N}^{-1}) \simeq \left(p_1^*(\mathcal{N}_2^{-1}) \otimes p_2^*(\mathcal{N}_2^{-1})\right)^{\otimes 3}.$$

Thus its line is a cube. Nevertheless K_{cyc} is not locally the cube of a meromorphic section.

More explicitly, on the orbifold cover near (P_7, P_7), choose coordinates u, v in which the three period maps have local degrees $1, 3, 2$. Up to holomorphic units,

$$K_0 = u - v, \qquad K_1 = u^3 - v^3, \qquad K_2 = u^2 - v^2, \tag{317}$$

and hence

$$K_{\mathrm{cyc}}(u,v) = \frac{\text{unit}}{(u+v)(u^2+uv+v^2)}. \tag{318}$$

The three off-diagonal coincidence branches occur with simple poles. The divisor of K_{cyc} is therefore not divisible by three.

Proof. The line identity follows from

$$\frac{(p_1^*\mathcal{L}_0 \otimes p_2^*\mathcal{L}_0)^2}{(p_1^*\mathcal{L}_1 \otimes p_2^*\mathcal{L}_1) \otimes (p_1^*\mathcal{L}_2 \otimes p_2^*\mathcal{L}_2)} = p_1^*(\mathcal{N}^{-1}) \otimes p_2^*(\mathcal{N}^{-1})$$

and Corollary 17.31. At the order-seven point, the local degrees of the uniformizing, σ_1-, and σ_2-period maps are respectively $1, 3, 2$. The wedge of their two Hodge lines is, up to a unit, the difference of the corresponding period coordinates. This proves (317). Factoring

$$(u^3 - v^3)(u^2 - v^2) = (u-v)^2(u+v)(u^2+uv+v^2)$$

gives (318). A meromorphic cube has all divisor multiplicities divisible by three, whereas the three displayed branches have multiplicity -1. □

The noncube divisor separates canonically into the two pairwise channels:

$$K_{\mathrm{cyc}} = \left(\frac{K_0}{K_2}\right)\left(\frac{K_0}{K_1}\right). \tag{319}$$

Near (P_7, P_7), these factors are, up to units,

$$\frac{1}{u+v}, \qquad \frac{1}{u^2+uv+v^2}.$$

On restriction to the diagonal their leading terms are the inverses of the two Higgs ratios Ξ_2 and Ξ_1; there the relation $\Xi_1 = \Xi_2^2$ is recovered. Away from the diagonal, the second denominator is not the square of the first. Thus the infinitesimal one-Hodge-line description and the stopped two-time description are genuinely different: the latter naturally retains two coordinates.

The two local denominators are nevertheless coprime in the projective tangent direction. The following calculation determines exactly what that observation gives, and why it does not globalize with bounded coefficients.

Proposition 17.11 (Local Koszul fusion and the outer-recurrence obstruction)**.** *In the orbifold coordinates of Proposition 17.10, put*

$$L(u,v) = u+v, \qquad Q(u,v) = u^2+uv+v^2.$$

Then

$$(L,Q) = (u+v, u^2) = (u+v, v^2) \quad \text{in } \mathbb{C}\{u,v\}. \tag{320}$$

Consequently the common zero scheme of the two pairwise divided differences is supported at (P_7, P_7) and has local intersection multiplicity two. Equivalently,

$$\mathrm{Res}_r(r+1, r^2+r+1) = 1, \tag{321}$$

so there is no common off-diagonal tangent branch.

There is no corresponding global Bézout estimate with uniformly bounded coefficients on the diagonal return groupoid. More precisely, there are a point $z \in \mathbb{H}$, integers $n_j \to \infty$, and points $w_j = A_3^{n_j} z$ such that

$$\begin{aligned} d_{\mathbb{H}}(z, w_j) &\longrightarrow \infty, \\ d_{\mathbb{H}}(F_i(z), F_i(w_j)) &\longrightarrow 0 \qquad (i = 1,2), \end{aligned} \tag{322}$$

where

$$A_3 = \begin{pmatrix} \alpha & \alpha \\ \lambda & \alpha \end{pmatrix}, \qquad \alpha = \lambda^2 - 1.$$

If

$$\mathscr{D}_i(z,w) = \frac{2\sinh(d_{\mathbb{H}}(F_i(z), F_i(w))/2)}{2\sinh(d_{\mathbb{H}}(z,w)/2)}, \qquad i = 1,2, \tag{323}$$

is the invariant norm of the pairwise divided difference K_i/K_0, then

$$\mathscr{D}_1(z, w_j),\ \mathscr{D}_2(z, w_j) \longrightarrow 0. \tag{324}$$

Hence any global identity which expresses 1 as a linear combination of the two divided differences must have coefficient norm tending to infinity along this sequence.

Proof. The elementary identities

$$Q - uL = v^2, \qquad Q - vL = u^2$$

show that (L, Q) contains both u^2 and v^2. Conversely, substitution of $v = -u$ gives $Q = u^2$, proving (320). The quotient is $\mathbb{C}\{u\}/(u^2)$, of length two. The resultant statement is the same calculation after setting $r = u/v$: the quadratic takes the value one at the root $r = -1$ of the linear factor.

Proposition 7.3 proves that A_3 is hyperbolic in the distinguished factor and elliptic of infinite order in both outer factors. Its two outer powers lie in a compact product of elliptic circles. Some sequence of positive powers therefore converges simultaneously to the identity. Equivariance gives

$$F_i(A_3^{n_j} z) = \rho_i(A_3)^{n_j} F_i(z) \longrightarrow F_i(z),$$

whereas central hyperbolicity gives $d_{\mathbb{H}}(z, A_3^{n_j} z) = n_j \ell_0(A_3) + O_z(1)$. This proves (322).

Equation (310), applied without the flat covectors, says precisely that the norm of K_i/K_0 is (323). Its numerator tends to zero and its denominator tends to infinity, proving (324). A linear combination with two bounded coefficient norms would then tend to zero and cannot equal one. □

> **What the resultant calculation settles.** Locally, the two off-diagonal branches do not share a hidden divisor: their Koszul intersection is the length-two scheme at the order-seven point. This is useful divisor bookkeeping. Globally, joint outer recurrence produces approximate double points arbitrarily far out in the central direction. Therefore the local Bézout coefficients acquire an unbounded metric when transported around the thin group. Controlling that coefficient growth is exactly the stopped affine/unit-height problem; the resultant does not bypass it. Notice also that u, v are complex period parameters, not algebraic integers, so (321) is not a finite-prime coprimality statement about $\Delta_{m,n}$.

The preceding bad sequence is generated by powers of one fixed element. That special case can be measured much more precisely than mere compactness of the outer factors suggests. The calculation also separates a genuine fixed-word phenomenon, which occurs for every odd triangle parameter, from the changing-word phenomenon detected by the $p = 11$ experiment.

Proposition 17.12 (Fixed-word recurrence is only polynomially deep)**.** *Let $h \in \mathrm{SL}_2(K)$ be hyperbolic at the distinguished embedding and elliptic of infinite order at both outer embeddings. Write $t = \mathrm{tr}(h)$, and define*

$$C_0(t) = 0, \qquad C_1(t) = 1, \qquad C_{n+1}(t) = tC_n(t) - C_{n-1}(t).$$

If $\ell_0(h) > 0$ is the central translation length, then

$$h^n = C_n(t)h - C_{n-1}(t)I, \tag{325}$$

$$\log\left|\mathrm{N}_{K/\mathbf{Q}}\, C_n(t)\right| = \frac{n\ell_0(h)}{2} + O_h(\log n). \tag{326}$$

In particular, for any column $Q \in K^2$ with $\det(Q, hQ) \neq 0$,

$$\log\left|\mathrm{N}\det(Q, h^n Q)\right| = \frac{n\ell_0(h)}{2} + O_{h,Q}(\log n). \tag{327}$$

Let $z \in \mathbb{H}$. If neither $F_1(z)$ nor $F_2(z)$ is the elliptic fixed point of the corresponding outer image of h, then

$$\mathfrak{T}(z, h^n z) = n\ell_0(h) + O_{h,z}(\log n), \tag{328}$$

where $\mathfrak{T}$ is defined by (311). Equivalently,

$$\log \mathscr{D}_i(z, h^n z) = -\frac{n\ell_0(h)}{2} + O_{h,z}(\log n), \qquad i = 1, 2, \tag{329}$$

along any simultaneous outer near-return. Thus the coefficient norm in any Bézout globalization of the two divided differences must pay at least $n\ell_0(h)/2 - O_{h,z}(\log n)$. *The unboundedness in Proposition 17.11 is therefore exponential in central length, up to a polynomial factor.*

Proof. Cayley–Hamilton gives (325). More explicitly, $C_n(t) = U_{n-1}(t/2)$, where U_j is the Chebyshev polynomial of the second kind. At the central embedding,

$$|C_n(\sigma_0 t)| = \frac{\sinh(n\ell_0(h)/2)}{\sinh(\ell_0(h)/2)}.$$

At an outer embedding write $\sigma_i(t) = 2\cos\theta_i$. Then

$$C_n(\sigma_i t) = \frac{\sin(n\theta_i)}{\sin\theta_i}. \tag{330}$$

The algebraic number $e^{2i\theta_i}$ is not a root of unity, because the outer element has infinite order. The Baker–Wüstholz lower bound for a nonzero linear form in logarithms [55] gives constants $c_i, C_i > 0$, depending only on h, such that

$$|e^{2in\theta_i} - 1| \geq c_i n^{-C_i}.$$

Consequently each outer logarithm in (330) is $O_h(\log n)$, both above and below. Adding the three embeddings proves (326). Alternation and (325) give the exact identity

$$\det(Q, h^n Q) = C_n(t)\det(Q, hQ),$$

which proves (327).

Central hyperbolicity gives $d(z, h^n z) = n\ell_0(h) + O_{h,z}(1)$. In an outer disk coordinate centered at the elliptic fixed point, the displacement of a point not at the center is comparable, with constants depending on that point, to $|\sin(n\theta_i)|$. The same Baker lower bound shows that the logarithm of this displacement is $O_{h,z}(\log n)$, both above and below. Substitute these estimates into (311) and (323). This proves (328) and (329). Finally, if $a_1\mathscr{D}_1 + a_2\mathscr{D}_2 = 1$, then at least one coefficient has norm at least $1/(2\max_i \mathscr{D}_i)$, giving the last assertion. □

The Chebyshev–Baker part is not cubic. If K' is any totally real field and $h \in \mathrm{SL}_2(K')$ is central-hyperbolic and elliptic of infinite order at every other real embedding, the same place-by-place proof gives

$$\log\left|\mathrm{N}_{K'/\mathbf{Q}}\, C_n(\operatorname{tr} h)\right| = \frac{n\ell_0(h)}{2} + O_h(\log n).$$

There are four bounded-rotation terms when $K' = K_{11}$ instead of two, but each contributes only $O_h(\log n)$. Fixed-word recurrence therefore has the same asymptotic form at $p = 7$ and $p = 11$.

There is a complementary boundary statement. It explains exactly how much of a long repeated block ordinary Roth theory can exclude.

Proposition 17.13 (A repeated hyperbolic block is Roth-critical)**.** *Let h satisfy the hypotheses of Proposition 17.12, and let ξ^+, ξ^- be its attracting and repelling fixed points at the distinguished embedding. Then*

$$[K(\xi^+) : K] = [K(\xi^-) : K] = 2.$$

Let $C \subset \mathbb{P}^1(\mathbb{R}) \setminus \{\xi^-\}$ be compact, let $z_n \in C \cap \mathbb{P}^1(K)$, and put $y_n = h^n z_n$. If the y_n are distinct, then for every $\epsilon > 0$,

$$n\ell_0(h) \leq (2+\epsilon)h_K(y_n) + O_{h,C,\epsilon}(1), \qquad h_K = [K : \mathbb{Q}]h_{\mathrm{abs}}. \tag{331}$$

Consequently a fixed hyperbolic word cannot be repeated for a length which is supercritical relative to the height of the resulting K-rational endpoint. The critical scale with constant 2 is not excluded.

Proof. The fixed points are the roots of a quadratic over K whose discriminant is t^2-4. If this discriminant were a square in the totally real field K, then its image at an outer embedding would be a square in $\mathbb{R}$. It is instead negative because the outer trace has absolute value less than two. Thus the fixed points have degree two over K.

North–south dynamics, uniformly for $z_n \in C$, gives

$$|h^n z_n - \xi^+| \ll_{h,C} e^{-n\ell_0(h)}.$$

Roth's theorem over K, at the distinguished real place and with the height normalized as in the statement, gives

$$-\log|y_n - \xi^+| \le (2+\epsilon)h_K(y_n) + O_{h,C,\epsilon}(1)$$

for distinct $y_n \in K$. Combining the two inequalities proves (331). □

What this removes, and what it leaves. The simultaneous outer recurrence used above is not exponentially accurate in its power n: Baker bounds it below by a power of n. Roth then rules out a rational Rosen tail which shadows the attracting quadratic fixed point for a supercritical repeated block. This disposes of the most coherent cyclic version of the obstruction. Two channels remain. First, a block at the exact H^{-2} scale is compatible with Roth, as anticipated by the continued-fraction constant calculation. Second, a general Rosen orbit uses changing words h_j, so its attracting quadratic targets and their heights also change. Neither the fixed-target Roth theorem nor the fixed-h Baker estimate is uniform in that moving target. The missing stopped inequality is precisely the statement that consecutive thin-orbit selection supplies such uniformity.

Why a formal cube root is not enough. Multiplying (313) by $\mathrm{N}(\Delta)$ makes its algebraic coefficient Δ^3. Taking one third of its logarithmic norm along a real segment is therefore numerically possible. Proposition 17.10 proves that this operation does not produce a holomorphic section with finite ideal (Δ): it has nontrivial branching around the outer coincidence divisors. Keeping those branches separately is precisely a vector-valued construction.

What the two-time calculation changes. The missing real transport can be captured without a moving Busemann centre: the determinant of period evaluations at the two endpoint times retains at least one half of the integrated Hodge defect. The remaining failure is now completely adelic. The contracting cyclic quotient has norm-one finite coefficient, while the product carrying the determinant ideal has the wrong real sign. A proof must fuse those two pieces in one vector-valued, one-divisor metric. This is the same structural demand as Problem 17.17, but the two-time formula gives an alternative analytic carrier in which the path comparison itself is already exact.

Problem 17.14 (Two-time one-divisor inequality)**.** *For consecutive reduced returns before recurrence, prove that there are $\eta_P > 0$ and C_P such that*

$$\log|\mathrm{N}\,\Delta_{m,n}| + \eta_P \mathfrak{T}(t_m, t_n) \le C_P, \tag{332}$$

unless the orbit has terminated or $\Delta_{m,n} = 0$. By (312), this inequality implies the heptagonal classification. Proposition 17.9 shows that it is not a formal consequence of either scalar factor of the two-time determinant; it is precisely a fusion statement between its radial and cyclic components.

What this advances. The algebraic covariant and its rank-drop locus are no longer hypothetical. They are the first Gauss–Manin Wronskians (296), and (298) joins the complete arithmetic shape of the transverse minor to the global Hodge divisor $3P_7$. After division by the two incidence values, (300) puts on the same line the norm ideal, the Hodge divisor, and

the two rational-frame Green potentials; it is the promised coupling of two successive states. Proposition 17.5 also removes return-word-dependent denominators from the determinant line. What is not yet proved is the uniform stopped comparison between this now explicit integral determinant line and its Hodge norm. That archimedean slope comparison, rather than another coordinate calculation, is the remaining Arakelov estimate.

The remaining all-place endpoint wall. Proposition 17.7 alone could not be integrated on $\mathcal{G}_{m,n}$, because that rational geodesic is not the stopped orbit segment. Proposition 8.12 now removes this ambiguity: the dynamically correct future–past pair accumulates the full defect on the actual path. Corollary 8.14 then compares it with the two-future integral Wronskian by the single rational shear $-A_{m,n}/\Delta_{m,n}$. Thus there is no longer a missing shadowing statement.

What remains is to convert the signed real boundary term of that shear into a proper all-place height. The denominator ideal is already $(\Delta_{m,n})$, and the numerator is the second integral coordinate $A_{m,n}$. A unit can nevertheless move size among the three real embeddings without changing either ideal. The required endpoint theorem must show that persistent absorption of $\mathfrak{C}(m,n)$ by this unit shape forces one of the two incidence coordinates to vanish, or forces recurrence of the affine frame. An arithmetic-holonomy overconvergence theorem is one possible way to prove that endpoint assertion, but overconvergence is no longer being used as a name for an unidentified path comparison.

Problem 17.15 (Stopped Wronskian slope inequality)**.** *Let $\gamma(s)$ be a noncuspidal ray and let $t_m < t_n$ be compact return times before the first repetition of a rational frame coset, and let $\mathfrak{C}(m,n)$ be the integrated Hodge defect (163). Prove that there are constants $\eta_P > 0$ and C_P, independent of m, n, such that*

$$\mathcal{P}_{m,n} \geq 3h([g_m P]) + 3h([g_n P]) + \eta_P \mathfrak{C}(m,n) - C_P. \tag{333}$$

The metrics and the integral lattice in this inequality are those of the Wronskians in (296); in particular, the finite part is the principal ideal $(\Delta_{m,n})$, not a separately chosen collection of local corrections.

The same assertion can now be stated without mentioning two different paths. For $\Delta_{m,n} \neq 0$, use $U = R_m^{\mathrm{dyn}}$ and $V = \widetilde{R}_{m,n}$ from Corollary 8.14. The function $\mathfrak{B}_{m,n}(t)$ was defined in (177); it is the time-dependent form of (167) and, by Proposition 8.16, its endpoint difference is the Galois-weighted cross ratio (179).

Problem 17.16 (All-place affine endpoint absorption)**.** *At the stopped intervals used in Problem 17.15, prove that there are $\eta_P > 0$ and C_P such that*

$$\mathfrak{B}_{m,n}(t_n) - \mathfrak{B}_{m,n}(t_m) \leq (1 - \eta_P)\mathfrak{C}(m,n) + C_P \tag{334}$$

unless $A_{m,n}\Delta_{m,n} = 0$ or a rational affine frame has already recurred. Equivalently, prove the alignment inequality (189). A bound at only one real embedding is not sufficient: the three endpoint lines are conjugates of the same pair $(a_{m,n}, b_{m,n})$. The maximal slope of the Galois completion of (203) is one possible tool, not an extra hypothesis. By Proposition 8.16, the left side is exactly $2\log|\mathrm{CR}_0| - \log|\mathrm{CR}_1| - \log|\mathrm{CR}_2|$.

The factor $1 - \eta_P$ is essential in this scalar formulation. A bounded endpoint term would be stronger and is not needed: the endpoint may absorb part of the Hodge curvature. Proposition 8.20 gives a different formulation in which no signed outer alignment estimate is requested.

Problem 17.17 (Rank-two stopped fusion)**.** *At the stopped intervals used above, and before repetition of a rational affine frame, prove that there are $\eta_P > 0$ and C_P such that*

$$\log|\,\mathrm{N}\,\Delta_{m,n}| + \eta_P \max\left\{X_0, \mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_n) - \mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_m)\right\} \le C_P \tag{335}$$

unless $A_{m,n}\Delta_{m,n} = 0$.

The two entries of the maximum must be treated as the two coordinates of the Galois-completed rank-two extension (203). In particular, the finite lattice of the rational splitting $V = Q_n/\Delta_{m,n}$ is to be inserted once, not once for each scalar entry. The word "slope" here means a stopped, Hodge-selected *slope, including the metric transport from t_m to t_n. It is not the ordinary static maximal Arakelov slope of the underlying two-column module. Proposition 17.20 below computes the latter and shows that it has the wrong symmetry.*

This is a sharply specified use of a stopped metric. It is not the claim that a singular-value gap controls every rational line, nor that an ordinary Harder–Narasimhan maximum chooses the desired cyclic character. The analytic input selecting a large coordinate has already been proved in (192); the missing statement is that the stopped transport on the one-divisor cyclic pair retains that selection after the three real places have been assembled. The condition that the interval quotients belong to the consecutive pointed orbit (176) is indispensable, because (335) is false for arbitrary primitive pairs in the ambient arithmetic group.

The first coordinate is completely algebraic. Proposition 17.34 constructs a norm-one CM-torus point $t_{m,n}$ for which recurrence is $t_{m,n} = 1$ and the three-place Green proximity is $\sum_i \log(2/|1 - t_{m,n,i}|)$. Thus the fusion problem is not missing an adelic model for the visual term. Its difficulty is to couple that torus coordinate to the controlled Wronskian coordinate along one consecutive thin orbit; ordinary Roth treats the former at the critical scale and does not perform this coupling.

Why the ordinary two-column slope is insufficient. The module

$$M_{m,n} = \mathcal{O}_K Q_m + \mathcal{O}_K Q_n \subset \mathcal{O}_K^2$$

has determinant ideal $(\Delta_{m,n})$, so it is the first natural candidate for the promised rank-two lattice. Its determinant degree, however, gives only the symmetric Green sum $\sum_i \Lambda_i$, equivalently the proximity $\mathcal{P}_{m,n}$. It does not give the cyclic difference $2r_0 - r_1 - r_2$ occurring in (190). Conversely, if one takes the projective line of the rational splitting $Q_n/\Delta_{m,n}$ and then saturates it, multiplication by $\Delta_{m,n}$ returns the line of Q_n; the record denominator has disappeared by the product formula. Thus the usual maximal-slope inequality for $M_{m,n}$ either retains the determinant with symmetric metrics or retains the cyclic direction after forgetting its affine framing. Problem 17.17 asks for the one object which retains both. This is why maximal slope is a description of the required mechanism, not a theorem which can simply be cited here.

Proposition 17.18 (The stopped one-divisor extension is explicit)**.** *Assume $\Delta_{m,n} \neq 0$, put $q = Q_m$, $v = Q_n$, and set*

$$M_{m,n} = \mathcal{O}_K q + \mathcal{O}_K v \subset \mathcal{O}_K^2.$$

Determinant with q gives an exact sequence of projective $\mathcal{O}_K$-modules

$$0 \longrightarrow \mathcal{O}_K q \longrightarrow M_{m,n} \xrightarrow{w \mapsto \det(q,w)} (\Delta_{m,n}) \longrightarrow 0. \tag{336}$$

Tensoring by the tautological line $\mathcal{O}_K q$ gives

$$0 \longrightarrow (\mathcal{O}_K q)^{\otimes 2} \longrightarrow (\mathcal{O}_K q) \otimes M_{m,n} \longrightarrow (\mathcal{O}_K q) \otimes (\Delta_{m,n}) \longrightarrow 0. \tag{337}$$

After identifying the last ideal with $\mathcal{O}_K q$ by division by $\Delta_{m,n}$, the rational lift of its distinguished generator is exactly

$$q \otimes \frac{v}{\Delta_{m,n}}. \tag{338}$$

Thus (337) is the integral model of the universal affine extension (203) for the stopped pair. Its quotient ideal is $(\Delta_{m,n})$, once and only once. Changing either balanced column by a unit gives an isomorphic framed sequence.

Proof. Every element of $M_{m,n}$ has the form $aq+bv$. Its determinant with q is $b\Delta_{m,n}$, so the image is $(\Delta_{m,n})$. Because $\Delta_{m,n} \neq 0$, the determinant vanishes exactly when $b = 0$, proving (336). Tensoring gives (337). The element $v/\Delta_{m,n}$ has determinant one with q, so it is precisely the rational splitting used in (168). Finally, unit rescaling acts by invertible changes on the subline and quotient ideal and hence preserves the isomorphism class of the framed sequence. □

Proposition 17.19 (Maximal slope is simultaneous approximation of three Hodge splittings)**.** *Use the stopped extension (337), put*

$$Q = Q_m, qquad V = Q_n/\Delta_{m,n}, qquad e_0 = Q \otimes Q, qquad e_1 = Q \otimes V,$$

and fix a time t. Every rational splitting line different from the subline Ke_0 is generated by

$$e(s) = e_1 + se_0 = Q \otimes (V + sQ), qquad s \in K. \tag{339}$$

At the i-th real place, set

$$\zeta_i(t) = -\frac{f_{V,i}(t)}{f_{Q,i}(t)} \in \mathbf{C}. \tag{340}$$

Then, with the determinant-one Hodge normalization,

$$\|e(s)\|_{i,t} = \|Q\|_{i,t} \frac{|f_{Q,i}(t)|}{\sqrt{\operatorname{Im} Z_i(t)}} |\sigma_i(s) - \zeta_i(t)|. \tag{341}$$

Relative to the integral generators $e_0 = Q \otimes Q$ and $Q \otimes Q_n$, the finite content of $e(s)$ is the fractional ideal

$$\mathfrak{c}_{m,n}(s) = (s, \Delta_{m,n}^{-1}). \tag{342}$$

Consequently the ordinary maximal Arakelov slope of the framed stopped extension is the supremum, over one $s \in K$ and the subline Ke_0, of the global degrees obtained by combining (341) at the three real places with (342). In particular, it is an adelic simultaneous-approximation problem for the three moving targets $\zeta_i(t)$, with the determinant ideal entered once.

Proof. A splitting of the last arrow in the generic fibre of (337) differs from the splitting e_1 by a unique element of its subline Ke_0, proving (339). The Hodge norm on a tensor product is multiplicative, and period evaluation is linear. Hence

$$f_{V+sQ,i} = f_{V,i} + \sigma_i(s) f_{Q,i} = f_{Q,i}\big(\sigma_i(s) - \zeta_i(t)\big),$$

which gives (341). In the integral basis $(e_0, Q \otimes Q_n)$, the two coefficients of $e(s)$ are s and $\Delta_{m,n}^{-1}$. This proves (342). The final statement is the definition of maximal slope as the largest degree of a saturated one-dimensional subspace. □

The preceding proposition identifies all rational splitting lines, but it does not yet identify the metric needed in the fusion problem. The next calculation makes the distinction exact. If a nonzero vector x in an integral module has coefficient ideal $\mathfrak{c}(x)$, then the saturated lattice on Kx is $\mathfrak{c}(x)^{-1}x$. With the convention

$$\widehat{\deg}_t(Kx) = \log \mathrm{N}\,\mathfrak{c}(x) - \sum_{i=0}^{2} \log \|x\|_{i,t}, \tag{343}$$

the expression is unchanged when x is multiplied by an element of $K^\times$.

Proposition 17.20 (The ordinary splitting degree has the wrong symmetry)**.** *For the splitting vector $e(s)$ in (339), put $W(s) = V + sQ$. Then*

$$\widehat{\deg}_t(Ke(s)) = \log \mathrm{N}(s, \Delta_{m,n}^{-1}) - \sum_{i=0}^{2} \Lambda\big(Z_i(t); [Q_i], [W(s)_i]\big). \tag{344}$$

For the dynamic splitting $s = -A_{m,n}/\Delta_{m,n}$, this becomes

$$\widehat{\deg}_t(Ke(s)) = -\log|\,\mathrm{N}\,\Delta_{m,n}| - \sum_{i=0}^{2} \Lambda\big(Z_i(t); [Q_{m,i}], [R^{\mathrm{dyn}}_{m,i}]\big) + O_P(1). \tag{345}$$

In particular, the ordinary maximal slope of (337) sees a symmetric sum of the three Green functions. Its variation does not equal the cyclic quantity $2\Delta r_0 - \Delta r_1 - \Delta r_2$ in (190). Thus ordinary maximal slope does not prove, or even reformulate, Problem 17.17.

Proof. The coefficient ideal of $e(s)$ is (s, Δ^{-1}) by (342). Moreover $\det(Q, W(s)) = 1$. Multiplicativity of the tensor norm and Lemma 8.21 therefore give

$$\log \|e(s)\|_{i,t} = \log \|Q_i\|_{i,t} + \log \|W(s)_i\|_{i,t} = \Lambda(Z_i(t); [Q_i], [W(s)_i]).$$

Substitution in (343) proves (344). For the dynamic splitting, Corollary 17.24 gives $W(s) = R^{\mathrm{dyn}}_m$ and coefficient ideal $\Delta^{-1}\mathcal{O}_K$, up to the fixed ideal δ_P^{-1}. This proves (345). The last assertion follows because every term in that equation occurs with coefficient one, whereas the controlled potential has coefficients $(2, -1, -1)$. □

Proposition 17.21 (The exact two-splitting angle identity)**.** *Keep the notation of Proposition 17.19, and let $s, s' \in K$ be distinct. Put*

$$\mathfrak{c}(s) = (s, \Delta_{m,n}^{-1}), \qquad \mathfrak{J}(s, s') = \left(\frac{s - s'}{\Delta_{m,n}}\right) \mathfrak{c}(s)^{-1}\mathfrak{c}(s')^{-1}.$$

Then $\mathfrak{J}(s, s')$ is an integral ideal. If $\theta_i(s, s'; t) \in [0, \pi/2]$ is the Hermitian angle between $W(s)_i$ and $W(s')_i$ at the i-th real place, then

$$\begin{aligned} &\widehat{\deg}_t(Ke(s)) + \widehat{\deg}_t(Ke(s')) - \widehat{\deg}_t(Ke_0) \\ &\qquad = -\log|\,\mathrm{N}\,\Delta_{m,n}| - \log \mathrm{N}\,\mathfrak{J}(s, s') + \sum_{i=0}^{2} \log \sin \theta_i(s, s'; t). \end{aligned} \tag{346}$$

In particular,

$$\widehat{\deg}_t(Ke(s)) + \widehat{\deg}_t(Ke(s')) - \widehat{\deg}_t(Ke_0) \le -\log|\,\mathrm{N}\,\Delta_{m,n}|. \tag{347}$$

Thus a static slope argument can make at most one rational splitting exceptional relative to the subline. In the stopped problem the second line does not have to be guessed: it is the controlled splitting $s' = 0$, paired with the dynamic splitting $s = -A_{m,n}/\Delta_{m,n}$. The remaining issue is to show that the two analytic channels in (192) make the sum of these two stopped *relative degrees large. The identity then forces one of the two incidence coordinates to vanish. Applying Minkowski's theorem twice to a single line cannot supply that stopped angle comparison.*

Proof. In the K-basis (Q, Q_n) of the stopped module, the coordinate vectors of $W(s) = Q_n/\Delta + sQ$ and $W(s')$ are (s, Δ^{-1}) and (s', Δ^{-1}). Their coordinate determinant is

$$\frac{s - s'}{\Delta}.$$

It belongs to the product of their coefficient ideals $\mathfrak{c}(s)\mathfrak{c}(s')$. This is exactly the assertion that $\mathfrak{J}(s, s')$ is integral.

At the i-th real place the determinant-one Hodge metric gives the area identity

$$|\sigma_i(s - s')| = \|W(s)_i\|_{i,t}\, \|W(s')_i\|_{i,t} \sin\theta_i(s, s'; t).$$

Apply the degree convention (343) to the two vectors and sum over the real places. The product formula and

$$(s - s') = (\Delta)\, \mathfrak{c}(s)\mathfrak{c}(s')\mathfrak{J}(s, s')$$

give

$$\widehat{\deg}_t(KW(s)) + \widehat{\deg}_t(KW(s')) = -\log|\,\mathrm{N}\,\Delta| - \log \mathrm{N}\,\mathfrak{J}(s, s') + \sum_i \log\sin\theta_i.$$

Tensoring both lines by KQ adds twice the degree of KQ, while $Ke_0 = K(Q \otimes Q)$ has exactly twice that degree. This proves (346). Finally $\mathrm{N}\,\mathfrak{J} \geq 1$ and $\sin\theta_i \leq 1$, proving (347). □

Corollary 17.22 (The canonical dynamic–controlled angle identity)**.** *Work in an integral determinant-one P-frame, so that $(A_{m,n}, \Delta_{m,n}) = \mathcal{O}_K$. Put*

$$s_{\mathrm{dyn}} = -A_{m,n}/\Delta_{m,n}, \qquad s_{\mathrm{ctl}} = 0.$$

If $A_{m,n}\Delta_{m,n} \neq 0$, then the two splitting lines are distinct,

$$\mathfrak{c}(s_{\mathrm{dyn}}) = \mathfrak{c}(s_{\mathrm{ctl}}) = \Delta_{m,n}^{-1}\mathcal{O}_K, \qquad \mathfrak{J}(s_{\mathrm{dyn}}, s_{\mathrm{ctl}}) = (A_{m,n}),$$

and

$$\begin{aligned}\widehat{\deg}_t(Ke(s_{\mathrm{dyn}})) + \widehat{\deg}_t(Ke(s_{\mathrm{ctl}})) - \widehat{\deg}_t(Ke_0)& \\ = -\log|\,\mathrm{N}(A_{m,n}\Delta_{m,n})| + \sum_{i=0}^{2} \log\sin\theta_i(s_{\mathrm{dyn}}, 0; t).&\end{aligned} \tag{348}$$

In particular, its variation over a stopped interval is exactly the angle boundary term

$$\begin{aligned}\Big[\widehat{\deg}_t(Ke(s_{\mathrm{dyn}})) + \widehat{\deg}_t(Ke(s_{\mathrm{ctl}})) - \widehat{\deg}_t(Ke_0)\Big]_{t_m}^{t_n}& \\ = \sum_{i=0}^{2} \log\frac{\sin\theta_i(s_{\mathrm{dyn}}, 0; t_n)}{\sin\theta_i(s_{\mathrm{dyn}}, 0; t_m)}.&\end{aligned} \tag{349}$$

Therefore the unresolved stopped fusion estimate is precisely an angle-transport assertion for two canonical lines, not an unspecified search for an auxiliary Arakelov subspace.

Proof. The content of the controlled splitting is $(0, \Delta^{-1}) = \Delta^{-1}\mathcal{O}_K$. Since $(A, \Delta) = \mathcal{O}_K$,

$$(-A/\Delta, \Delta^{-1}) = \Delta^{-1}(-A, 1) = \Delta^{-1}\mathcal{O}_K.$$

Substitution in the definition of $\mathfrak{J}$ gives

$$\mathfrak{J}(-A/\Delta, 0) = (-A/\Delta^2)(\Delta)(\Delta) = (-A).$$

Now apply Proposition 17.21. The finite terms are independent of t, so subtraction at the two endpoints gives (349). □

The static two-line shortcut is rejected. It is tempting to lower-bound the left side of (348) by $-\log|\,\mathrm{N}\,A_{m,n}|$ plus a positive fraction of the Hodge defect. The exact identity shows that this proposed inequality is simply stopped fusion made stronger by the extra nonpositive term $\sum_i \log\sin\theta_i$; it supplies no independent mechanism. More fundamentally, (349) says that ordinary degree variation is only a boundary-angle variation. It contains none of the interior Schwarz–Pick curvature. The required positive term must come from the Hodge twist and its stopped Bott–Chern metric on the one-divisor cyclic pair below. The canonical pair identifies the endpoint wall exactly, but does not cross it.

There is nevertheless an exact integral carrier for the cyclic character. It is not the preceding rank-two module by itself. Let τ generate $\mathrm{Gal}(K/\mathbf{Q})$, numbered so that $\sigma_i = \sigma_0 \circ \tau^i$, and put

$$\mathbb{W}_{\mathrm{cyc}} = \mathrm{Sym}^2 V \otimes \tau V \otimes \tau^2 V.$$

For this calculation write $q = Q_m$, $v = Q_n$, $\Delta = \det(q, v)$, and $V' = v/\Delta$.

Proposition 17.23 (The one-divisor cyclic Asai pair)**.** *Define two decomposable vectors in* $\mathbb{W}_{\mathrm{cyc}}$ *by*

$$\begin{aligned} U(q,v) &= q^{\otimes 2} \otimes \tau(V') \otimes \tau^2(V'), \\ W(q,v) &= (V')^{\otimes 2} \otimes \tau(q) \otimes \tau^2(q). \end{aligned} \tag{350}$$

At the distinguished Hodge metric one has the exact identity

$$\log \frac{\|W(q,v)\|_t}{\|U(q,v)\|_t} = \mathfrak{A}^{\mathrm{ctl}}_{m,n}(t). \tag{351}$$

If q, v are primitive, their coefficient ideals are

$$\mathfrak{c}(U) = (\tau\Delta\,\tau^2\Delta)^{-1}, \qquad \mathfrak{c}(W) = \Delta^{-2}, \qquad \mathrm{N}\,\mathfrak{c}(U) = \mathrm{N}\,\mathfrak{c}(W) = |\,\mathrm{N}\,\Delta|^{-2}. \tag{352}$$

After the common rescaling

$$\widetilde{U} = \Delta U, \qquad \widetilde{W} = \Delta W, \tag{353}$$

the norm ratio remains (351), while

$$\mathfrak{c}(\widetilde{U}) = \frac{\Delta}{\tau\Delta\,\tau^2\Delta}\mathcal{O}_K, \qquad \mathfrak{c}(\widetilde{W}) = \Delta^{-1}\mathcal{O}_K, \qquad \mathrm{N}\,\mathfrak{c}(\widetilde{U}) = \mathrm{N}\,\mathfrak{c}(\widetilde{W}) = |\,\mathrm{N}\,\Delta|^{-1}. \tag{354}$$

Thus $(\widetilde{U}, \widetilde{W})$ is an explicit two-coordinate integral object in which the controlled channel is a norm ratio and the finite determinant divisor is shared once.

The static global degree still has no strict sign. At the three arithmetic embeddings the three logarithmic norm ratios are the cyclic characters

$$2r_i - r_{i+1} - r_{i+2} \qquad (i \bmod 3),$$

whose sum is zero. Since the two content ideals in (354) have equal norm, the global Arakelov degrees of the two coordinate lines are equal. The positive term in Problem 17.17 must therefore come from stopped Hodge transport—equivalently, from the common $3P_7$ Green–Higgs divisor and its Bott–Chern endpoint term—and not from the ordinary static Harder–Narasimhan polygon.

Proof. At the distinguished embedding, the two tensor norms are

$$\|U\| = \|q_0\|^2\|V_1'\|\|V_2'\|, \qquad \|W\| = \|V_0'\|^2\|q_1\|\|q_2\|.$$

Their logarithmic ratio is $2\log(\|V_0'\|/\|q_0\|) - \sum_{i=1}^{2}\log(\|V_i'\|/\|q_i\|)$, which is (190); the factors $\sqrt{\operatorname{Im} Z_i(t)}$ cancel in each quotient. This proves (351).

Primitivity of q, v and multiplicativity of tensor contents give the first two ideals in (352). Since $\Delta\tau\Delta\tau^2\Delta = \mathrm{N}\,\Delta \in \mathbf{Q}$, both have norm $|\mathrm{N}\,\Delta|^{-2}$. Multiplication by Δ gives the two ideals in (354); their norms are both $|\mathrm{N}\,\Delta|^{-1}$. Finally, applying τ cyclically permutes which Hodge factor is doubled. Summing the three displayed characters gives zero, and the finite degrees agree by (354). □

Corollary 17.24 (The dynamic splitting uses the whole determinant denominator)**.** *For*

$$s = s_{m,n} = -A_{m,n}/\Delta_{m,n},$$

one has

$$e(s) = Q_m \otimes R_m^{\mathrm{dyn}}, \tag{355}$$

$$\mathfrak{c}_{m,n}(s) = \Delta_{m,n}^{-1}\mathcal{O}_K \quad \textit{up to the fixed cusp-normalization ideal } (\delta_P)^{-1}. \tag{356}$$

At the real places,

$$|\sigma_i(s) - \zeta_i(t)| = \left|\frac{f_{R_m^{\mathrm{dyn}},i}(t)}{f_{Q_m,i}(t)}\right|. \tag{357}$$

Thus the same global splitting line carries the complete finite divisor and the actual-path Hodge potential. Proposition 17.20 explains why this fact alone does not turn that potential into an ordinary Arakelov slope.

Proof. Equation (169) gives $V + sQ = R_m^{\mathrm{dyn}}$, proving (355). Moreover,

$$(s, \Delta^{-1}) = \Delta^{-1}(-A, 1).$$

If $A \in \mathcal{O}_K$, the last ideal is exactly $\Delta^{-1}\mathcal{O}_K$; in the cusp-normalized frame $A \in \delta_P^{-1}\mathcal{O}_K$, it differs by an ideal depending only on P. Finally, the definition (340) and $f_{V+sQ} = f_{R_m^{\mathrm{dyn}}}$ give (357). □

Proposition 17.25 (The stopped splitting is an Arakelov ray-class lift)**.** *Work in an integral determinant-one P-frame and write the first column of an interval quotient as $(A, \Delta)^t$, with $\Delta \neq 0$. Then $(A, \Delta) = \mathcal{O}_K$, so*

$$[\bar{A}] \in \mathscr{R}_{(\Delta)} = (\mathcal{O}_K/(\Delta))^\times / \operatorname{im}(\mathcal{O}_K^\times)$$

is defined. Changing the integral splitting replaces A by $A + t\Delta$, while changing its unit normalization multiplies A by a global unit. Hence the displayed class is intrinsic to the framed extension. Since K has class number one, the usual ray-class exact sequence identifies $\mathscr{R}_{(\Delta)}$ with the finite ray class group of modulus (Δ).

For the dynamic splitting $s = -A/\Delta$, the pair consisting of s and its finite content $(s, \Delta^{-1}) = \Delta^{-1}\mathcal{O}_K$ is an adelic representative of this finite ray-class coordinate, while the three quantities

$$|\sigma_i(s) - \zeta_i(t)|$$

in (341) are its archimedean coordinates in the stopped Hodge metric. Thus the one-divisor stopped extension is an Arakelov lift of the ray-Levi class, not merely a rational approximation to three unrelated numbers.

Proof. The two entries of the first column of a determinant-one integral matrix generate the unit ideal, proving that $\bar{A}$ is invertible modulo (Δ). The two changes of frame have the stated effects, which are exactly the equivalence relations defining the quotient $\mathscr{R}_{(\Delta)}$. The standard ray-class exact sequence is

$$\mathcal{O}_K^\times \longrightarrow (\mathcal{O}_K/(\Delta))^\times \longrightarrow \mathrm{Cl}_{(\Delta)}(K) \longrightarrow \mathrm{Cl}(K) \longrightarrow 1.$$

The last group is trivial here. The claims about finite content and Hodge norm are respectively (356) and (341). □

The closest classical number-theoretic structure. Arakelov class groups and number-field infrastructure combine an ideal with its logarithmic unit-shape coordinate; reduced Arakelov divisors give controlled representatives of that compact extension [19]. The proposition shows that this is genuinely the right number-theoretic comparison. It does not prove the stopped inequality. Classical infrastructure treats a fixed Arakelov class (or a fixed modulus) and permits arbitrary reduction by units. Here the modulus $(\Delta_{m,n})$ moves with the interval, its residue class is the affine coordinate of the same consecutive thin word, and the metric is transported by the Picard–Fuchs system. The missing statement is therefore a Hodge-selected, varying-ray-class non-drift theorem. This explains precisely in what sense the problem is a number-theory–Hodge hybrid.

The exact Roth comparison. If the three numbers $\zeta_i(t)$ were fixed algebraic targets, Proposition 17.19 would put the problem in the scope of the adelic Subspace Theorem, with the usual two-dimensional critical exponent. They are instead period ratios of the double-heptagon Picard–Fuchs system and move with the stopped return. Their motion is exactly what contributes the positive Hodge defect. Thus a stronger numerical exponent in ordinary Roth is not the missing input. A sufficient Roth-type theorem must be relative and holonomic: it must compare the change of the targets (340) with the conformal/Hodge charge of their common differential system, while retaining the finite content (342). The arithmetic-holonomy route is aimed at precisely this statement.

What remains after this construction. Neither the rank-two lattice, its filtration, its quotient ideal, nor its rational splitting is now conjectural. Equip (337) at the three real places with the Hodge metrics and retain its full Galois orbit. The Green coordinate measures the relative size of its two columns, while Proposition 17.70 measures the shape of the rational splitting. Problem 17.17 is exactly the assertion that the *stopped transport metric* captures the larger of those two quantities along consecutive thin-orbit intervals. The static metric is now known not to suffice. No additional algebraic lattice remains to be invented; the outstanding construction is the Bott–Chern/stopped metric comparison on the one-divisor cyclic pair.

Proposition 17.26 (Rank-two fusion is sufficient)**.** *Problem 17.17 implies Conjecture 1.1.*

Proof. Suppose that $[P]$ is noncuspidal and that no rational affine frame recurs. Choose stopped intervals containing arbitrarily many compact returns. Strict Schwarz–Pick on a fixed compact flow box makes $\mathfrak{C}(m,n)$ unbounded. Proposition 8.20 gives

$$\max\left\{X_0, \mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_n) - \mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_m)\right\} \geq \frac{1}{3}\mathfrak{C}(m,n).$$

If $A_{m,n} = 0$, the endpoint is a cusp by (171); if $\Delta_{m,n} = 0$, the rational Borel frame recurs. Otherwise (335) bounds the displayed maximum uniformly, contradicting the preceding inequality. Hence a frame recurs, and (210) together with Proposition 7.1 gives the cusp or special hyperbolic alternative. □

Proposition 17.27 (Endpoint absorption is the remaining part of the Wronskian inequality)**.** *Problem 17.16 implies Problem 17.15, with the cusp incidence $A_{m,n} = 0$ and the recurrence incidence $\Delta_{m,n} = 0$ removed as the two equality alternatives.*

Proof. Apply Proposition 8.12 to the actual dynamic pair $(Q_m, R_m^{\mathrm{dyn}})$. Corollary 8.14 identifies the normalized jet of $(Q_m, \widetilde{R}_{m,n})$ with the two-future jet $J_i^{m,n}$. Proposition 8.13 says that subtracting the endpoint difference $\mathfrak{B}_{m,n}(t_n) - \mathfrak{B}_{m,n}(t_m)$ is the entire change in passing from the dynamic cyclic metric to that two-future metric. Therefore (334) leaves at least $\eta_P \mathfrak{C}(m,n) - C_P$ in the two-future Wronskian norm. Finally, the Green–Wronskian factorization (300), evaluated in the fixed compact return charts, converts this metric gain into (333). Its finite determinant is exactly $(\Delta_{m,n})$; no additional denominator is introduced. The two excluded vanishing cases are the incidences described after (171). □

Proposition 17.28 (The stopped Wronskian inequality closes the proof)**.** *Problem 17.15 implies Conjecture 1.1.*

Proof. The integrand in (163) is strictly positive by Schwarz–Pick. On a compact flow box meeting the chosen return section it has a positive minimum, and every compact return spends a uniform positive time in that box. Hence, if the ray is noncuspidal and has no repeated frame coset, one can choose $m < n$ for which $\mathfrak{C}(m,n)$ exceeds the constant in (333).

For such a pair, $\Delta_{m,n} \neq 0$. The product-formula identity (212) gives

$$\mathcal{P}_{m,n} \leq 3h([g_m P]) + 3h([g_n P]),$$

because $\Delta_{m,n}$ is a nonzero algebraic integer. This contradicts (333). Therefore a frame coset repeats. Equation (210) then gives a projectively nonidentity element of $G_7 \cap B_P(K)$; Proposition 7.1 gives the cusp or special hyperbolic alternative. □

Why arithmetic holonomy is relevant, and what it must prove. The periods $\ell(\omega_i)$ satisfy the Picard–Fuchs equation of the double-heptagon family, and their first derivatives are again holonomic. Thus Proposition 17.4 supplies the kind of fixed finite-dimensional first-jet system to which arithmetic holonomy methods apply. The arithmetic holonomy theorem of Calegari–Dimitrov–Tang packages coefficient denominators and archimedean analytic continuation in a single Bost slope inequality [57]. This is structurally closer to (333) than a stronger irrationality exponent for one scalar target.

It is not yet a black-box application. Proposition 17.5 verifies at determinant level that the logarithmic first-jet lattice has no word-dependent denominator beyond $(\Delta_{m,n})$. What remains is to prove, uniformly over the return groupoid, that the resulting rank-two slope sees one of the two coordinates selected by Proposition 8.20. Their determinant denominator is common and must be charged once. The third condition which appeared in the earlier holonomy proposal—a description of the rank-drop locus—is now settled by (297): away from P_7, rank drops exactly when the frame cosets agree. The zeros at P_7 have total order three and are exactly the Hodge charge which the slope calculation is meant to retain. This leaves the concrete fusion estimate (335). The scalar cross-ratio estimate (334) remains a sufficient alternative, rather than an unspecified appeal to Roth's theorem.

The rank-two holonomy inequality which would suffice. For a return frame Q_r, let

$$f_{r,i}(t) = \ell_{r,i}(\omega_i(t)).$$

The two functions $f_{m,i}, f_{n,i}$ lie in the same fixed rank-two Picard–Fuchs module; only their integral coefficient vectors vary. Equation (297) proves

$$\dim_{K(t)}\langle f_{m,i}, f_{n,i}\rangle = 2 \quad \Longleftrightarrow \quad \Delta_{m,n} \neq 0. \tag{358}$$

This is the rank alternative which a holonomy argument would have to turn into a quantitative contradiction. The standard theorem does not perform that step merely because the two coefficient vectors came from a long thin-group word.

Proposition 17.29 (The direct holonomy bound is coefficient-blind)**.** *Fix one of the three rank-two Picard–Fuchs modules and an admissible continuation map φ. In the arithmetic holonomy dimension bound, the denominator term $\tau(\mathbf{b})$ and the boundary energy $\mathscr{G}(\varphi)$ depend on the fixed differential module and on φ, but not on the integral coefficient matrix expressing $f_{m,i}, f_{n,i}$ in a fixed basis. If $\Delta_{m,n} \neq 0$, the pair is a basis of the same module. Therefore the ordinary holonomy theorem necessarily gives*

$$\frac{\mathscr{G}(\varphi)}{\log|\varphi'(0)| - \tau(\mathbf{b})} \geq 2 \tag{359}$$

whenever its denominator is positive. In particular, applying the standard theorem to the two moving period combinations cannot by itself force $\Delta_{m,n} = 0$.

Proof. The holonomy theorem bounds the $K(t)$-dimension of the space of functions having a prescribed coefficient-denominator type and a common meromorphic continuation through φ [58, Theorem 2.5.1]. An integral change of basis does not worsen that denominator type, and every solution of the fixed Picard–Fuchs equation has the same continuation domain after pullback to the simply connected source of φ. When $\Delta_{m,n} \neq 0$, equation (358) supplies two independent functions satisfying those same hypotheses. Their dimension is two, so the right side of the dimension bound must be at least two. The coefficient matrix, and hence the fact that it is a consecutive thin-orbit matrix, never enters the displayed quotient. □

To see the quantitative target, write $\mathscr{G}(\varphi)$ for the boundary-energy numerator in the arithmetic holonomy dimension bound and $\tau(\mathbf{b})$ for the denominator cost of the two-row first-jet lattice. In rank two that bound has the shape

$$2 \leq \frac{\mathscr{G}(\varphi)}{\log|\varphi'(0)| - \tau(\mathbf{b})} \tag{360}$$

whenever the two functions are independent. Hence the strict reverse inequality

$$2\big(\log|\varphi'(0)| - \tau(\mathbf{b})\big) > \mathscr{G}(\varphi)$$

could force $\Delta_{m,n} = 0$, but Proposition 17.29 proves that the unmodified hypotheses can never yield this strict reverse inequality. The all-place extension sought here must use the three continuations belonging to the same K-system and the cyclic Wronskian metric. The ambient differential system is fixed, not moving; the word changes the two integral sections and the continuation branch. Proposition 8.20 supplies the positive analytic quantity before holonomy is invoked. A holonomy proof must identify its boundary energy with the maximal slope of the one framed two-section lattice and thereby prove (335). This is narrower than constructing a word-dependent normal function and explains exactly how the holonomy dimension theorem could yield the classification.

Noncircularity check. A long triangle-group word is only monodromy of the same rank-two Picard–Fuchs system. It does not by itself enlarge the conformal radius of that system. Moreover, all moving functions $f_{r,i}$ already lie in its two-dimensional solution space, so the unmodified holonomy rank bound merely permits the known rank two. The strict reverse inequality must come from *relative overconvergence of the stopped rational two-section lattice*: an incidence must cancel the first analytic obstruction, or the maximal-slope line must acquire the central Green gain. Proposition 8.20 proves that one of those alternatives carries at least $\mathfrak{C}/3$; it does not prove that the arithmetic holonomy inequality captures that alternative. One may not obtain the last assertion simply by assigning a continuation map to a long word.

Proposition 17.30 (The neutral cyclic line is trivial)**.** *In fact the last identity in* (295) *has a square root:*

$$\mathcal{Q}_2 = \mathcal{L}_0 \otimes \mathcal{L}_1 \otimes \mathcal{L}_2^{-2} \simeq \mathcal{O}_C \tag{361}$$

as an orbifold line on the compactified triangle curve. Its Hodge metric is uniformly comparable with the constant metric on the right. Hence the neutral Hodge character

$$\Theta_{\mathrm{neut}}^{\mathrm{Hod}} = \Theta_1^{\mathrm{Hod}} - 2\Theta_2^{\mathrm{Hod}} \tag{362}$$

is a bounded coboundary on the reduced return groupoid.

Proof. An orbifold line on the compactification of a curve of signature $(2, 7, \infty)$ is determined by its coarse degree and by its characters at the points of orders two and seven. The degree of $\mathcal{Q}_2$ is

$$\frac{5}{28} + \frac{1}{28} - 2\frac{3}{28} = 0.$$

At the order-two point the three rank-two systems have the same local monodromy, because S has rational entries, so the character $1 + 1 - 2$ is zero. At the order-seven point their projective rotation labels are $1, 3, 2$; the character of $\mathcal{Q}_2$ is therefore $1 + 3 - 2 \cdot 2 = 0 \pmod 7$. The logarithmic residues at the cusp also cancel. Thus both orbifold characters and the coarse degree vanish, which proves (361).

The Hodge norm of a local tautological frame of $\mathcal{L}_i$ is, up to a fixed normalization, $(\operatorname{Im} F_i(z))^{1/2}$, with $F_0(z) = z$. Consequently the squared norm on $\mathcal{Q}_2$ is comparable near the cusp to

$$\frac{\operatorname{Im} z \, \operatorname{Im} F_1(z)}{(\operatorname{Im} F_2(z))^2}.$$

The signed cusp expansions $\operatorname{Im} F_i(z) = (|\lambda_i|/\lambda_0) \operatorname{Im} z + O(1)$ show that this ratio tends to a positive constant. It is smooth and positive on the remaining compact orbifold. The Hodge and constant metrics are therefore uniformly comparable. In a global trivialization, the automorphy in (362) is the difference of the logarithm of this bounded metric potential at the two ends of a reduced arrow. This is the asserted bounded coboundary. □

The triviality just proved has a stronger consequence than boundedness of one metric combination. It gives an actual cube root of the determinant Hodge defect which appeared in Lemma 17.50.

Corollary 17.31 (One fundamental Hodge defect)**.** *There are canonical orbifold-line isomorphisms, after fixing the trivialization in Proposition 17.30,*

$$\mathcal{N}_1 \simeq \mathcal{N}_2^{\otimes 2}, \qquad \mathcal{N} \simeq \mathcal{N}_2^{\otimes 3}. \tag{363}$$

After compatible nonzero rescalings of the Higgs ratios, one also has

$$\Xi_1 = \Xi_2^2, \qquad \Xi = \Xi_2^3. \tag{364}$$

Thus the effective divisors $2P_7$ and $3P_7$ are respectively the second and third powers of the one fundamental divisor P_7; they are not independent sources of curvature.

Proof. The identity $\mathcal{Q}_2 \simeq \mathcal{O}_C$ is $\mathcal{L}_0\mathcal{L}_1\mathcal{L}_2^{-2} \simeq \mathcal{O}_C$. Hence

$$\mathcal{L}_1\mathcal{L}_0^{-1} \simeq (\mathcal{L}_2\mathcal{L}_0^{-1})^2,$$

which is the first identity in (363); taking determinants gives the second. Under these identifications, Ξ_1 and Ξ_2^2 are holomorphic sections of the same line and both have divisor $2P_7$. Their ratio is therefore a nowhere-vanishing function on the compact orbifold, hence a constant. Absorb that constant into the chosen trivialization. Multiplication by Ξ_2 proves the cubic identity. □

Proposition 17.32 (An integral equation for the full order-seven orbit)**.** *Let $\eta + \eta^{-1} = \lambda$, with $\eta = e^{\pi i/7}$, and put*

$$Q_7(u, v) = u^2 - \lambda uv + v^2 = \mathrm{Nm}_{K(\eta)/K}(u - v\eta).$$

For $R = ST$ and a row vector X, one has

$$Q_7(XR) = Q_7(X). \tag{365}$$

Let $\zeta = -\eta^{-1}$, the fixed point of R in the upper half-plane, and write

$$g = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in G_7.$$

Define the polar form and denominator

$$\begin{aligned} D_g &= Q_7(c, d) = c^2 - \lambda cd + d^2, \\ E_g &= 2ac - \lambda(ad + bc) + 2bd. \end{aligned} \tag{366}$$

Then

$$\mathrm{Re}(g\zeta) = \frac{E_g}{2D_g}, \qquad \mathrm{Im}(\sigma_i(g)\zeta_i) = \frac{\sqrt{4 - \lambda_i^2}}{2\sigma_i(D_g)} \quad (i = 0, 1, 2), \tag{367}$$

where ζ_i is the upper root of $Z^2 + \lambda_i Z + 1$. Every $\sigma_i(D_g)$ is positive.

For $x = p/q \in K$, with $p, q \in \mathcal{O}_K$, put

$$A_g(P) = 2D_g p - E_g q \in \mathcal{O}_K. \tag{368}$$

If

$$X_g = a - b\eta, \qquad Y_g = c - d\eta, \qquad Z_g = pY_g - qX_g,$$

then the horizontal incidence has the exact CM factorization

$$A_g(P) - q(\eta - \eta^{-1}) = 2\overline{Y_g}Z_g. \tag{369}$$

Then

$$\prod_{i=0}^{2} \frac{|\sigma_i(x) - \mathrm{Re}(\sigma_i(g)\zeta_i)|}{\mathrm{Im}(\sigma_i(g)\zeta_i)} = \frac{|\,\mathrm{N}\, A_g(P)|}{\sqrt{7}\,|\,\mathrm{N}\, q|}, \tag{370}$$

$$A_g(P)^2 + (4 - \lambda^2)q^2 = 4D_g Q_7(pc - qa, pd - qb). \tag{371}$$

In the last formula $(pc - qa, pd - qb) = (-q, p)g$ is exactly the transported remainder row.

Proof. The element $u - v\eta$ is carried by right multiplication by R to $\eta(u - v\eta)$; this proves (365). Since $\zeta + \bar{\zeta} = -\lambda$ and $\zeta\bar{\zeta} = 1$, direct rationalization of $(a\zeta + b)/(c\zeta + d)$ gives (367). The quadratic form Q_7 is positive definite at every real embedding because $|\lambda_i| < 2$.

The first identity in (370) follows from

$$x - \frac{E_g}{2D_g} = \frac{A_g(P)}{2qD_g}$$

at the three embeddings and from $\mathrm{N}(4 - \lambda^2) = 7$. Finally, the determinant-one Gram identity

$$E_g^2 + (4 - \lambda^2) = 4Q_7(a, b)Q_7(c, d)$$

gives, after completing the square,

$$A_g(P)^2 + (4 - \lambda^2)q^2 = 4D_g\big(D_g p^2 - E_g pq + Q_7(a, b)q^2\big).$$

The expression in parentheses is $Q_7(pc - qa, pd - qb)$, proving (371).

Equivalently, with the notation in the statement,

$$X_g\overline{Y_g} - \overline{X_g}Y_g = \eta - \eta^{-1}.$$

Since $Z_g = pY_g - qX_g$, multiplication by $2\overline{Y_g}$ gives

$$2\overline{Y_g}Z_g = 2pD_g - q\big(E_g + \eta - \eta^{-1}\big),$$

which is (369). □

Proposition 17.33 (The order-seven orbit is a thin orbit of CM bases)**.** *One has $\mathcal{O}_L = \mathcal{O}_K[\eta]$. For every $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in \mathrm{SL}_2(\mathcal{O}_K)$, the pair*

$$\mathcal{B}_g = (X_g, Y_g) = (a - b\eta, c - d\eta)$$

is an $\mathcal{O}_K$-basis of $\mathcal{O}_L$, characterized by the orientation identity

$$X_g\overline{Y_g} - \overline{X_g}Y_g = \delta. \tag{372}$$

Conversely, every ordered $\mathcal{O}_K$-basis of this form satisfying (372) comes from a unique element of $\mathrm{SL}_2(\mathcal{O}_K)$.

Right multiplication by $R = ST$ rotates the complete basis:

$$\mathcal{B}_{gR} = (\eta X_g, \eta Y_g). \tag{373}$$

If (p, q) is primitive, then $Z_g = pY_g - qX_g$ is primitive as an $\mathcal{O}_K$-coordinate in this basis. Hence the orbit $G_7\zeta$ is the image of the thin G_7-suborbit of oriented CM bases modulo the common rotation $(X, Y) \mapsto (\eta X, \eta Y)$; the ambient $\mathrm{SL}_2(\mathcal{O}_K)$-orbit consists of all oriented bases.

Proof. The cyclotomic identity $\eta = -\zeta_7^4$ gives $L = \mathbb{Q}(\zeta_7)$ and $\mathcal{O}_L = \mathcal{O}_K[\eta]$. Relative to the basis $(1, \eta)$, the coefficient determinant of (X_g, Y_g) is a unit. Direct expansion gives

$$(a - b\eta)(c - d\eta^{-1}) - (a - b\eta^{-1})(c - d\eta) = (ad - bc)(\eta - \eta^{-1}) = \delta,$$

which proves both the basis assertion and its converse. Equation (365) was obtained from the stronger linear identity

$$(u, v)R\begin{pmatrix} 1 \\ -\eta \end{pmatrix} = \eta(u - v\eta).$$

Apply it to both rows of g to obtain (373). Finally, the coordinates of Z_g in (X_g, Y_g) are $(-q, p)$, which are primitive by hypothesis. The last assertion is now just the distinction between the subgroup G_7 and the ambient basis-change group. □

Proposition 17.34 (The two-frame Green function is one CM torus coordinate)**.** *Let $q = (q_1, q_2)^t$ and $v = (v_1, v_2)^t$ be two nonzero columns in K^2, with $\Delta = \det(q, v) \neq 0$. Put*

$$Z_q = q_1 - q_2\eta, \qquad Z_v = v_1 - v_2\eta, \qquad t(q,v) = \frac{Z_q\overline{Z_v}}{\overline{Z_q}Z_v} \in L^\times.$$

Then

$$t(q,v)\overline{t(q,v)} = 1, \qquad |\widetilde{\sigma}(t(q,v))| = 1 \quad (\widetilde{\sigma} : L \hookrightarrow \mathbf{C}), \tag{374}$$

$$t(q,v) - 1 = \frac{\delta\Delta}{\overline{Z_q}Z_v}. \tag{375}$$

For the upper root η_i of $X^2 - \lambda_i X + 1$, one has the exact local identity

$$\Lambda(\eta_i; [\sigma_i q], [\sigma_i v]) = \log \frac{2}{|1 - \widetilde{\sigma}_i(t(q,v))|}. \tag{376}$$

Consequently

$$\sum_{i=0}^{2} \Lambda(\eta_i; [\sigma_i q], [\sigma_i v]) = 3\log 2 - \frac{1}{2}\log 7 - \log|\,\mathrm{N}\,\Delta| + \frac{1}{2}\log|\,\mathrm{N}\,Q_7(q_1, q_2)| + \frac{1}{2}\log|\,\mathrm{N}\,Q_7(v_1, v_2)|. \tag{377}$$

In particular,

$$t(q,v) = 1 \quad \Longleftrightarrow \quad \Delta = 0 \tag{378}$$

when the formula is extended to the coincident-line case. For $(q, v) = (Q_m, Q_n)$, the torus coordinate is unchanged by the balancing units and packages the central Green recurrence channel into one global algebraic number.

Proof. CM conjugation immediately proves (374). Direct expansion gives

$$Z_q\overline{Z_v} - \overline{Z_q}Z_v = (q_1v_2 - q_2v_1)(\eta - \eta^{-1}) = \Delta\delta,$$

which is (375). At the embedding $\widetilde{\sigma}_i$, the two factors in its denominator have absolute values $|Z_{q,i}|$ and $|Z_{v,i}|$. Since $\operatorname{Im}\eta_i = |\delta_i|/2$, the definition of Λ gives

$$\exp\Lambda(\eta_i; [q_i], [v_i]) = \frac{|Z_{q,i}|\,|Z_{v,i}|}{(|\delta_i|/2)|\Delta_i|} = \frac{2}{|1 - t_i|},$$

proving (376). Finally, $|Z_{q,i}|^2 = \sigma_i(Q_7(q_1, q_2))$, and similarly for v, while $\prod_i |\delta_i| = \sqrt{7}$. Summing the logarithm proves (377). Equation (378) follows from (375); scaling either column by an element of $K^\times$ cancels from $t(q, v)$. □

Why this is useful, and why it is not yet the proof. The order-seven divisor is not merely a formal Hodge class. At its algebraic lift, the complete three-place Green proximity is exactly proximity of one norm-one CM-torus point to the identity, and (377) enters the determinant norm once. On a fixed compact return set, changing the three basepoints from the period points to η_i changes the sum of Green functions by $O(1)$, uniformly in the two rational lines. Thus the Green channel in Proposition 8.20 has a fully algebraic global model.

This does not make ordinary Roth sufficient. The target 1 and the approximants $t(Q_m, Q_n)$ lie in the same fixed number field, and the approximants are not known to belong to one fixed finitely generated multiplicative subgroup. The product formula in (375) is therefore at the critical scale. What the thin group must add is a restriction on the consecutive sequence of these torus points, or an arithmetic-holonomy estimate for the companion controlled-Wronskian channel. The proposition identifies the global object on which such a theorem has to act; it does not replace that theorem by a scalar approximation claim.

Proposition 17.35 (The two congruence incidences of the full thin group)**.** *Put $L = K(\eta)$, write a bar for the nontrivial element of $\mathrm{Gal}(L/K)$, and set $\delta = \eta - \eta^{-1}$. Assume that (p, q) is primitive in $\mathcal{O}_K^2$. With the notation of Proposition 17.32, one has*

$$\begin{aligned} (Y_g, \overline{Y_g}) &\mid (\delta), \\ (Z_g, \overline{Z_g}) &\mid (\delta), \\ (Y_g, Z_g) &\mid (q\delta). \end{aligned} \tag{379}$$

The middle assertion is an assertion about ideal support; equivalently, no prime outside (δ) divides both entries.

Consequently every prime of L outside $(q\delta)$ which divides $Y_g Z_g$ lies over a prime of K which splits in L/K. If $\mathfrak{p} = \mathfrak{P}\overline{\mathfrak{P}}$ is such a split prime, then in $\mathbb{P}^1(\mathcal{O}_L/\mathfrak{P})$

$$\begin{aligned} \mathfrak{P} \mid Y_g \quad &\Longleftrightarrow \quad \overline{g}\,\overline{\zeta} = \infty, \\ \mathfrak{P} \mid Z_g \quad &\Longleftrightarrow \quad \overline{g}\,\overline{\zeta} = \overline{x}. \end{aligned} \tag{380}$$

Here bars over g, ζ, x in (380) mean reduction modulo $\mathfrak{P}$, rather than CM conjugation.

For all but finitely many primes $\mathfrak{p}$ of K, strong approximation gives

$$G_7 \bmod \mathfrak{p} = \mathrm{SL}_2(\mathcal{O}_K/\mathfrak{p}).$$

At every such split prime, the two conditions in (380) are therefore two complete, distinct point-stabilizer cosets in the full finite projective orbit.

Proof. The determinant-one identity already used above is

$$X_g\overline{Y_g} - \overline{X_g}Y_g = \delta,$$

so a common prime divisor of Y_g and $\overline{Y_g}$ divides δ. Write

$$Z_g = r_g - s_g\eta, \qquad (r_g, s_g) = (pc - qa, pd - qb) = (-q, p)g.$$

Because $g \in \mathrm{SL}_2(\mathcal{O}_K)$ and (p, q) is primitive, (r_g, s_g) is primitive. A prime outside (δ) which divided both Z_g and $\overline{Z_g}$ would, by subtraction, divide s_g, and then r_g, a contradiction. Finally,

$$\overline{Y_g}Z_g - Y_g\overline{Z_g} = -q\delta,$$

which proves the last coprimality.

If a good prime $\mathfrak{P}$ were inert, then it would be fixed by CM conjugation. Divisibility of Y_g, or of Z_g, would contradict the first, respectively second, coprimality. Ramified primes divide the relative discriminant and hence (δ). Thus every remaining prime is split.

Since $\zeta = -\eta^{-1}$, direct cancellation gives

$$g\zeta = \frac{a - b\eta}{c - d\eta} = \frac{X_g}{Y_g}, \qquad Z_g = Y_g(p - qg\zeta).$$

The determinant identity prevents X_g and Y_g from vanishing simultaneously at a good prime, and $(Y_g, Z_g) \mid (q\delta)$ separates the two displayed incidences. This proves (380). The last assertion is the standard strong-approximation theorem for the finitely generated Zariski-dense subgroup $G_7 < \mathrm{SL}_2(K)$, followed by the transitive action of SL_2 on the projective line. □

Corollary 17.36 (The Steinberg coefficient selected by $3P_7$)**.** *Retain a good split prime from Proposition 17.35, write $k_\mathfrak{p} = \mathcal{O}_K/\mathfrak{p}$, and put $Q_\mathfrak{p} = |k_\mathfrak{p}|$. On $\ell^2(\mathbb{P}^1(k_\mathfrak{p}))$, set*

$$e_y = \delta_y - \frac{1}{Q_\mathfrak{p}+1}\mathbf{1}.$$

These vectors span the Steinberg representation $\mathrm{St}_\mathfrak{p} = \mathbf{1}^\perp$. For every $g \in G_7$ and every residue point y,

$$\mathbf{1}_{\{\overline{g}\,\overline{\zeta}=y\}} = \frac{1}{Q_\mathfrak{p}+1} + \left\langle \mathrm{St}_\mathfrak{p}(\overline{g})e_{\overline{\zeta}}, e_y \right\rangle. \tag{381}$$

In particular, Y_g and Z_g are integral lifts of the two incidence coefficients in (381) with $y = \infty$ and $y = \overline{x}$, respectively.

Proof. The permutation representation on the finite projective line is $\mathbf{1} \oplus \mathrm{St}_\mathfrak{p}$, and $\mathrm{St}_\mathfrak{p}(\overline{g})e_z = e_{\overline{g}z}$. Thus

$$\langle e_{\overline{g}\overline{\zeta}}, e_y \rangle = \mathbf{1}_{\{\overline{g}\overline{\zeta}=y\}} - \frac{1}{Q_\mathfrak{p}+1}.$$

The last assertion is (380). □

Proposition 17.37 (The global norm-one-torus ledger)**.** *Set $\delta = \eta - \eta^{-1}$ and define*

$$t_g(P) = \frac{A_g(P) + q\delta}{A_g(P) - q\delta} = \frac{Y_g\overline{Z_g}}{\overline{Y_g}Z_g} \in L^\times. \tag{382}$$

Then

$$t_g\overline{t_g} = 1, \qquad |\widetilde{\sigma}(t_g)| = 1$$

for every complex embedding $\widetilde{\sigma} : L \hookrightarrow \mathbb{C}$. Let

$$h_L^{\mathrm{fin}}(t) = \sum_{\mathfrak{P}} \max\{0, v_\mathfrak{P}(t)\} \log \mathrm{N}\,\mathfrak{P}$$

be the unnormalized finite logarithmic height. There is a constant depending only on the fixed primitive pair $P = (p, q)$ such that

$$\begin{aligned} h_L^{\mathrm{fin}}(t_g) &= \log|\,\mathrm{N}_{L/\mathbb{Q}}(Y_gZ_g)| + O_P(1) \\ &= \log|\,\mathrm{N}_{K/\mathbb{Q}}(D_gQ_7(pc-qa, pd-qb))| + O_P(1) \\ &= \log\left|\mathrm{N}_{K/\mathbb{Q}}\big(A_g(P)^2 + (4-\lambda^2)q^2\big)\right| + O_P(1). \end{aligned} \tag{383}$$

Because all archimedean absolute values of t_g are one, $h_L^{\mathrm{fin}}(t_g)$ is its full Weil height, apart from the chosen normalizing factor. If t_g is a unit of $\mathcal{O}_L$, then it is a root of unity.

Proof. CM conjugation of (369) gives

$$A_g(P) + q\delta = 2Y_g\overline{Z_g},$$

which proves (382). The norm-one and archimedean assertions follow immediately.

At a prime outside $(q\delta)$, the three coprimalities in (379) say that exactly one member of each conjugate pair can occur in each factor, and that Y_g and Z_g cannot use the same member. Therefore the contribution of that conjugate pair to $h_L^{\mathrm{fin}}(t_g)$ is exactly its contribution to $\log \mathrm{N}(Y_g Z_g)$. At the fixed primes dividing $q\delta$, any cancellation is bounded by the valuations of $q\delta$, again by (379). This proves the first line of (383). The second follows from $Y_g\overline{Y_g} = D_g$ and

$$Z_g\overline{Z_g} = Q_7(pc - qa, pd - qb),$$

and the third from (371); the factor 4 is absorbed by $O_P(1)$. Finally, a unit t_g is an algebraic integer all of whose conjugates have modulus one, so Kronecker's theorem makes it a root of unity. □

Corollary 17.38 (Height of the elliptic incidence coordinate)**.** *If h_K and h_L denote unnormalized logarithmic projective heights, then*

$$h_L(t_g) = 2h_K([A_g(P) : q]) + O_P(1). \tag{384}$$

Equivalently,

$$h_L(t_g) = \sum_{i=0}^{2} \log\left(A_{g,i}^2 + |\widetilde{\sigma}_i(\delta)|^2 q_i^2\right) + O_P(1),$$

where $A_{g,i} = \sigma_i(A_g(P))$ and $q_i = \sigma_i(q)$.

Proof. The Cayley transform

$$[A : q] \longmapsto [A + q\delta : A - q\delta]$$

is a fixed projective linear transformation over L, so it changes height by $O_P(1)$. Extending a K-point to the quadratic extension L doubles the unnormalized height. The local displayed formula follows from $\max(|u|, |v|)^2 \asymp |u|^2 + |v|^2$. □

Corollary 17.39 (Angular cross-ratio interpretation)**.** *Put $w_{g,i} = \sigma_i(g)\zeta_i$ and $x_i = \sigma_i(x)$. For either extension of σ_i to L,*

$$\widetilde{\sigma}_i(t_g) = \frac{x_i - \overline{w_{g,i}}}{x_i - w_{g,i}}, \qquad \frac{|\widetilde{\sigma}_i(t_g) + 1|}{|\widetilde{\sigma}_i(t_g) - 1|} = \frac{|x_i - \operatorname{Re} w_{g,i}|}{\operatorname{Im} w_{g,i}}. \tag{385}$$

Thus t_g is the adelic angular coordinate of the segment joining the order-seven orbit point $g\zeta$ to the rational boundary point x, and approximation of $g\zeta$ to the vertical geodesic ending at x is approximation of t_g to the root of unity -1.

Proof. Since $Z_g = Y_g(p - qg\zeta) = qY_g(x - g\zeta)$, substitution in (382) gives the first formula. Writing $w = u + iv$ gives

$$\left|\frac{x - \bar{w}}{x - w} + 1\right| = \frac{2|x - u|}{|x - w|}, \qquad \left|\frac{x - \bar{w}}{x - w} - 1\right| = \frac{2v}{|x - w|},$$

which proves the second. □

Proposition 17.40 (Cayley-state heights)**.** *Define the fixed Cayley parametrization on K-points*

$$c_\eta : \mathbb{P}^1(K) \longrightarrow \left(\mathrm{Res}^1_{L/K}\, \mathbb{G}_m\right)(K), \qquad c_\eta(y) = \frac{1 + y\eta}{1 + y\eta^{-1}}.$$

For $g \in G_7$, put

$$u_g = \frac{Z_g}{\overline{Z_g}}, \qquad v_g = \frac{Y_g}{\overline{Y_g}}.$$

Then

$$u_g = c_\eta(g^{-1}x), \qquad v_g = c_\eta(g^{-1}\infty), \qquad t_g = \frac{v_g}{u_g}. \tag{386}$$

Moreover, for unnormalized logarithmic heights,

$$\begin{aligned} h_L(u_g) &= \log|\,\mathrm{N}_{K/\mathbb{Q}}\, Q_7(pc-qa, pd-qb)| + O_P(1) = 2h_K(g^{-1}x) + O_P(1),\\ h_L(v_g) &= \log|\,\mathrm{N}_{K/\mathbb{Q}}\, D_g| + O(1) = 2h_K(g^{-1}\infty) + O(1),\\ h_L(t_g) &= h_L(u_g) + h_L(v_g) + O_P(1). \end{aligned} \tag{387}$$

If $u_g = u_h$, then hg^{-1} fixes x; if $v_g = v_h$, then hg^{-1} fixes ∞.

Proof. Write

$$(r_g, s_g) = (pc-qa, pd-qb) = (-q, p)g.$$

Since

$$g^{-1}x = -\frac{s_g}{r_g}, \qquad g^{-1}\infty = -\frac{d}{c},$$

the definitions of $Z_g = r_g - s_g\eta$ and $Y_g = c - d\eta$ give the first two identities in (386); the third is (382).

The coprimalities $(Z_g, \overline{Z_g}) \mid (\delta)$ and $(Y_g, \overline{Y_g}) \mid (\delta)$ show that the finite heights of u_g and v_g equal the logarithmic norms of Z_g and Y_g, up to fixed ramified terms. Taking relative norms gives the first expressions in (387). The fixed projective linear map c_η changes height by $O(1)$, and extension from K to L doubles unnormalized height. The last line is Proposition 17.37. Finally, c_η is injective. Thus $u_g = u_h$ gives $g^{-1}x = h^{-1}x$, whence $hg^{-1}x = x$; the cusp statement is identical. □

Proposition 17.41 (The Cayley pair is one full-group orbit)**.** *Let*

$$C_\eta = \begin{pmatrix} \eta & 1 \\ \eta^{-1} & 1 \end{pmatrix}, \qquad \widehat{g} = C_\eta g^{-1} C_\eta^{-1} \in \mathrm{PGL}_2(L).$$

Then the collection $\widehat{G}_7 = \{\widehat{g} : g \in G_7\}$ is the Cayley conjugate of the complete G_7-action on $\mathbb{P}^1(K)$. It preserves $\big(\mathrm{Res}^1_{L/K}\,\mathbb{G}_m\big)(K)$, and at every complex embedding it preserves the unit circle. If

$$u_* = c_\eta(x), \qquad v_* = c_\eta(\infty) = \eta^2,$$

then

$$(u_g, v_g) = (\widehat{g}u_*, \widehat{g}v_*). \tag{388}$$

In particular,

$$\mathrm{Stab}_{G_7}(x) = \{g \in G_7 : \widehat{g}u_* = u_*\}.$$

Proof. The matrix C_η represents the fractional linear map c_η. Hence

$$c_\eta(g^{-1}y) = \widehat{g}\, c_\eta(y)$$

for every $y \in \mathbb{P}^1(K)$, which gives (388). CM conjugation sends $c_\eta(y)$ to its inverse, so these points have relative norm one. At a complex embedding, C_η carries the real projective line to the unit circle; conjugating a real Möbius transformation by it therefore preserves that circle. The stabilizer identity follows from the injectivity of c_η. □

Proposition 17.42 (The Cayley heights count the full pointed congruence tower)**.** *Let S_0 be the fixed set of primes of L dividing δ. For $\mathfrak{P} \notin S_0$, write*

$$\iota_{\mathfrak{P},m}(g;y) = \mathbf{1}_{\{g\zeta \equiv y \pmod{\mathfrak{P}^m}\}} \qquad (y = x, \infty).$$

Then, with the usual projective interpretation of these congruences,

$$\begin{aligned} h_L(u_g) &= \sum_{\mathfrak{P}\notin S_0} \sum_{m\geq 1} \iota_{\mathfrak{P},m}(g;x) \log \mathrm{N}\,\mathfrak{P} + O_P(1), \\ h_L(v_g) &= \sum_{\mathfrak{P}\notin S_0} \sum_{m\geq 1} \iota_{\mathfrak{P},m}(g;\infty) \log \mathrm{N}\,\mathfrak{P} + O_P(1). \end{aligned} \tag{389}$$

Every nonzero summand lies over a prime of K split in L/K.

Proof. Use the primitive coordinates

$$g\zeta = [X_g : Y_g], \qquad x = [p : q], \qquad Z_g = pY_g - qX_g.$$

Outside δ, the first two coprimalities in Proposition 17.35 make both displayed coordinate pairs primitive. Hence the projective determinant formula gives

$$v_{\mathfrak{P}}(Z_g) = \sum_{m\geq 1} \iota_{\mathfrak{P},m}(g;x), \qquad v_{\mathfrak{P}}(Y_g) = \sum_{m\geq 1} \iota_{\mathfrak{P},m}(g;\infty).$$

The coprimalities with the CM conjugates show that the positive valuation of $u_g = Z_g/\overline{Z_g}$, respectively $v_g = Y_g/\overline{Y_g}$, is exactly the displayed valuation. The common valuations at primes in S_0 are bounded in terms of δ, and so contribute $O(1)$. Summing proves (389); splitness is Proposition 17.35. □

Corollary 17.43 (Whole-orbit Cayley closing criterion)**.** *Let (g_n) be an infinite sequence of projectively distinct elements of G_7. If*

$$\sup_n \big(h_L(t_{g_n}) - h_L(v_{g_n})\big) < \infty, \tag{390}$$

then x has a nonidentity stabilizer in G_7.

Proof. Equation (387) turns (390) into a uniform bound for $h_L(u_{g_n})$. Northcott's theorem leaves only finitely many possible u_{g_n}, because they all belong to the fixed field L. Hence $u_{g_m} = u_{g_n}$ for some $m \neq n$, and Proposition 17.40 says that $g_n g_m^{-1}$ fixes x. Projective distinctness makes this stabilizer nonidentity. □

The corollary is formal, not the missing thin-group theorem. Its value is that it identifies the exact relative height which such a theorem must control. The term $h_L(v_g)$ is the entire parabolic or cusp payment; after subtracting it, no letter-by-letter quantity remains:

$$h_L(t_g) - h_L(v_g) = h_L\big(c_\eta(g^{-1}x)\big) + O_P(1).$$

Thus a useful exploitation of thinness must act on the projectively distinct *whole orbit* of the pair (x, ∞). In geometric language it should be a relative domination theorem: arbitrarily long parabolic excursions are charged to v_g, while each return through the thick part controls the residual state

u_g. Uniform estimates for isolated generators do not state this, and ordinary expansion of finite quotients does not imply it for the single trajectory selected by the stopping rule.

This proposition is the arithmetic form of the phrase "the divisor $3P_7$." The lifts of P_7 are not an unspecified collection of charged branches. They are the single orbit $G_7\zeta$, and their horizontal incidence with the rational vertical leaf is the algebraic integer $A_g(P)$. Formula (371) also joins this orbit equation to the complete two-coordinate remainder state. Thus the curved part of the stopped problem can be stated as an Arakelov comparison between two explicit norm forms.

The factorization (369) locates the S-unit issue exactly. Outside the fixed set of primes dividing $2q(\eta-\eta^{-1})$, the two conjugate factors $A_g(P)\pm q(\eta-\eta^{-1})$ are coprime. Their prime ideals are distributed between the cusp factor Y_g and the complete state factor Z_g. Thus the flat escape is not a collection of unrelated primes attached to Rosen letters; it is the ideal distribution in one CM factorization attached to the full order-seven orbit. Bounded support turns its quotient into an ordinary S-unit equation. The unrestricted case asks for a global statement preventing this ideal distribution from being recycled at successive first norm records.

Proposition 17.35 makes "recycled" a presentation-free term. It does not mean that a prime is first noticed in a particular generator. It means that, along the one reduced orbit $g\zeta$, successive norm records revisit the two congruence sections $g\zeta=\infty$ and $g\zeta=x$. At almost every possible prime these are full Borel cosets in a surjective reduction of the thin group. The remaining estimate must therefore concern a pointed trajectory through these complete finite quotients. Ordinary strong approximation proves that each quotient is complete; expansion controls averages over group balls. Neither statement alone controls the one trajectory selected by the continued-fraction stopping rule, which is the precise extra uniformity still required.

Corollary 17.36 also explains why the order-seven point cannot be replaced by a generic target. The Hodge defect $3P_7$ supplies the source vector $e_{\overline{\zeta}}$, while the rational Borel supplies the target $e_{\overline{x}}$. There is only one nonconstant finite representation to control. If Hodge-balanced shells avoid the second incidence, their pointed Steinberg coefficient takes the extremal value $-1/(Q_{\mathfrak{p}}+1)$. The required thin-group theorem is therefore an inverse statement: persistence of that extremal coefficient over coherent growing split levels must force a characteristic-zero stabilizer of x. This is the same stabilizer that gives the cusp or special alternative, now read directly from the full order-seven orbit.

Proposition 17.37 packages the entire finite divisor into one Arakelov height. The torus $\mathrm{Res}^1_{L/K}\mathbb{G}_m$ is compact at every real place, so its height is entirely finite-place mass. The unrestricted S-unit problem is precisely the possible growth of $h_L(t_g)$ along first records. Bounded height gives finitely many t_g by Northcott, while height zero gives only roots of unity. Thus the remaining flat-direction theorem can be stated globally: unbounded torus height on a coherent Hodge-stopped thin orbit must either pay the positive $3P_7$ degree or align with one characteristic-zero cyclic stabilizer.

This also states exactly what scalar Roth misses. Formula (370) controls the product $\prod_i|A_{g,i}|$, hence the sum of three signed logarithms. Formula (384) controls instead

$$\sum_i \log\max\bigl(|A_{g,i}|,|\widetilde{\sigma}_i(\delta)q_i|\bigr).$$

The difference is the amount that size is redistributed among the three Galois places. It is exactly the neutral Asai direction. Consequently a stronger constant in C/q^2, or even a fixed-target Roth theorem, cannot supply the missing estimate unless it also controls this projective torus height.

There is still a genuine step here. The product formula (370) controls the sum of the three local elliptic proximities, while a first remainder record selects a maximal slope before that sum is taken.

The neutral coordinate (391) measures precisely the possible redistribution among the three factors. Consequently the CM norm identity removes any justification for generator-by-generator charges, but it does not by itself interchange "maximum" and "sum."

This ledger sharpens the Galois-descent obstruction. The positive conjugate degree is not an unspecified analytic loss: after squaring, it is exactly the fixed effective divisor $3P_7$, while the third cyclic component is not merely degree zero: it is the trivial line in Proposition 17.30. In particular the inverse section Υ_1^{-1} is a holomorphic section of $\mathcal{Q}_1^2$ vanishing on $3P_7$. Thus a maximal-slope proof may charge positive Galois compensation only to the order-seven height, rather than to arbitrary unit drift. What remains is to compare that fixed elliptic divisor with the moving reciprocal incidence divisor at a first norm record. This is a substantially narrower arithmetic intersection problem than bounding three unrelated real norms.

This also separates the two parts of that remaining comparison. If

$$u_i(r) = \log|\sigma_i(r)| - \log|\sigma_0(r)|, \qquad i = 1, 2,$$

then the neutral coordinate is

$$u_{\mathrm{neut}}(r) = u_1(r) - 2u_2(r) = \log\left|\frac{\sigma_0(r)\sigma_1(r)}{\sigma_2(r)^2}\right|. \tag{391}$$

There is no unbounded Hodge curvature hidden in this coordinate: Proposition 17.30 removes it by a bounded change of metric. Thus an unbounded value of (391) is purely arithmetic unit or ideal escape. The complementary cyclic coordinate is the one charged by $3P_7$. The full maximal-slope problem is therefore a coupling of one curved Hodge direction with one flat S-unit direction, rather than two unspecified outer losses. The ordinary S-unit theorem controls the flat direction once the relevant multipliers lie in one fixed finite-rank group. The remaining global task is to prove directly that the order-seven height and the moving incidence record prevent unbounded escape in the norm-one torus.

> **The first genuinely thin-group input.** The presentation $C_2 * C_7$, the two Galois projective rotation classes, and their simultaneous occurrence in one integral monodromy representation are properties of the full thin group. Formula (292) cannot be recovered from estimates for isolated generator moves. It is the global reason that the required height has two coordinates of unequal slopes.

Lemma 17.44 (The determinant character is not proper)**.** *On the unit logarithm plane $t_0+t_1+t_2 = 0$, the pairwise character*

$$(t_0, t_1, t_2) \longmapsto (t_1 - t_0, t_2 - t_0) \tag{392}$$

is a linear isomorphism onto $\mathbf{R}^2$, and hence is proper on the unit lattice. Its sum

$$(t_1 - t_0) + (t_2 - t_0) = -3t_0,$$

which is the character of $\det \mathcal{E}_{\mathrm{pair}} = \mathcal{N}$, is not proper on that rank-two lattice.

Proof. The inverse to (392), with coordinates (u_1, u_2), is

$$t_0 = -\frac{u_1 + u_2}{3}, \qquad t_1 = \frac{2u_1 - u_2}{3}, \qquad t_2 = \frac{-u_1 + 2u_2}{3}.$$

This proves the first assertion. Choose multiplicatively independent units ϵ, η. The ratio $\log|\sigma_0\epsilon| / \log|\sigma_0\eta|$ is irrational; otherwise a nontrivial product $\epsilon^m\eta^{-n}$ would have distinguished absolute value one, hence would be ± 1, contrary to independence. The additive group generated by the two distinguished logarithms is therefore dense in $\mathbf{R}$. It contains a sequence tending to zero whose exponent vectors tend to infinity. Along that sequence $-3t_0$ stays bounded while $(t_1 - t_0, t_2 - t_0)$ diverges. □

This lemma identifies the precise loss in every scalar Arakelov degree. The degree of $\mathcal{N}$ controls the sum of two slopes, while a unit can send one pairwise component to infinity and the other to minus infinity with nearly unchanged sum. A canonical sufficient strengthening is therefore a rank-two successive-minima or maximal-slope statement for an integral realization of $\mathcal{E}_{\mathrm{pair}}$, rather than merely an upper bound for its determinant degree. If both successive minima are bounded, lattice discreteness gives repetition. If equality persists in one rational line, its holonomy is a unit eigenline and gives the special alternative. This does not rule out a scalar proof supplemented by a separate rigidity theorem saying that escape in the kernel of the determinant character is necessarily periodic. Projective states are already insensitive to a common unit rescaling; what the scalar degree fails to control is whether successive rescalings come from one group holonomy.

There is also a representation-theoretic no-go statement which prevents a shortcut through a different flat line. Let $H = \mathrm{Res}_{K/\mathbf{Q}} \mathrm{SL}_2$. Over $\mathbf{R}$, a split maximal torus has weight lattice $\mathbf{Z}^3$, and the Galois group cyclically permutes the three coordinates. The pairwise Hodge characters are

$$\chi_1 = (-1, 1, 0), \qquad \chi_2 = (-1, 0, 1),$$

and their determinant is $(-2, 1, 1)$. None is Galois invariant.

Proposition 17.45 (No flat rational pairwise line)**.** *No nontrivial one-dimensional* $\mathbf{Q}$*-subquotient of an algebraic representation of* H *has any of the characters* χ_1, χ_2*, or* $(-2, 1, 1)$*. In particular, the pairwise Hodge line cannot be replaced by a hidden flat integral line carrying the same weight. Its smallest Galois-stable datum contains the full cyclic orbit of the character.*

Proof. The connected group H is semisimple, so every algebraic character $H \to \mathbb{G}_m$ is trivial. Equivalently, every one-dimensional subquotient of an algebraic representation has zero torus weight. Even at the level of the split torus, a $\mathbf{Q}$-defined one-dimensional weight must be fixed by the cyclic Galois action; the three displayed asymmetric weights are not. Their cyclic orbits have three elements. □

The Hodge filtration avoids this obstruction because it is a reduction of the flat H-bundle to a Borel, and characters of that Borel do define nonflat line bundles. Their nonzero degrees are precisely the curvature being used. The price is unavoidable: integrality lives in the ambient flat representation, while contraction lives in its Hodge-graded pieces. The three Asai components below are the concrete Galois-stable package which retains both sides.

There is, however, a useful integral realization which uses every element of the group and retains the complete boundary point. It is more informative than any estimate for the six Rosen letters. Let $V = \mathcal{O}_K^2$, let τ generate $\mathrm{Gal}(K/\mathbf{Q})$, and write ${}^{\tau^i}V$ for the same module with its coefficients twisted by τ^i. Define

$$\mathbb{W}_i = V \otimes_{\mathcal{O}_K} {}^{\tau^i}V, \qquad \rho_i(g) = g \otimes \tau^i(g), \qquad i = 1, 2. \tag{393}$$

For a primitive column $P \in V$, put

$$w_i(P) = P \otimes \tau^i(P) \in \mathbb{W}_i. \tag{394}$$

Proposition 17.46 (The full group supplies an integral contracting state)**.** *The tensors* $w_i(P)$ *are primitive and satisfy*

$$\rho_i(g) w_i(P) = w_i(gP) \qquad (g \in G_7). \tag{395}$$

Moreover,

$$[\rho_i(g) w_i(P)] = [w_i(P)] \quad \Longleftrightarrow \quad g[P] = [P]. \tag{396}$$

Along the distinguished geodesic ending at $[P]$, the largest Hodge weight occurring in the first real realization of $w_i(P)$ is

$$-a_0 + a_i \le 0. \tag{397}$$

It is strictly negative on every fixed compact part of the triangle quotient. Consequently the adapted norm of this real tensor is bounded at all reduced returns, and tends to zero if the accumulated compact-part length tends to infinity.

Proof. Primitivity follows because the content ideal of a tensor is the product of the two content ideals. Equivariance is immediate. If two nonzero decomposable tensors in $V \otimes^{\tau^i} V$ are proportional, their first factor lines are equal. This proves (396); the converse is clear.

The vector P is the contracting boundary line in the uniformizing factor. The other factor can have either Hodge weight, so its largest possible weight is $+a_i$. This gives (397). The modular map $F_i : \mathbb{H} \to \mathbb{H}$ is distance nonincreasing, so $a_i \le a_0$. On a fixed compact subset it is uniformly strictly contracting because it is not an automorphism. Integration along the geodesic gives the final assertion. □

This proposition resolves two issues which a scalar remainder cannot resolve. The construction is a representation of the entire thin group, not a list of generator inequalities, and its projective tensor separates the complete boundary state. Thus repetition gives the required Borel intersection immediately. It also removes the need for a quantitative comparison between a long cusp block and the total time: every compact return contributes some strict contraction, however sparse the returns are.

There is an exact reason why this fixed tensor does not already prove the norm-record estimate. The alternating form on each rank-two factor gives a bilinear form B_i on $\mathbb{W}_i$. If $r_g = \det(P, gP)$, then

$$B_i\big(w_i(P), \rho_i(g) w_i(P)\big) = r_g \tau^i(r_g). \tag{398}$$

Its finite ideal is $(r_g)(\tau^i r_g)$, and multiplying the two pairings gives

$$B_1 B_2 = r_g \, \mathrm{N}(r_g). \tag{399}$$

Thus the fixed, contracting object has a zero along the incidence divisor. The stopped Arakelov line constructed in Section 13 needs the opposite finite divisor $(r_g)^{-1}$. Passing from (399) to that reciprocal divisor is not a change of metric: it requires division by a moving integral evaluation. This is the finite-sign form of the flat-line/record-line obstruction.

That division can nevertheless be made canonically at the level of the complete state. Put

$$\mathbb{W} = \mathrm{Sym}^2 V \otimes {}^{\tau} V \otimes {}^{\tau^2} V, \qquad \rho_{\mathbb{W}}(g) = \mathrm{Sym}^2(g) \otimes \tau(g) \otimes \tau^2(g). \tag{400}$$

For a primitive state $X = (u, v)^t$ with $v \neq 0$, define

$$\mathbf{U}(X) = \frac{X^{\otimes 2} \otimes \tau(X) \otimes \tau^2(X)}{v} \in \mathbb{W} \otimes K. \tag{401}$$

Proposition 17.47 (Canonical transport of the reciprocal incidence)**.** *The content ideal of $\mathbf{U}(X)$ is exactly $(v)^{-1}$. If*

$$M_i(X) = \max\{|\sigma_i u|, |\sigma_i v|\},$$

then the tensor sup norm at the distinguished realization satisfies

$$\log \|\sigma_0 \mathbf{U}(X)\|_\infty = 2 \log M_0(X) + \log M_1(X) + \log M_2(X) - \log |\sigma_0 v| = \mathscr{E}(X). \tag{402}$$

If $A \in G_7$, $X' = AX = (u', v')^t$, and $v' \neq 0$, then

$$\rho_{\mathbb{W}}(A)\mathbf{U}(X) = \frac{v'}{v}\,\mathbf{U}(X'). \tag{403}$$

Finally, the projective tensor $[\mathbf{U}(X)]$ determines $[X]$. In particular, repetition of normalized tensors gives a projective stabilizer.

Proof. The content of $X^{\otimes 2} \otimes \tau X \otimes \tau^2 X$ is the product of four unit content ideals, hence is $\mathcal{O}_K$. Division by v proves the first claim. The largest tensor coordinate at σ_0 is $M_0(X)^2 M_1(X) M_2(X)/|\sigma_0 v|$. Taking its logarithm and using the product formula gives exactly the definition of $\mathscr{E}$, which proves (402).

The numerator in (401) is equivariant for $\rho_{\mathbb{W}}$. Replacing its denominator v by the new denominator v' gives (403). If two tensors are projectively equal, their $\operatorname{Sym}^2 V$-factor lines are equal; the Veronese map is injective, so the original projective lines are equal. □

The numerator of $\mathbf{U}$ is the tensor product of the two pairwise states in Proposition 17.46. Its largest distinguished Hodge weight is $-2a_0 + a_1 + a_2$, so it contracts along every nonconstant part of the geodesic. Equation (403) shows exactly what can undo this contraction: the one scalar v'/v. At a special return $X' = \epsilon X$, one has $\mathbf{U}(X') = \mathrm{N}(\epsilon)\mathbf{U}(X) = \pm\mathbf{U}(X)$; thus the normalization retains the required neutral unit holonomy.

This removes an ambiguity in the earlier formulation of the obstruction. An integral ambient tensor, its reciprocal incidence lattice, and its transport all exist explicitly. What is not yet proved is the stopped maximum principle for the scalar cocycle in (403). A new norm record controls $|\mathrm{N}(v'/v)|$, but not the distinguished value $|\sigma_0(v'/v)|$; the two outer values can absorb the difference. Conjecture 12.6 is precisely the assertion that the Hodge contraction prevents this absorption from recurring indefinitely, apart from the neutral holonomy just described.

17.2 Why the affine framing cannot be discarded

There is a small but consequential distinction between the tensor and its projective line. It prevents one from applying a theorem about the projective flat bundle and claiming that the record divisor has thereby been controlled.

Proposition 17.48 (Projectivization forgets the record ideal)**.** *For every primitive $X = (u, v)^t$, with $v \neq 0$,*

$$[\mathbf{U}(X)] = \left[X^{\otimes 2} \otimes \tau X \otimes \tau^2 X\right] \quad \text{in } \mathbb{P}(\mathbb{W}). \tag{404}$$

Consequently the projective point contains no information about the fractional ideal $(v)^{-1}$. That ideal belongs to the affine datum

$$\mathfrak{F}(X) = \left((v)^{-1},\ \mathbf{U}(X) \in \mathbb{W} \otimes_{\mathcal{O}_K} (v)^{-1}\right), \tag{405}$$

not to $[\mathbf{U}(X)]$ alone. Multiplying the affine lift by v removes its reciprocal content and changes the real norms by exactly the amount which the product formula assigns to $(v)^{-1}$.

Proof. Equation (404) follows because division by $v \in K^\times$ is scalar multiplication in $\mathbb{W} \otimes K$. The numerator is primitive, whereas Proposition 17.47 gives the content of the displayed affine lift as $(v)^{-1}$. Clearing that content by v returns the numerator. At the three real places the logarithmic change is $\sum_i \log|\sigma_i(v)| = \log|\mathrm{N}\, v|$, which is the negative of the finite-degree change under the convention used here. Thus the archimedean and finite changes cancel, as asserted. □

This is more than a warning about notation. The equality locus in Proposition 17.54 is projective, but the inequality to be proved is affine. The correct geometric object is therefore the projective normalized state together with a chosen integral lattice in its cone. Equivalently, it is a line in the normalized Asai representation equipped with the record framing $(v)^{-1}$. A positivity theorem for the projective flat bundle can control directions and equality cases; it cannot, without an additional framed-slope assertion, bound $\log|\mathrm{N}\, v|$.

There is an exact integral description of that framing. It also fixes the direction in which a slope argument would have to run. Write

$$N(X) = X^{\otimes 2} \otimes \tau X \otimes \tau^2 X, \qquad \mathcal{R}_v = v^{-1}\mathcal{O}_K,$$

and let $\mathbb{W}_{\mathcal{O}}$ denote the integral lattice in (400).

Proposition 17.49 (The record lattice is a saturated quotient)**.** *For every primitive $X = (u, v)^t$, evaluation on $\mathbf{U}(X)$ defines a surjection of fractional $\mathcal{O}_K$-modules*

$$q_X : \mathbb{W}_{\mathcal{O}}^{\vee} \longrightarrow \mathcal{R}_v, \qquad \ell \longmapsto \ell\big(\mathbf{U}(X)\big). \tag{406}$$

If $A \in G_7$, $X' = AX = (u', v')^t$, and neither denominator vanishes, then

$$q_{X'}(\ell) = \frac{v}{v'} q_X\big(\ell \circ \rho_{\mathbb{W}}(A)\big). \tag{407}$$

Multiplication by v/v' is an exact lattice isomorphism

$$\mu_{v,v'} : \mathcal{R}_v \xrightarrow{\sim} \mathcal{R}_{v'}. \tag{408}$$

With the standard real metrics on the two record lines, its logarithmic archimedean norm satisfies

$$\sum_{i=0}^{2} \log|\sigma_i(v/v')| = \log|\mathrm{N}\, v| - \log|\mathrm{N}\, v'|. \tag{409}$$

Thus the increase of the Arakelov degree of the record line is exactly the negative archimedean degree of its transport map.

Proof. The coordinates of $\mathbf{U}(X) = N(X)/v$ generate the fractional ideal $(v)^{-1}$, by Proposition 17.47. The image of (406) is therefore all of $\mathcal{R}_v$. Equation (403) is equivalent to

$$\mathbf{U}(X') = \frac{v}{v'} \rho_{\mathbb{W}}(A)\mathbf{U}(X).$$

Evaluation by ℓ proves (407). Since $(v/v')v^{-1}\mathcal{O}_K = v'^{-1}\mathcal{O}_K$, the bottom map is an exact lattice isomorphism. Finally, the last formula is the product formula at the three real embeddings. □

The natural tensor norm at all three real embeddings can also be computed without an estimate. This calculation separates the ordinary symmetric height from the one genuinely asymmetric Hodge term. Number indices cyclically, and put

$$E_i(X) = \log\|\sigma_i \mathbf{U}(X)\|_{\infty}.$$

Lemma 17.50 (Symmetric size and transverse Hodge imbalance)**.** *For every primitive state with* $v \neq 0$*,*

$$E_i(X) = 3h(X) + \eta_i, \tag{410}$$

$$E_0(X) + E_1(X) + E_2(X) = 12h(X) - \log|\mathrm{N}\, v|, \tag{411}$$

$$2E_0(X) - E_1(X) - E_2(X) = 2\eta_0 - \eta_1 - \eta_2. \tag{412}$$

In particular, if $S = E_0 + E_1 + E_2$ *and* $\Delta = 2E_0 - E_1 - E_2$*, then*

$$\mathscr{E}(X) = E_0(X) = \frac{S(X) + \Delta(X)}{3}. \tag{413}$$

Proof. At the i-th real embedding, the tensor numerator has norm $M_i^2 M_{i+1} M_{i+2}$, and the denominator has absolute value $|\sigma_i(v)|$. Since $\sum_k \log M_k = 3h(X)$, this gives

$$E_i = 3h(X) + \log M_i - \log|\sigma_i(v)| = 3h(X) + \eta_i.$$

Summing and using $\eta_0 + \eta_1 + \eta_2 = 3h(X) - \log|\mathrm{N}\, v|$ proves the second identity. The third and fourth are linear combinations of the first. □

This proposition prevents a second possible misuse of ordinary slope theory. The record line is not an unexpectedly positive subline of the Asai lattice; it is a saturated quotient of its dual. Moreover, the height of the transition between two such quotients is already the full record jump. Consequently a generic quotient-slope inequality, or the formal fact that $\mathcal{R}_{v'}$ is the maximal-degree coordinate at a first record, merely returns (409). The needed strict term can only come from the fact that q_X is the particular decomposable quotient selected by the Hodge endpoint and the Rosen stopping rule. This is the precise meaning of a *Hodge-selected quotient* below. Lemma 17.50 adds that the natural Galois-symmetric tensor norm sees S, not the desired E_0. The missing conversion is exactly the transverse character Δ, whose Hodge-theoretic incarnation, up to the sign convention for norms, is $\mathcal{N}^{-1}$; its square is the effective line $\mathcal{N}^{-2} \simeq \mathcal{O}(3P_7)$. Thus the order-seven divisor is not an optional source of extra negativity: it is the unique displayed term which converts symmetric lattice control into the distinguished record energy.

The same cancellation appears if one tries to use the strongest ordinary Roth or Subspace Theorem directly on a cusp row. The following identity uses every finite place at once and shows that the difficulty is exactly critical, rather than a loss caused by a weak version of Roth.

Proposition 17.51 (The full adelic two-form product is exactly critical)**.** *Choose* $Q \in \mathcal{O}_K^2$ *such that* $H_P = (P, Q) \in \mathrm{SL}_2(\mathcal{O}_K)$*. For a nonzero row* $w \in K^2$*, put*

$$L_P(w) = wP, \qquad L_Q(w) = wQ, \qquad H(w) = \prod_z \|w\|_z,$$

where z *ranges over all places of* K*, the absolute values satisfy the product formula, and* $\|\cdot\|_z$ *is the max norm. If neither linear form vanishes, then*

$$\prod_z \frac{|L_P(w)|_z\, |L_Q(w)|_z}{\|w\|_z^2} = H(w)^{-2}. \tag{414}$$

In particular, the divisibility of the remainder $L_P(w)$ *at all its finite prime ideals does not create a Roth surplus in this product: its finite contribution is canceled exactly by the three real absolute values of the same algebraic number.*

Proof. Both $L_P(w)$ and $L_Q(w)$ are nonzero elements of K. Hence

$$\prod_z |L_P(w)|_z = \prod_z |L_Q(w)|_z = 1$$

by the product formula. The denominator in (414) is $H(w)^2$. This proves the identity. When $w \in \mathcal{O}_K^2$ is primitive, all finite max norms of w are one. A prime dividing $L_P(w)$ then makes the finite part smaller, but its contribution reappears in $\prod_{i=0}^{2} |\sigma_i L_P(w)|$; there is no unaccounted gain. □

For two variables the Roth threshold is the power 2. Thus (414) explains why replacing Roth's theorem by its optimal standard form does not settle the problem. A strict inequality must use the nonparallel Hodge weights: one central incidence has to be counted twice while the two outer incidences are counted once. Proposition 17.48 explains the other half of the problem: that extra central copy must remain attached to the affine record lattice, because projectivization cancels it. These two facts leave one viable form of a Roth–Hodge proof: a slope theorem for the framed affine object $\mathfrak{F}(X_j)$ at first-record times, with the projective zero-defect locus treated separately.

17.3 The relative integral-vector target

The preceding construction permits a precise extension of Filip's theorem to be stated without inventing another representation. Use the post-transition convention

$$X_j = (r_j, r_{j+1})^t.$$

Call j a record time if $|\mathrm{N}\, r_{j+1}| > \max_{k \le j} |\mathrm{N}\, r_k|$. The finite content of $\mathbf{U}(X_j)$ is then the reciprocal of the maximal coordinate ideal selected by the stopping rule.

Conjecture 17.52 (Relative integral-vector closing)**.** *There is a constant B_7, depending only on the heptagonal curve, with the following property. For every primitive initial state P and every record time j, at least one of the following holds:*

1. $$\mathscr{E}(X_j) \le 6h(P) + B_7; \tag{415}$$
2. *a remainder has vanished by time $j+1$;*
3. *two projective states among $X_0, \dots, X_j$ are equal.*

This is a quantitative equality-case version of the integral-vector theorem. The first alternative says that the movement of the reciprocal integral lattice has not overcome the transverse Hodge contraction. The second is the parabolic conclusion. The third is the neutral conclusion: the intervening holonomy preserves one rational line and is therefore special or parabolic. Unlike a statement of strict domination, Conjecture 17.52 retains the special closed geodesics which the classification is required to retain.

Theorem 17.53 (The relative integral-vector theorem is sufficient)**.** *Conjecture 17.52 implies the heptagonal cusp classification.*

Proof. Suppose that an orbit has neither vanished nor repeated. At every record time, only the first alternative of the conjecture is then possible. Since all η_i are nonnegative, the definition of $\mathscr{E}$ gives

$$\log |\mathrm{N}\, r_{j+1}| \le \mathscr{E}(X_j) \le 6h(P) + B_7.$$

Every nonrecord norm is bounded by an earlier record norm. Hence all remainder norms are bounded. Proposition 9.10 therefore forces termination or repetition, a contradiction. □

Proposition 17.54 (The equality locus has already been identified)**.** *If $m < n$ and*

$$[\mathbf{U}(X_m)] = [\mathbf{U}(X_n)] \quad \text{in } \mathbf{P}(\mathbb{W}),$$

then the group element carrying X_m to X_n fixes the initial endpoint projectively. Conversely, on a special period with multiplier $\varepsilon \in \mathcal{O}_K^\times$,

$$\mathbf{U}(\varepsilon X) = \mathrm{N}(\varepsilon)\mathbf{U}(X) = \pm\mathbf{U}(X).$$

Thus the zero-defect projective lines of the normalized state are exactly the lines which can produce the required unit-holonomy alternative.

Proof. Proposition 17.47 says that the projective tensor determines the projective state. Equality of the two tensors therefore gives $[X_m] = [X_n]$, and conjugating the intervening word by the prefix at time m gives a nonidentity stabilizer of the initial endpoint. The converse is the last calculation in Lemma 17.55, with $m = 2$. □

This separates discovery from proof. The integral object is $\mathbf{U}$; its projective equality locus is correct; its finite lattice is the reciprocal record ideal; and its distinguished Hodge numerator has weight $-2a_0 + a_1 + a_2$. These statements are proved. The one assertion not supplied by the construction is (415). A proof must show that the scalar in (403) cannot move indefinitely in its shape torus while its field norm is being selected by first records. The real Hodge cocycle splits that motion into one curved and one flat character; Proposition 17.59 below shows that their Galois completion is the one K-valued coordinate α_2. Controlling its full height is the relative maximal-slope estimate, and not another search for a contracting tensor.

There is a concrete four-part route to that estimate.

1. Work on the affine cone over the projective flat bundle carrying the rational endpoint, with the reciprocal record ideal retained as the framing in (405). The projective bundle alone forgets precisely this datum. Equivalently, retain the saturated quotient q_X in (406), not just its projective kernel.

2. Extend the normalized Asai lattice over the toroidal cusp boundary. Its height jump is already prescribed by the content $(r_{j+1})^{-1}$.

3. Use the fundamental Green–Higgs factor with divisor P_7. Corollary 17.31 proves that the factor with divisor $2P_7$ is its square and the determinant factor with divisor $3P_7$ is its cube.

4. Prove a stopped inequality coupling the Hodge-selected quotient, the curved cocycle $\mathcal{A}_2$, and the full adelic height of its arithmetic coordinate $\alpha_2(m) = \tau^2(m)/m$, stronger than the formal Harder–Narasimhan maximum of the record direct sum. Proposition 17.49 shows that the latter contains no gain. Proposition 17.59 shows that this one K-valued coordinate recovers the analytically flat character $\nu(m)$ from its Galois conjugates. A stopped degree bound for norm plus full α_2-height therefore gives Problem 17.65. Persistent zero defect must force the coordinate to come from one recurrent projective state; Proposition 17.54 then supplies the stabilizer and its unit holonomy.

The first three items have explicit models in this manuscript. The fourth is the remaining theorem. It is narrower than an arbitrary Harder–Narasimhan statement: it concerns one displayed representation, one fundamental divisor P_7, one explicit K-valued shape coordinate, and the first-record sections of one continued-fraction correspondence.

The construction cannot be improved merely by taking more copies of the central factor. For an integer $m \geq 0$, with the evident interpretation when $m = 0$, put

$$\mathbf{U}_m(X) = \frac{X^{\otimes m} \otimes \tau(X) \otimes \tau^2(X)}{v^{m-1}}. \tag{416}$$

Lemma 17.55 (Central tensor powers pay for their own contraction)**.** *For every $m \geq 0$,*

$$\log \|\sigma_0 \mathbf{U}_m(X)\|_\infty = \log |\operatorname{N} v| + m\eta_0 + \eta_1 + \eta_2 = 3h(X) + (m-1)\eta_0, \tag{417}$$

$$\rho_m(A)\mathbf{U}_m(X) = \left(\frac{v'}{v}\right)^{m-1} \mathbf{U}_m(X'). \tag{418}$$

If $X' = \epsilon X$ with $\epsilon \in \mathcal{O}_K^\times$, then $\mathbf{U}_m(X') = \mathrm{N}(\epsilon)\mathbf{U}_m(X) = \pm\mathbf{U}_m(X)$.

Proof. The tensor norm is $M_0^m M_1 M_2 / |\sigma_0 v|^{m-1}$. Substitution of $M_i = |\sigma_i v| e^{\eta_i}$ and the product formula gives (417). Equivariance of the numerator and the change from v^{m-1} to $(v')^{m-1}$ give (418). Finally $\epsilon^m \tau(\epsilon)\tau^2(\epsilon)/\epsilon^{m-1} = \mathrm{N}(\epsilon)$. □

Thus increasing m strengthens the worst Hodge weight from $-2a_0 + a_1 + a_2$ to $-ma_0 + a_1 + a_2$, but it simultaneously raises the moving incidence factor to the power $m - 1$. The two changes are the same identity in (417); no symmetric-power trick creates a stopping surplus. Any improvement must use additional geometry which changes Hodge weight without increasing the moving pole. The effective order-seven Higgs divisor is the only such mechanism identified in this manuscript.

It does not yet permit lattice discreteness to be invoked, for a precise reason. The first real component of an $\mathcal{O}_K$-lattice is not discrete by itself. The other two components are its Galois conjugates. With $\deg \mathcal{L}_0 = 5/28$, $\deg \mathcal{L}_1 = 1/28$, and $\deg \mathcal{L}_2 = 3/28$, the three conjugate Hodge characters and their degrees are

	character	degree
$\mathbb{W}_1$	$(-1,1,0),\ (0,-1,1),\ (1,0,-1)$	$-1/7,\ 1/14,\ 1/14$
$\mathbb{W}_2$	$(-1,0,1),\ (1,-1,0),\ (0,1,-1)$	$-1/14,\ 1/7,\ -1/14.$

(419)

Each row has total degree zero. Tensoring the two selected negative components gives the doubled-central character $(-2,1,1)$, but its cyclic conjugates have degrees $-3/14, 3/14, 0$. Thus Galois descent necessarily restores exactly as much positive degree as the distinguished component loses.

This is the structural obstruction, rather than a shortage of contraction. For an arbitrary $\mathcal{O}_K$-integral sequence, a distinguished component can tend to zero while a conjugate component tends to infinity; powers of a unit do exactly this. In the desired special case that compensation is not noise. If $gP = \epsilon P$, then

$$\rho_i(g) w_i(P) = \epsilon \tau^i(\epsilon) w_i(P),$$

so the projective tensor repeats while its three real norms are rescaled by a unit. What remains to prove is therefore an orbit-specific statement: unit compensation for the two tensors in (393) must either stay in a bounded projective set or come from one cyclic holonomy fixing $[P]$. The radial remainder norm and the three-component Asai shape below are precisely the finite and archimedean data needed to formulate that statement.

There is a second trap in the language of principal bundles. Let $C = G_7 \backslash \mathbb{H}$, and let $\mathscr{F}$ be the flat bundle whose fiber is the product of the three real flag varieties. The Hodge flag

$$z \longmapsto (z, F_1(z), F_2(z))$$

is equivariant and therefore gives an honest Borel reduction of $\mathscr{F}$ over C. The rational flag $\Delta_K(x) = (\sigma_0 x, \sigma_1 x, \sigma_2 x)$ is not an equivariant section over C. It descends only after passing to the stabilizer cover

$$C_x = \mathrm{Stab}_{G_7}(x)\backslash\mathbb{H}. \tag{420}$$

Indeed, constancy of that section in a flat trivialization is equivalent to invariance of $\Delta_K(x)$ under the deck group.

For a boundary point of a Fuchsian group the stabilizer is either trivial, parabolic cyclic, or hyperbolic cyclic. Consequently the theorem sought here says exactly that C_x is never the universal cover: it is a cusp cylinder or a hyperbolic cylinder. This observation rules out a tempting but circular use of ordinary Harder–Narasimhan theory. The rational Borel needed for such a comparison is not a reduction on the compact base; in the case that must be excluded it exists only on $\mathbb{H}$. What is needed is instead a boundary, or stopped, instability theorem comparing the Hodge reduction with the flat rational flag along the orbit in (420). The incidence remainders are its generalized minors, and the two characters χ_1, χ_2 are its two numerical slopes. A compact-base stability theorem which starts with both reductions has already assumed the essential conclusion.

The integral realization of these two coordinates is already present in the three Asai components of (267). Choose the two Galois automorphisms τ_i by $\sigma_i = \sigma_0 \circ \tau_i$. For a nonzero incidence remainder r, set

$$a_i(r) = \frac{\tau_i(r)}{r} \in K^\times, \qquad i = 1, 2. \tag{421}$$

Equation (272) gives

$$(a_i(r)) = \frac{I(\mathcal{Z}_i)}{I(\mathcal{Z}_0)}. \tag{422}$$

Moreover

$$\log|\sigma_0 a_i(r)| = \log|\sigma_i(r)| - \log|\sigma_0(r)|. \tag{423}$$

Thus the direct sum $\mathcal{Z}_0 \oplus \mathcal{Z}_1 \oplus \mathcal{Z}_2$, modulo common scale, is an integral realization of the two-dimensional norm-one torus in (392). This construction is stable under the Galois action and the full Asai representation of G_7.

The distinction from the scalar tensor is exact. The line $\mathcal{Z}_1 \otimes \mathcal{Z}_2/s$ has ideal $(r)^{-1}$, so it remembers the radial norm needed in Proposition 13.2; but its archimedean character is only the determinant $u_1 + u_2$. Projectivizing all three components retains the ordered pair (u_1, u_2). This supplies a canonical stopped candidate with a radial line and a rank-two Asai shape, rather than one line asked to do both jobs.

Proposition 17.56 (Radial norm plus Asai shape is proper)**.** *Fix $B, C > 0$. There are only finitely many nonzero $r \in \mathcal{O}_K$ satisfying*

$$|\,\mathrm{N}(r)| \le B, \qquad |\log|\sigma_i(r)| - \log|\sigma_0(r)|| \le C \quad (i = 1, 2). \tag{424}$$

Consequently, if these bounds hold for both entries of every state in an infinite sequence of reduced Rosen returns, two complete states repeat and the endpoint is special. If a remainder vanishes, the endpoint is a cusp.

Proof. Put $t_i = \log|\sigma_i(r)|$, $s = \log|\,\mathrm{N}(r)|$, and $u_i = t_i - t_0$. Since r is a nonzero algebraic integer, $0 \le s \le \log B$. The inverse formulas are

$$t_0 = \frac{s - u_1 - u_2}{3}, \qquad t_1 = \frac{s + 2u_1 - u_2}{3}, \qquad t_2 = \frac{s - u_1 + 2u_2}{3}.$$

Thus all three real conjugates of r lie in a fixed bounded interval. The Minkowski image of $\mathcal{O}_K$ is a lattice in $\mathbf{R}^3$, so only finitely many r occur. Apply this to both coordinates of a state. Repetition gives a nontrivial element of the stabilizer Borel by Proposition 7.1; its alternatives were classified in the proof of that proposition. □

This proposition gives a more exact target than "use the thin group." One sufficient route is a stopped maximal-slope estimate for the projectivized three-component Asai lattice, while the incidence line bounds the common radial norm. A determinant estimate proves only the second assertion unless it is accompanied by rigidity along its kernel. A word calculation which does not descend to the Asai torus proves neither. The remaining comparison is between the two positive curvature forms in (462) and the two integral ideal ratios in (422).

The Higgs-ratio section makes this identification algebraic. Maximality of the uniformizing Higgs field gives $\mathcal{L}_0^2 \simeq \Omega^1_{\log}$, and hence

$$\Omega^1_{\log} \otimes \mathcal{M}^{-2} = \mathcal{N}^{-2}. \tag{425}$$

The section Ξ is therefore a section of $\mathcal{N}^{-2}$. Its divisor is three times the order-seven orbifold point, so, as orbifold line bundles,

$$\mathcal{N}^{-2} \simeq \mathcal{O}(3P_7), \qquad \mathcal{N}^{2} \simeq \mathcal{O}(-3P_7). \tag{426}$$

The degree of the first identity is $3/7$, agreeing with $-2\deg\mathcal{N}$. This agreement is an internal check, not the proof of the identity: the proof uses the Higgs-field construction of Ξ and its explicit divisor.

Equations (288)– (426) locate the real obstruction. The cusp-framed relative incidence class supplies the integral value r_j and its boundary Green function, but has zero curvature on the open curve. The line $\mathcal{N}$ supplies strict interior contraction, but is a Hodge line, not a flat integral local system. The missing theorem must construct their tensor product as one evaluated integral line on the Rosen correspondence. Taking the divisor from the first object and the norm estimate from the second without constructing this tensor product is precisely the invalid step behind the earlier scalar arguments.

There is a natural candidate tensor product on the Rosen return groupoid. Writing it down removes one ambiguity, but does not yet construct the transport morphism needed for the proof. On the universal projective flat bundle $\widetilde{\mathcal{Y}} = \mathbb{H} \times \mathbf{P}(V^\vee)$, let D_c be the hyperplane $\ell(c) = 0$. For a prefix g, put $c_g = g^{-1}e$. The quotient

$$s_g(\ell) = \frac{\ell(c_g)}{\ell(e)} \tag{427}$$

is the canonical meromorphic section of $\mathcal{O}(D_{c_g} - D_e)$. It is independent of the representative of the cusp because the peripheral group fixes e. At the horizontal endpoint covector it gives

$$s_g(\ell_P) = \frac{\ell_P(g^{-1}e)}{\ell_P(e)} = \frac{r_g}{r_0}. \tag{428}$$

Thus its finite divisor is exactly $(r_g) - (r_0)$; the second term is a fixed part of the initial height budget. The divisor is well defined projectively, but its degree at a selected record is not: to retain that degree one must pull s_g to the tautological affine cone and keep its integral trivialization. This is the same framing distinction as Proposition 17.48.

The incidence section has a further interpretation which uses the whole group and was hidden by the remainder notation. Write

$$g = \begin{pmatrix} a & b \\ c & d \end{pmatrix}, \qquad j_g(x) = cx + d.$$

Lemma 17.57 (Incidence is the denominator cocycle)**.** *On the affine chart $x = p/q$, one has*

$$s_g(\ell_P) = \frac{r_g}{r_0} = \frac{cp + dq}{q} = j_g(x). \tag{429}$$

Consequently

$$j_{gh}(x) = j_g(hx) j_h(x). \tag{430}$$

If x is fixed by g, then $j_g(x) = \mu \in \mathcal{O}_K^\times$, and the Galois Hodge character of this period is

$$\Theta_g^{\mathrm{rat}}(x) = \log|\sigma_1\mu| + \log|\sigma_2\mu| - 2\log|\sigma_0\mu| = -3\log|\sigma_0\mu|. \tag{431}$$

It vanishes for a parabolic period and records the unit direction of a special period.

Proof. Since $g^{-1}e = (d, -c)^t$ and $\ell_P = (-q, p)$,

$$r_g = \ell_P(g^{-1}e) = -(cp + dq), \qquad r_0 = -q,$$

which proves (429). The second identity is the standard denominator identity for composition of fractional linear maps. At a fixed point, $(x, 1)^t$ is an eigenvector with eigenvalue $\mu = j_g(x)$; integrality and determinant one make μ a unit. The last equality follows from the product formula for that unit. □

This lemma locates the missing transport more accurately. The algebraic multiplier is not absent: it is already the incidence r_g/r_0. What is absent is a proved comparison between its rational-flag Hodge character and the Hodge character of the modular embedding. To write that comparison, set

$$\rho_i(g) = \begin{pmatrix} a_{i,g} & b_{i,g} \\ c_{i,g} & d_{i,g} \end{pmatrix}, \qquad x_i^\dagger = \varepsilon_i \sigma_i(x), \qquad j_{i,g}^\dagger(w) = c_{i,g} w + d_{i,g}, \qquad F_0(z) = z.$$

The sign normalization was chosen precisely so that

$$j_{i,g}^\dagger(x_i^\dagger) = \sigma_i(j_g(x)). \tag{432}$$

For $i = 1, 2$, define the two relative characters

$$\Theta_{i,g}^{\mathrm{rat}}(x) = \log|\sigma_i(j_g(x))| - \log|\sigma_0(j_g(x))|, \tag{433}$$

$$\Theta_{i,g}^{\mathrm{Hod}}(z) = \log|j_{i,g}^\dagger(F_i z)| - \log|j_{0,g}^\dagger(z)|, \tag{434}$$

$$\mathcal{A}_{i,g}(z, x) = \Theta_{i,g}^{\mathrm{rat}}(x) - \Theta_{i,g}^{\mathrm{Hod}}(z). \tag{435}$$

Their determinant characters are the sums

$$\Theta_g^{\mathrm{rat}} = \Theta_{1,g}^{\mathrm{rat}} + \Theta_{2,g}^{\mathrm{rat}}, \tag{436}$$

$$\Theta_g^{\mathrm{Hod}} = \Theta_{1,g}^{\mathrm{Hod}} + \Theta_{2,g}^{\mathrm{Hod}} = \log|j_{1,g}^\dagger(F_1 z)| + \log|j_{2,g}^\dagger(F_2 z)| - 2\log|j_{0,g}^\dagger(z)|, \tag{437}$$

$$\mathcal{A}_g = \mathcal{A}_{1,g} + \mathcal{A}_{2,g} = \Theta_g^{\mathrm{rat}}(x) - \Theta_g^{\mathrm{Hod}}(z). \tag{438}$$

Equivariance of the F_i and (430) show that $\mathcal{A}$ is an exact additive cocycle on the rational-flag bundle over the triangle surface. In fact each $\mathcal{A}_i$ is separately an exact additive cocycle. Geometrically it is the automorphy of the ratio between the tautological rational-flag line of character $(-2, 1, 1)$ and the Hodge line $\mathcal{N}$. This is a canonical, group-wide candidate for the norm of the normalized transport. In fact its Green function and curvature can be computed exactly.

Corollary 17.31 now separates this stopped cocycle into one curved coordinate and one flat arithmetic coordinate. Set

$$\Theta^{\mathrm{rat}}_{\mathrm{neut},g}(x) = \Theta^{\mathrm{rat}}_{1,g}(x) - 2\Theta^{\mathrm{rat}}_{2,g}(x), \tag{439}$$

$$\Theta^{\mathrm{Hod}}_{\mathrm{neut},g}(z) = \Theta^{\mathrm{Hod}}_{1,g}(z) - 2\Theta^{\mathrm{Hod}}_{2,g}(z). \tag{440}$$

Proposition 17.30 supplies a smooth potential $B : \mathbb{H} \to \mathbf{R}$, bounded on a fixed cusp-adapted fundamental cross-section, such that

$$\Theta^{\mathrm{Hod}}_{\mathrm{neut},g}(z) = B(gz) - B(z). \tag{441}$$

Consequently the relative cocycles satisfy the exact identity

$$\mathcal{A}_g = 3\mathcal{A}_{2,g} + \Theta^{\mathrm{rat}}_{\mathrm{neut},g}(x) - B(gz) + B(z). \tag{442}$$

In particular, between reduced returns the final two terms differ from the rational neutral character by $O_7(1)$.

The neutral character itself has no hidden Hodge content. If $m = j_g(x) \in K^\times$ and τ is numbered by $\sigma_i = \sigma_0 \circ \tau^i$, then

$$\begin{aligned}\Theta^{\mathrm{rat}}_{\mathrm{neut},g}(x) &= \log\left|\frac{\sigma_0(m)\sigma_1(m)}{\sigma_2(m)^2}\right| \\ &= \log|\sigma_0\nu(m)|, \qquad \nu(m) = \frac{m\tau(m)}{\tau^2(m)^2} = \frac{\mathrm{N}_{K/\mathbf{Q}}(m)}{\tau^2(m)^3}.\end{aligned} \tag{443}$$

Lemma 17.58 (Radial, curved, and flat coordinates are complete)**.** *Put*

$$a_i = \log|\sigma_i(m)|, \qquad s = a_0 + a_1 + a_2 = \log|\,\mathrm{N}(m)|,$$

and set

$$c = \Theta^{\mathrm{rat}}_{2,g} = a_2 - a_0, \qquad f = \Theta^{\mathrm{rat}}_{\mathrm{neut},g} = a_0 + a_1 - 2a_2.$$

Then

$$a_0 = \frac{s-f}{3} - c, \qquad a_1 = \frac{s+2f}{3} + c, \qquad a_2 = \frac{s-f}{3}. \tag{444}$$

Thus the field norm, the one curved rational character, and the one flat rational character recover all three real sizes of the multiplier.

Proof. The definition of f gives $f = s - 3a_2$, so $a_2 = (s-f)/3$. The definition of c then gives $a_0 = a_2 - c$, and $a_1 = s - a_0 - a_2$. These are the three displayed formulas. □

This elementary inversion is the numerical core of the hybrid formulation. At a new record the stopping rule constrains the radial quantity s. The Hodge line $\mathcal{N}_2$ compares the curved quantity c with a metric of positive curvature. The quantity f has no curvature and is the logarithm of $\nu(m)$. Bounding signed values of s, c, f at one real place is still not a height theorem: a Weil height also contains their positive parts at every archimedean and finite place. The next two propositions give the all-place version and show exactly what must be controlled.

Proposition 17.59 (The fundamental Hodge character is a proper arithmetic shape coordinate)**.** *Define*

$$\alpha_2(m) = \frac{\tau^2(m)}{m} \in K^\times. \tag{445}$$

Then

$$\Theta^{\mathrm{rat}}_{2,g}(x) = \log|\sigma_0\alpha_2(m)|, \qquad \nu(m) = \alpha_2(m)^{-2}\tau(\alpha_2(m))^{-1}. \tag{446}$$

The first identity identifies α_2 with the rational character paired with the fundamental Hodge line $\mathcal{N}_2$. Modulo rational scalars, the map $m \mapsto \alpha_2(m)$ induces the isogeny

$$\bar{\alpha}_2 = \tau^2 - 1 : \mathbb{T}^{\mathrm{sh}}_K \longrightarrow \mathbb{T}^{\mathrm{sh}}_K, \qquad \deg \bar{\alpha}_2 = 3. \tag{447}$$

Consequently the full height of $\alpha_2(m)$, not merely the one real logarithm displayed above, controls the entire non-radial shape of m. In particular, the flat coordinate $\nu(m)$ is not an independent arithmetic direction: it is a fixed monomial in the Galois orbit of the fundamental curved coordinate.

Proof. The first identity is

$$\log|\sigma_0(\tau^2 m/m)| = \log|\sigma_2 m| - \log|\sigma_0 m|.$$

If $a = \alpha_2(m)$, then

$$a^{-2}\tau(a)^{-1} = \left(\frac{m}{\tau^2 m}\right)^2 \left(\frac{\tau m}{m}\right) = \frac{m\tau m}{(\tau^2 m)^2} = \nu(m),$$

which proves (446). Rational scalars cancel from α_2, so it descends to the shape torus. On its rank-two character lattice the eigenvalues of $\tau^2 - 1$ are $\zeta^2 - 1$ and $\zeta - 1$, where $\zeta^3 = 1$ and $\zeta \neq 1$. Their product is 3. Hence the endomorphism is an isogeny of degree three. A finite morphism pulls an ample height back to an ample height; the asserted properness follows. □

This identifies the arithmetic object which must be put under the Hodge metric. Its finite divisor is simply

$$(\alpha_2(m)) = \frac{(\tau^2 m)}{(m)}, \tag{448}$$

the same ratio of Galois-paired Asai component ideals that appears in (422). Its distinguished real norm is the character of $\mathcal{N}_2$, and its other two real norms are its Galois conjugates. Thus the finite divisor, the curved Hodge norm, and the full shape height belong to one K-line; no auxiliary flat line has to be invented. The remaining issue is to prove a stopped degree bound for that line along the thin orbit.

There is also no indexing ambiguity in this formulation. For a nonzero state denominator v, put

$$\mathfrak{a}_2(v) = \frac{\tau^2(v)}{v}.$$

If $X' = AX$ and $m = v'/v$ is the scalar in (403), then

$$\alpha_2(m) = \frac{\mathfrak{a}_2(v')}{\mathfrak{a}_2(v)}. \tag{449}$$

Thus the fundamental multiplier is exactly the jump of the integral Asai shape coordinate across the same transition; it is not evaluated one step before or after the record line. Problem 17.65 asks for bounded all-place height of this jump along a record subsequence.

Proposition 17.60 (The neutral divisor is the cyclic ideal Laplacian)**.** *Let $m \in K^\times$, and exclude the unique rational prime ramified in $K/\mathbf{Q}$. If a rational prime is inert, its valuation in $\nu(m)$ is zero. If it splits as*

$$p\mathcal{O}_K = \mathfrak{p}_0\mathfrak{p}_1\mathfrak{p}_2, \qquad \mathfrak{p}_j = \tau^j\mathfrak{p}_0,$$

put $e_j = v_{\mathfrak{p}_j}(m)$ *and* $n_j = v_{\mathfrak{p}_j}(\nu(m))$. *Then*

$$(n_0, n_1, n_2) = (e_0 + e_2 - 2e_1,\ e_1 + e_0 - 2e_2,\ e_2 + e_1 - 2e_0). \tag{450}$$

In particular,

$$n_0 + n_1 + n_2 = 0, \qquad (n_0, n_1, n_2) = 0 \Longleftrightarrow e_0 = e_1 = e_2.$$

Consequently $\nu(m)$ *is a unit if and only if the fractional ideal* (m) *is Galois invariant, apart from the fixed ramified prime. Equivalently,*

$$(m) = (a)\mathfrak{p}_7^k \tag{451}$$

for some $a \in \mathbf{Q}^\times$ *and an integer* k, *where* $\mathfrak{p}_7$ *is the ramified prime. Modulo the relation* $\mathfrak{p}_7^3 = (7)$, *one may take* $k \in \{0, 1, 2\}$.

Proof. For any prime $\mathfrak{p}$,

$$v_{\mathfrak{p}}(\tau^j m) = v_{\tau^{-j}\mathfrak{p}}(m).$$

At an inert prime all three primes on the right are the same, so the coefficient $1 + 1 - 2$ makes the valuation of $\nu(m)$ zero. At a split prime, substitution gives (450). Its kernel over $\mathbf{Z}^3$ is the diagonal. This proves the local assertions.

Away from the ramified prime, a fractional ideal is Galois invariant exactly when its exponents are constant on every split orbit; inert primes are already invariant. Such an ideal is extended from a fractional ideal of $\mathbf{Q}$. The only remaining factor is a power of the unique ramified prime, which proves (451). The converse follows by applying $1 + \tau - 2\tau^2$ to a Galois-invariant divisor. □

Proposition 17.61 (The neutral coordinate is a proper torus isogeny)**.** *Let*

$$\mathbb{T}_K^{\mathrm{sh}} = \big(\mathrm{Res}_{K/\mathbf{Q}}\,\mathbb{G}_m\big)/\mathbb{G}_m,$$

where the denominator is the diagonal rational-scalar torus. The map $m \mapsto \nu(m)$ *is unchanged by* $m \mapsto am$ *for* $a \in \mathbf{Q}^\times$, *and therefore induces an isogeny*

$$\bar{\nu} : \mathbb{T}_K^{\mathrm{sh}} \longrightarrow \mathbb{T}_K^{\mathrm{sh}}.$$

In multiplicative notation,

$$[\nu(m)] = [\tau^2(m)^{-3}] \quad \text{in } \mathbb{T}_K^{\mathrm{sh}}, \tag{452}$$

so $\bar{\nu} = [-3] \circ \tau^2$ *and*

$$\deg \bar{\nu} = 9. \tag{453}$$

For any τ*-invariant symmetric ample toric height* h_{sh},

$$h_{\mathrm{sh}}(\bar{\nu}(t)) = 3h_{\mathrm{sh}}(t) + O_K(1). \tag{454}$$

Consequently bounded height of $\nu(m)$ *is equivalent to bounded non-radial shape of* m *modulo* $\mathbf{Q}^\times$; *it controls both unit directions and all Galois-asymmetric ideal directions simultaneously.*

Proof. The first assertion follows either directly from the definition or from

$$\nu(m) = \mathrm{N}_{K/\mathbf{Q}}(m)\tau^2(m)^{-3}.$$

The norm is rational and disappears in the quotient torus, proving (452). The character lattice of $\mathbb{T}_K^{\mathrm{sh}}$ has rank two. Cubing acts on it with determinant 3^2, while τ^2 is an automorphism; this proves (453). Choose a symmetric ample toric compactification and average its boundary divisor over the order-three action of τ. The resulting height is τ-invariant up to $O_K(1)$. Standard functoriality of heights under the cubing map gives (454). Since an isogeny is finite, the resulting height is proper modulo its finite kernel. □

Corollary 17.62 (Radial norm plus fundamental Hodge shape is a proper multiplier height)**.** *The homomorphism of* $\mathbf{Q}$*-tori*

$$\mathfrak{M}_2 : \mathrm{Res}_{K/\mathbf{Q}}\,\mathbb{G}_m \longrightarrow \mathbb{G}_m \times \mathbb{T}_K^{\mathrm{sh}}, \qquad m \longmapsto \big(\mathrm{N}_{K/\mathbf{Q}}(m), \bar{\alpha}_2([m])\big) \tag{455}$$

is an isogeny of degree 9*. Consequently, for compatible ample heights there are constants* $c_1, c_2 > 0$ *and* C_K *such that*

$$c_1 h_K(m) - C_K \le h_{\mathbf{Q}}\big(\mathrm{N}(m)\big) + h_{\mathrm{sh}}\big(\bar{\alpha}_2([m])\big) \le c_2 h_K(m) + C_K. \tag{456}$$

Replacing $\bar{\alpha}_2$ *by* $\bar{\nu}$ *gives the earlier degree-*27 *isogeny and an equivalent proper height.*

Proof. On the diagonal scalar torus, the norm map is cubing and hence has degree 3. On the quotient by that torus, Proposition 17.59 gives an isogeny of degree 3. The character-lattice determinant of (455) is therefore $3 \cdot 3 = 9$. The pullback of an ample height under a finite morphism is ample, and all ample heights on a fixed projective compactification are comparable. This gives (456). The final assertion follows from Proposition 17.61; equivalently use the monomial identity (446) and finiteness of both shape isogenies. □

Corollary 17.63 (The multiplier height is the complete state height)**.** *Let* $X = (u, v)^t \in K^2$ *with* $u \neq 0$*, and put* $m = v/u$*. Then*

$$h_K(m) = h([1 : m]) = h([u : v]). \tag{457}$$

Consequently the height in the middle of (456) *is bounded if and only if the complete projective state* $[u : v]$ *has bounded height.*

Proof. The two projective points are equal because $[u : v] = [1 : v/u]$. The first equality is the definition of the ordinary height of a field element. The final assertion follows from Corollary 17.62. □

Theorem 17.64 (Closing from bounded record-multiplier height)**.** *Let*

$$x_j = \frac{r_{j+1}}{r_j} \in K^\times$$

be the nonterminal Rosen state ratios. Suppose that either there are only finitely many norm records, or there is a constant C_P *and infinitely many record indices* j *for which*

$$h_{\mathbf{Q}}\big(\mathrm{N}(x_j)\big) + h_{\mathrm{sh}}\big(\bar{\alpha}_2([x_j])\big) \le C_P. \tag{458}$$

Then the initial endpoint is a cusp or a special hyperbolic fixed point.

Proof. If there are only finitely many records, all remainder norms are bounded. Proposition 9.10 then gives termination or repetition.

In the second case, Corollary 17.62 and Northcott's theorem leave only finitely many possible values of x_j at the indicated record indices. Two of them are equal. The Rosen map is deterministic on its state ratio, so the expansion from the first occurrence onward is periodic. Equivalently, the intervening projectively nonidentity word fixes the initial endpoint. Proposition 7.1 gives the cusp or special alternative. □

Problem 17.65 (Stopped multiplier compactness)**.** *For every primitive initial state P, prove that either its Rosen orbit has only finitely many norm records, or there is a constant C_P such that infinitely many record indices satisfy (458). Equivalently, when the set $\mathcal{R}(P)$ of record indices is infinite, prove*

$$\liminf_{j\in\mathcal{R}(P)} \left\{ h_{\mathbf{Q}}\big(\mathrm{N}(x_j)\big) + h_{\mathrm{sh}}\big(\bar{\alpha}_2([x_j])\big) \right\} < \infty.$$

One height in three languages. For a nonterminal complete state $X = (u, v)^t$, put $m = v/u$. The quantity which has to return to a bounded set can now be written in three ways:

$$\begin{aligned} \text{projective state:} \quad & h([u : v]) = h_K(m), \\ \text{radial plus shape:} \quad & h_{\mathbf{Q}}(\mathrm{N}\, m) + h_{\mathrm{sh}}([\tau^2(m)/m]) \asymp h_K(m), \\ \text{Hodge compensation at compact returns:} \quad & \mathcal{D}_P(t) = 6h([g_t P]) + O(1). \end{aligned}$$

The first equality is elementary projective geometry. The second is the degree-nine torus isogeny and is algebraic number theory. The third is the comparison of the integral Asai norm with the Hodge metric. The missing assertion is not a comparison among these quantities: those comparisons are proved. It is the thin-group recurrence statement that the particular sequence selected by one rational flag returns to a bounded set for this height, unless its rational Borel already has a stabilizer.

Where the three structures enter. The torus isogeny is algebraic number theory: it proves that the norm and the full α_2-height together lose no part of the multiplier. The identity $\mathcal{N} \simeq \mathcal{N}_2^3$ is Hodge theory: it proves that all nonzero curvature is concentrated in the one P_7-direction. Neither statement proves Problem 17.65. Its proof must use the fact that the multipliers arise from one rational flag under the whole thin group G_7. Arbitrary elements of $K^\times$ can escape along units forever, so no field-theoretic inequality can supply the conclusion without this orbit restriction. Conversely, calculations for individual Rosen generators do not see the global torus height. The missing theorem is therefore a genuine hybrid: a Hodge height-jump or relative-domination statement for the tautological denominator cocycle of the thin G_7-orbit, with recurrent Borel orbits as its equality case.

The exact wall. The Galois-completed Hodge degrees of the fundamental coordinate are

$$-\frac{1}{14}, \qquad \frac{1}{7}, \qquad -\frac{1}{14};$$

this is the second row of (419). Their sum is zero. The two negative entries are the genuine Hodge contraction, while the middle conjugate can pay for both of them. A Weil height takes positive parts at all places and therefore sees precisely that compensating conjugate. Consequently neither total Arakelov degree nor ordinary semistability has a strict sign. The theorem still needed is orbit-specific: if the positive conjugate repeatedly compensates the two negative conjugates at first-record times, then the rational flag must recur and hence have a nontrivial G_7-stabilizer. This is the point at which the present argument stops; all earlier tensor, divisor, and torus constructions have now been made on the same indexed line.

This criterion includes the bounded-record-norm route, because finitely many records are its first alternative. Its second alternative asks only for a bounded subsequence of complete state multipliers; it does not ask for an a priori bound on every earlier remainder. On a special period, unit holonomy can make individual real sizes of integral representatives grow, although their field norms repeat and hence create only finitely many records. The criterion uses exactly the scalar already present in (403): the rational height of its norm is the radial part, and the height of α_2 is the entire non-radial part. Formula (446) then recovers ν automatically. Thus a proof may target

(458) directly instead of first proving (415). The remaining task is still substantive, but it is now a compactness statement for one explicit transport scalar rather than an absolute bound for the accumulated remainder divisor.

> **Elementary closing lemma, not the missing estimate.** The isogeny calculation and the application of Northcott are standard algebraic number theory. Their role is to identify the correct global quantity and to prove that bounding it is sufficient. They do not prove that the Rosen–Hodge dynamics supplies the bound in (458). Problem 17.65, with the periodic unit-holonomy equality case retained, is the substantive theorem still to be established.

Thus all positive curvature in the stopped transport is carried by the single P_7-line $\mathcal{N}_2^{-1}$. At the level of one distinguished real norm, the other term is the flat cube character $\nu(m)$; on a special period it records the permitted unit holonomy. At the algebraic level, however, Proposition 17.59 shows that this is not a second independent variable: it is a monomial in the Galois orbit of $\alpha_2(m)$, the arithmetic coordinate paired with $\mathcal{N}_2$. Proposition 17.60 explains what its distinguished flat logarithm measures—Galois-asymmetric ideal movement—and Proposition 17.61 shows that its full height is proper. The unresolved step is therefore sharper than "control a separate flat direction": upgrade the one-place Hodge control of α_2 to an all-place stopped height bound for the same K-valued coordinate. Galois descent then controls ν as well.

For $w \in \mathbb{H}$ and $\xi \in \mathbf{R}$, put

$$\Phi(w,\xi) = \log \frac{(\operatorname{Im} w)^2}{|w-\xi|^2}.$$

The other affine chart gives the definition at $\xi = \infty$. Set

$$G_{i,x}(z) = \Phi(F_i z, x_i^\dagger) - \Phi(z, x_0^\dagger), \qquad i = 1,2, \tag{459}$$

and hence

$$G_x(z) = \Phi(F_1 z, x_1^\dagger) + \Phi(F_2 z, x_2^\dagger) - 2\Phi(z, x_0^\dagger) = G_{1,x}(z) + G_{2,x}(z). \tag{460}$$

Proposition 17.66 (Explicit relative Green potential)**.** *For every $g \in G_7$,*

$$G_{i,gx}(gz) - G_{i,x}(z) = 2\mathcal{A}_{i,g}(z,x), \tag{461}$$

$$dd^c G_{i,x} = 2\big(\omega_{\mathrm{hyp}} - F_i^*\omega_{\mathrm{hyp}}\big) > 0. \tag{462}$$

Consequently,

$$G_{gx}(gz) - G_x(z) = 2\mathcal{A}_g(z,x). \tag{463}$$

If the hyperbolic form is normalized by $\omega_{\mathrm{hyp}} = -dd^c \log \operatorname{Im} w$, then

$$dd^c G_x = 2\big(2\omega_{\mathrm{hyp}} - F_1^*\omega_{\mathrm{hyp}} - F_2^*\omega_{\mathrm{hyp}}\big). \tag{464}$$

This form is strictly positive on the open triangle surface. Moreover, $\Phi(w,\xi) \le 0$, with equality exactly when $\operatorname{Re} w = \xi$.

Proof. For a real determinant-one Möbius transformation,

$$\operatorname{Im}(gw) = \frac{\operatorname{Im} w}{|j_g(w)|^2}, \qquad |gw - g\xi| = \frac{|w-\xi|}{|j_g(w) j_g(\xi)|}.$$

Therefore

$$\Phi(gw, g\xi) - \Phi(w,\xi) = 2\big(\log|j_g(\xi)| - \log|j_g(w)|\big).$$

Apply this identity first with weights $(-1,1,0)$ and $(-1,0,1)$ to obtain (461); summing gives (463). On $\mathbb{H}$, the function $\log|w-\xi|^2$ is harmonic. Hence $dd^c\Phi = -2\omega_{\mathrm{hyp}}$, which gives (462) and, after summing, (464). Strict Schwarz–Pick contraction makes each of the two component forms positive. Finally $|w-\xi|^2 \ge (\operatorname{Im} w)^2$, with equality precisely on the vertical geodesic above ξ. □

The identity

$$\Phi(w,\xi) = \log \operatorname{Im} w - \log \frac{|w-\xi|^2}{\operatorname{Im} w}$$

shows that this is exactly the Hodge potential minus the angular Busemann function of Lemma 14.2. Thus the real transport and its strict curvature are now explicit, rather than conjectural.

Proposition 17.67 (Integral model of the denominator transport)**.** *On the action groupoid of G_7 on $\mathbf{P}^1_K$, the finite part of the transport (463) is the standard integral tautological line. In the chart $q \neq 0$, its transition under $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ is*

$$\frac{q(gP)}{q(P)} = \frac{cp+dq}{q} = j_g(x) = \frac{r_g}{r_0}. \tag{465}$$

At every finite place v, $g \in \mathrm{SL}_2(\mathcal{O}_{K_v})$ preserves the tautological lattice. Hence the finite transition divisor is exactly $(r_g) - (r_0)$. Together with the archimedean metric G_x, these transitions define an adelically metrized line on the open rational-flag groupoid.

Proof. The second coordinate of $g(p,q)^t$ is $cp+dq$, proving (465). An integral unimodular matrix carries a primitive local vector to a primitive local vector, so the standard $\mathcal{O}_{K_v}$-lattice in the tautological line is preserved. The finite divisor assertion follows. Finally, (430) and (463) are respectively the finite and archimedean cocycle conditions needed to glue the metrized line. □

The metric also has the correct normalization in the standard cusp. Equivariance under T and positivity of the signed cusp widths give the usual modular-embedding Fourier expansion [40, Proposition 1.1]

$$F_i(z) = \frac{|\lambda_i|}{\lambda_0} z + B_{i,0} + O\left(e^{-2\pi \operatorname{Im} z/\lambda_0}\right) \qquad (\operatorname{Im} z \longrightarrow \infty). \tag{466}$$

This is the usual cusp expansion for a modular embedding. In particular, for every fixed finite $\xi \in \mathbf{R}$,

$$\Phi(F_i(z),\xi) = O\left((\operatorname{Im} z)^{-1}\right), \qquad \Phi(z,\xi) = O\left((\operatorname{Im} z)^{-1}\right).$$

The same statement in the other affine chart holds for $\xi = \infty$. Consequently

$$G_x(z) \longrightarrow 0 \qquad (\operatorname{Im} z \longrightarrow \infty). \tag{467}$$

Thus the open integral model and its canonical standard-cusp metric are now constructed. At a cusp $g^{-1}\infty$, the transition to this normalized chart is (463); its algebraic part is exactly r_g/r_0 by (465).

There is an explicit holomorphic form of the same construction. In the period coordinates above, the Higgs-ratio section Ξ of $\mathcal{N}^{-2}$ is represented, up to a constant, by $F_1'(z)F_2'(z)$. Normalize that constant by the leading terms in (466). For a rational flag x, define

$$\Psi_{i,x}(z) = \frac{F_i'(z)(z-x_0^\dagger)}{F_i(z)-x_i^\dagger}, \qquad i=1,2, \tag{468}$$

and

$$\Psi_x(z) = \Psi_{1,x}(z)\Psi_{2,x}(z) = \frac{F_1'(z)(z-x_0^\dagger)}{F_1(z)-x_1^\dagger}\,\frac{F_2'(z)(z-x_0^\dagger)}{F_2(z)-x_2^\dagger}. \tag{469}$$

Proposition 17.68 (Exact automorphy under the whole thin group)**.** *For every $g \in G_7$, the individual factors satisfy*

$$\Psi_{i,gx}(gz) = \Psi_{i,x}(z)\frac{j^\dagger_{0,g}(z)}{j^\dagger_{i,g}(F_i(z))}\frac{j^\dagger_{i,g}(x_i^\dagger)}{j^\dagger_{0,g}(x_0^\dagger)}, \qquad i = 1, 2. \tag{470}$$

Consequently

$$\frac{\Psi_{gx}(gz)}{\Psi_x(z)} = \frac{j^\dagger_{0,g}(z)^2}{j^\dagger_{1,g}(F_1(z))j^\dagger_{2,g}(F_2(z))}\frac{\sigma_1(j_g(x))\sigma_2(j_g(x))}{j_g(x)^2}. \tag{471}$$

In absolute logarithms, the right side is exactly $\mathcal{A}_g(z, x) = \Theta^{\mathrm{rat}}_g(x) - \Theta^{\mathrm{Hod}}_g(z)$. Thus it is a multiplicative cocycle of the complete G_7-action, not a collection of multipliers assigned to a preferred alphabet.

If $gx = x$, then $j_g(x) = \mu \in \mathcal{O}_K^\times$. The second factor in (471) is the unit character

$$\frac{\sigma_1(\mu)\sigma_2(\mu)}{\mu^2}, \tag{472}$$

whose logarithmic modulus is $-3\log|\sigma_0\mu|$. It is trivial for a parabolic period and is precisely the nontrivial unit holonomy of a special hyperbolic period.

Proof. The derivative of the equivariance relation and the elementary Möbius difference formula give

$$F_i'(gz) = F_i'(z)\frac{j^\dagger_{0,g}(z)^2}{j^\dagger_{i,g}(F_i(z))^2}, \qquad \frac{gz - gx}{\rho_i(g)F_i(z) - \rho_i(g)x_i^\dagger} = \frac{z - x}{F_i(z) - x_i^\dagger}\frac{j^\dagger_{i,g}(F_i(z))j^\dagger_{i,g}(x_i^\dagger)}{j^\dagger_{0,g}(z)j^\dagger_{0,g}(x_0^\dagger)}.$$

Their product proves (470); multiply over $i = 1, 2$ and use (432) to obtain (471). The cocycle assertion follows from composition. The last statement is Lemma 17.57 and the product formula. □

Proposition 17.69 (Normalized logarithmic-derivative factorization)**.** *With compatible Hodge frames,*

$$\log|\Psi_{i,x}(z)|^2 = G_{i,x}(z) + \log\|\Xi_i(z)\|^2, \qquad i = 1, 2. \tag{473}$$

Summing gives

$$\log|\Psi_x(z)|^2 = G_x(z) + \log\|\Xi(z)\|^2. \tag{474}$$

The zeros of $\Psi_{1,x}$ and $\Psi_{2,x}$ are respectively the lifts of $2P_7$ and P_7, independently of x. On the vertical geodesic $z = x_0^\dagger + iy$,

$$|\Psi_{i,x}(z)| \le 1 \quad (i = 1, 2), \qquad |\Psi_x(z)| \le 1, \tag{475}$$

with strict inequality in the interior, and each factor tends to one as $y \to \infty$.

Proof. Differentiating the modular equivariance gives

$$F_i'(gz) = \frac{j^\dagger_{0,g}(z)^2}{j^\dagger_{i,g}(F_i(z))^2}F_i'(z).$$

Thus F_i' has the automorphy of $\mathcal{N}_i^{-2}$; this is the local expression of Ξ_i. Substitution of the Hodge norm of $\mathcal{N}_i^{-2}$ and the definition of $G_{i,x}$ gives (473). Since the denominators in (468) never vanish

in $\mathbb{H}$, the zero divisors are those in (292). Multiplication gives the determinant identity and the divisor $3P_7$.

For $z = x_0^\dagger + iy$, Schwarz–Pick gives

$$|F_i'(z)| \le \frac{\operatorname{Im} F_i(z)}{\operatorname{Im} z}.$$

Moreover

$$|z - x_0^\dagger| = \operatorname{Im} z, \qquad |F_i(z) - x_i^\dagger| \ge \operatorname{Im} F_i(z).$$

Each $\Psi_{i,x}$ therefore has modulus at most one. Strictness follows from strict Schwarz–Pick. The cusp expansion gives limit one for each normalized factor. □

Proposition 17.70 (The controlled channel is a divisor-free ratio)**.** *Let $x, y \in \mathbf{P}^1(K)$, and define Ψ_x, Ψ_y by (469). In compatible affine charts,*

$$\frac{\Psi_y(z)}{\Psi_x(z)} = \left(\frac{z - y_0^\dagger}{z - x_0^\dagger}\right)^2 \prod_{i=1}^{2} \frac{F_i(z) - x_i^\dagger}{F_i(z) - y_i^\dagger}. \tag{476}$$

If $f_{x,i}, f_{y,i}$ are the corresponding period evaluations, then

$$\log\left|\frac{\Psi_y(z)}{\Psi_x(z)}\right| = 2\log\left|\frac{f_{y,0}(z)}{f_{x,0}(z)}\right| - \sum_{i=1}^{2} \log\left|\frac{f_{y,i}(z)}{f_{x,i}(z)}\right|, \tag{477}$$

up to the constant fixed by the projective representatives. Both Ψ_x and Ψ_y vanish on the same divisor $3P_7$, so their ratio has neither a zero nor a pole in $\mathbb{H}$. For $(x, y) = ([Q_m], [Q_n])$, the difference of (477) between t_m and t_n is exactly the controlled gain in (191).

Proof. The factors $F_1'F_2'$ cancel from the quotient of the two definitions (469), giving (476). A period evaluation of a real boundary column is, up to its fixed projective scalar, the corresponding linear factor $z - x_0^\dagger$ or $F_i(z) - x_i^\dagger$. Taking absolute logarithms proves (477). Proposition 17.69 gives the common zero divisor $3P_7$, while none of the boundary linear factors vanishes in the upper half-plane. The final assertion is the definition (190); scaling Q_n by $\Delta_{m,n}^{-1}$ changes only a constant, which cancels between the two times. □

Proposition 17.71 (Exact Fubini–Study audit of the rank-two pencil)**.** *Let $\mathcal{L}$ be the automorphy line in which the sections Ψ_x live, equipped with its Hodge metric, and put*

$$\mathbf{\Psi}_{x,y} = (\Psi_x, \Psi_y) \in \mathcal{L}^{\oplus 2}, \qquad A_{x,y} = \log\left|\frac{\Psi_y}{\Psi_x}\right|.$$

Give the direct sum the orthogonal metric. Away from the common zero divisor $D = 3P_7$, one has the exact soft-maximum identity

$$\log\|\mathbf{\Psi}_{x,y}\|^2 = \log\|\Psi_x\|^2 + \log\bigl(1 + e^{2A_{x,y}}\bigr). \tag{478}$$

If $c_1(\mathcal{L}, h) = -dd^c \log\|e\|^2$ in a local holomorphic frame and $\omega_{\mathrm{FS}} = dd^c \log(1 + |z|^2)$, then

$$dd^c \log\|\mathbf{\Psi}_{x,y}\|^2 = -c_1(\mathcal{L}, h) + \delta_D + \left(\frac{\Psi_y}{\Psi_x}\right)^* \omega_{\mathrm{FS}}. \tag{479}$$

Thus the pencil counts the Hodge divisor once and adds one positive Fubini–Study term.

This positivity does not, by itself, imply a stopped lower bound proportional to the controlled gain. Indeed, if $A(t_n) - A(t_m) = H$, the variation contributed by the last term in (478) is

$$\log(1 + e^{2A(t_n)}) - \log(1 + e^{2A(t_m)}). \tag{480}$$

For $A(t_m) = -2H$ and $A(t_n) = -H$, this tends to zero as $H \to \infty$, although the controlled gain is H. Consequently no positive constant times the gain follows from (479) without an additional estimate on the initial affine/unit shape.

The same conclusion holds for the singular toric, or sup, metric, but for a more precise reason than divisor duplication. In the local factorization used below, define

$$\|\mathbf{\Psi}_{x,y}\|_\infty = \max(\|\Psi_x\|, \|\Psi_y\|).$$

Then

$$\log \|\mathbf{\Psi}_{x,y}\|_\infty^2 = \log |a|^2 + \log \|e\|^2 + \max(\log |u_x|^2, \log |u_y|^2), \tag{481}$$

$$dd^c \log \|\mathbf{\Psi}_{x,y}\|_\infty^2 = -c_1(\mathcal{L}, h) + \delta_D + dd^c \max(\log |u_x|^2, \log |u_y|^2). \tag{482}$$

The last term is a positive current supported on $|u_x| = |u_y|$. In particular, the common divisor D is still counted only once. At a finite place the identical statement is the elementary factorization

$$\max(|\Delta u_x|_v, |\Delta u_y|_v) = |\Delta|_v \max(|u_x|_v, |u_y|_v),$$

so the common determinant ideal is also counted once.

Nevertheless the directed endpoint term relative to the first coordinate is only

$$\log \|\mathbf{\Psi}_{x,y}\|_\infty^2 - \log \|\Psi_x\|^2 = 2 \max(0, A_{x,y}). \tag{483}$$

For $A(t_m) = -2H$, $A(t_n) = -H$, both endpoint values in (483) vanish although the controlled gain is H. Thus the toric metric fixes the common-divisor ledger but not the directed affine baseline.

Proof. In a local frame write $\Psi_x = a u_x e$ and $\Psi_y = a u_y e$, where $a = 0$ cuts out the common divisor and u_x, u_y are nowhere zero. The norm identity is the elementary equality

$$\|\Psi_x\|^2 + \|\Psi_y\|^2 = \|\Psi_x\|^2 \left(1 + \left|\frac{\Psi_y}{\Psi_x}\right|^2\right).$$

Poincaré–Lelong applied to a, followed by $dd^c \log(|u_x|^2 + |u_y|^2)$, gives (479). Substitution of the two displayed endpoint values proves the final assertion. □

The audit distinguishes two notions of positivity. The Fubini–Study form measures motion of the projective ratio in the compact line $\mathbf{P}^1$; it cannot measure an arbitrarily long excursion toward zero or infinity with linear cost. The toric sup metric retains linear distance only on the side of the wall which it currently dominates. Equations (481) and (482) correct an earlier, overly pessimistic diagnosis: passing to the toric metric does *not* duplicate either common divisor. Its failure is directional. A gain which occurs entirely on the nondominant side is invisible, as (483) shows. The remaining theorem is therefore an orbit-specific bound on wall backtracking, or equivalently on the affine/unit baseline. It is not a missing static divisor identity.

A second reason scalar addition fails. The determinant ideal is shared by the two rational columns, and the order-seven divisor is shared by the two logarithmic-derivative sections. Thus adding separate scalar estimates duplicates both pieces of the Arakelov ledger. The quotient in (477) is harmonic; the positive curvature belongs to the rank-two pencil before the quotient is taken. The remaining fusion theorem must apply maximal slope to that pencil, not assign an independent copy of $3P_7$ to each of its two coordinates.

Proposition 17.72 (One outer rational horoball forces a cusp)**.** *Let $x = [P] \in \mathbf{P}^1(K)$, and in a finite affine chart put $z_t = x_0^\dagger + ie^{-t}$. If x is not a cusp and $t_n \to \infty$ are returns to a fixed compact part of the triangle surface, then for each $i = 1, 2$ there is a constant $c_{P,C} > 0$ such that*

$$\frac{|F_i(z_{t_n}) - x_i^\dagger|^2}{\operatorname{Im} F_i(z_{t_n})} \geq c_{P,C}. \tag{484}$$

Consequently, if for one outer embedding this quotient tends to zero along a compact-return subsequence, then x is a cusp. In particular, the conclusion holds if, for some i, the modular embedding has the rational Galois boundary value with finite nonzero angular derivative:

$$F_i(z_t) \longrightarrow x_i^\dagger, \qquad \frac{\operatorname{Im} F_i(z_t)}{e^{-t}} \longrightarrow \delta_i \in (0, \infty). \tag{485}$$

The same statements hold in the affine chart at infinity.

Proof. For a fixed real line ξ, the function

$$w \longmapsto \log \frac{|w - \xi|^2}{\operatorname{Im} w}$$

is its rank-two Hodge norm and is a hyperbolic Busemann function. It is therefore one-Lipschitz. Schwarz–Pick gives $d_{\mathbb{H}}(F_j(z_0), F_j(z_t)) \leq t + O(1)$, so, for either outer factor,

$$\log \|\sigma_j(P)\|_{j,t} \leq t + O_P(1). \tag{486}$$

At a compact return the nonzero integral Asai vector has norm bounded below, while its central factor has logarithmic norm $-t + O_P(1)$ by (131). Applying (486) to the other outer factor gives

$$\log \|\sigma_i(P)\|_{i,t_n} \geq -O_{P,C}(1).$$

In the chosen affine chart this is precisely (484). Thus entry into arbitrarily deep horoballs based at $x_i^\dagger$ contradicts integral-lattice discreteness.

Finally, Julia–Carathéodory applied to (485) gives

$$F_i(z_t) - x_i^\dagger = i\delta_i e^{-t} + o(e^{-t}),$$

so the quotient in (484) is $\delta_i e^{-t} + o(e^{-t})$. It tends to zero along every compact return. Hence a noncuspidal endpoint is impossible. □

Condition (485) is also the equality case of the pointwise Schwarz factor:

$$|\Psi_{i,x}(z_t)| \longrightarrow 1.$$

The real conclusion is stronger than that angular-derivative corollary. A noncuspidal counterexample must avoid, at every compact return, both outer horoballs based at the two rational conjugate

lines. We call this the *double-horoball mismatch*. It is consistent with a special period: under an outer embedding the rational eigenline can be the repelling, rather than the forward Hodge, line. The remaining equality theorem must show that persistent double-horoball mismatch either creates enough state height to be stopped by the Arakelov record, or coheres into exactly such a periodic unit eigenline.

The last sentence can be made quantitative. It reveals that the equality corridor is not an unspecified failure of contraction: it is motion in one explicit cone of the logarithmic unit lattice.

Proposition 17.73 (Unit-cone coordinates on the equality corridor)**.** *Put*

$$B_i(t) = \log \|\sigma_i(P)\|_{i,t} \qquad (i = 0, 1, 2).$$

Let t_n be compact returns, let g_n reduce them to one compact set, and choose ε_n so that $\varepsilon_n g_n P$ is unit-balanced as in Proposition 8.5. If $\mathcal{D}_P(t_n) = o(t_n)$, then

$$2\log|\sigma_0(\varepsilon_n)| = t_n + o(t_n), \tag{487}$$

$$2\log|\sigma_i(\varepsilon_n)| = -B_i(t_n) + o(t_n) \qquad (i = 1, 2). \tag{488}$$

Every accumulation point

$$\theta_i = \lim \frac{B_i(t_n)}{t_n}$$

therefore satisfies

$$0 \le \theta_1, \theta_2 \le 1, \qquad \theta_1 + \theta_2 = 1, \tag{489}$$

and the normalized logarithmic unit vectors converge along the same subsequence to

$$\frac{2}{t_n}\big(\log|\sigma_0\varepsilon_n|, \log|\sigma_1\varepsilon_n|, \log|\sigma_2\varepsilon_n|\big) \longrightarrow (1, -\theta_1, -\theta_2). \tag{490}$$

Proof. At a reduced point in a fixed compact set, the squared Hodge norm and the square of a fixed Euclidean norm are uniformly comparable. Unit balancing and $\mathcal{D}_P(t_n) = 6h([g_nP]) + O(1) = o(t_n)$ therefore give

$$2\log|\sigma_i(\varepsilon_n)| + B_i(t_n) = o(t_n) \qquad (i = 0, 1, 2).$$

The central identity $B_0(t_n) = -t_n + O_P(1)$ proves (487); the other two identities prove (488). Proposition 17.72 gives $B_i(t_n) \ge -O(1)$, while the individual Busemann bound (486) gives $B_i(t_n) \le t_n + O(1)$. Finally

$$B_1(t_n) + B_2(t_n) = t_n + o(t_n)$$

is the definition of the equality corridor. This proves (489); the product formula for ε_n proves the same sum relation in (490). □

Corollary 17.74 (The permitted special units lie in one cone)**.** *Suppose that a hyperbolic $h \in G_7$ fixes $[P]$, write $hP = \varepsilon P$, and orient h so that $[P]$ is its attracting fixed point in the distinguished embedding. Then*

$$|\sigma_0(\varepsilon)| > 1, \qquad |\sigma_1(\varepsilon)| < 1, \qquad |\sigma_2(\varepsilon)| < 1. \tag{491}$$

Along the corresponding closed geodesic the equality-corridor slope is

$$\theta_i = -\frac{\log|\sigma_i(\varepsilon)|}{\log|\sigma_0(\varepsilon)|} \qquad (i = 1, 2). \tag{492}$$

Proof. Attraction gives $|\sigma_0(\varepsilon)| > 1$. Reducing successive periods of the closed geodesic and using Hodge equivariance gives $B_i(nL) = -2n \log |\sigma_i(\varepsilon)| + O(1)$, where $L = 2 \log |\sigma_0(\varepsilon)|$. The outer lower bounds in Proposition 17.72 imply $|\sigma_i(\varepsilon)| \leq 1$. Equality would make a real conjugate of the algebraic unit equal to 1 or -1, hence would force $\varepsilon = \pm 1$, contrary to hyperbolicity. Both inequalities are strict. Dividing the displayed formula for $B_i(nL)$ by nL proves (492); the sum is one by the product formula. □

> **What this adds.** The Levi part of the equality case is now explicit. Balancing does not merely remove the units: it produces a sequence of actual unit-lattice points whose directions lie asymptotically on the compact segment $(1, -\theta, \theta - 1)$, $0 \leq \theta \leq 1$. A special closed geodesic is a rational ray in this lattice, with slope (492). Proposition 8.9 shows that this is necessary but not sufficient for closing: the transition equation also contains the additive Borel coordinate, while the rational-state coset can continue to drift. The missing theorem must make that coset recur; once it does, the permitted neutral Levi holonomy is exactly the rational unit ray above. This is more precise than asking for unspecified "control of units," and it uses the whole Galois orbit of the group.

Equivalently, Poincaré–Lelong and (462) give the two identities

$$dd^c\big(G_{1,x} + \log \|\Xi_1\|^2\big) = \delta_{2P_7}, \qquad dd^c\big(G_{2,x} + \log \|\Xi_2\|^2\big) = \delta_{P_7}.$$

Their sum is

$$dd^c\big(G_x + \log \|\Xi\|^2\big) = \delta_{3P_7}.$$

Thus the compact Hodge charge is globally the Green function of the full G_7-orbit of the order-seven point. The charged Rosen letters merely record crossings of a chosen tessellation near this divisor; they do not create the charge.

There is a crucial limitation. Inequality (475) is pointwise. It cannot be multiplied over different return times: by (474), the relative contribution between two returns is a difference of two values of $\log |\Psi_x|$, and hence telescopes. A trajectory may pass through the thick part repeatedly while the boundary value resets by motion in the rank-two unit frame. This is the exact reason that strict Schwarz–Pick contraction, even after the correct global transport has been constructed, is not yet the stopped height bound.

The remaining arithmetic statement is narrower: prove that these local extensions glue on the relevant toroidal or Deligne extension with an effective height jump at the first norm-record stopping regions, and convert that height jump into the stopped all-place degree inequality. Positive curvature and the local cusp expansion do not alone imply this global stopped boundary inequality.

Let p_1, p_2 be the source and target maps of a branch of the return groupoid. The formal Hodge-twisted line on that branch is

$$\mathscr{I}_g = \mathcal{O}(D_{c_g} - D_e) \otimes p_2^*\mathcal{N} \otimes p_1^*\mathcal{N}^{-1}. \tag{493}$$

The first factor has the evaluation (428) and zero interior curvature. The second factor has the right Hodge curvature as a difference of fibers, but it does not come with the required algebraic evaluation section. This is not a technical distinction. There are two possible interpretations, and neither supplies the missing map automatically.

If an arrow is regarded as a deck transformation, equivariance of the Hodge metric identifies the two fibers isometrically; that identification contains no accumulated Schwarz–Pick contraction. If the source and target are regarded as two points along the geodesic, the Gauss–Manin connection transports the ambient flat local systems, but it does not transport the Hodge sublines $\mathcal{L}_i$, and

therefore does not identify the two fibers of $\mathcal{N}$. Choosing an orthogonal projection would produce an analytic map, but its coefficient is a period and has no proved integral model. Merely tensoring the two fibers does not choose that map.

Consequently (493) specifies the source, target, divisor, and desired metric. At the level of the one-dimensional Hodge-graded line, the relative cocycle (438) still does not arise from an integral algebraic morphism. There are now two equivalent ways to retain the required information: construct a section

$$\tau_g \in p_2^*\mathcal{N} \otimes p_1^*\mathcal{N}^{-1}, \qquad \mathfrak{t}_g = s_g \otimes \tau_g \in \mathscr{I}_g, \tag{494}$$

in the associated graded, or keep a higher-rank flat tensor whose extremal Hodge weight is $\mathcal{N}$. Proposition 17.47 carries out the second construction: its content is $(r_g)^{-1}$, its norm is the pairwise energy, and its transport is the explicit scalar $r_{g'}/r_g$ in (403). Thus an unspecified integral projection is no longer part of the remaining problem.

What remains is stopped positivity for this explicit higher-rank transport. The numerator contracts in its distinguished Hodge norm, but the scalar $r_{g'}/r_g$ can transfer size to the other two embeddings. The first-record condition bounds its field norm without bounding its two unit-shape coordinates. A stopped arithmetic maximum principle must show that this compensation is bounded or periodic. Thus the obstruction is no longer the discovery or integral realization of a cocycle; it is the maximal-slope estimate for the concrete cocycle already displayed.

The notation in (494) must be interpreted carefully. There cannot be an independent flat transport in $\mathcal{N}$: $\mathcal{N}$ is an associated-graded Hodge line, and its deck-equivariant metric has zero period around a closed orbit, whereas the desired defect has a strictly positive period on every hyperbolic orbit. An integral realization must be extracted from a higher-rank flat object. The additive Hodge-graded model is

$$\mathbb{E}_\star = \operatorname{Sym}^2 \mathbb{V}_0 \oplus (\mathbb{V}_1 \otimes \mathbb{V}_2).$$

Its two highest Hodge lines are

$$A = \mathcal{L}_0^{\otimes 2}, \qquad B = \mathcal{L}_1 \otimes \mathcal{L}_2, \qquad \operatorname{Hom}(A, B) = \mathcal{N}.$$

Thus a lift of $\mathcal{A}$ compares two projections of *one and the same flat vector*, one to A and one to B, rather than choose a map between unrelated line fibers. Arithmetic integrality requires keeping the full Galois orbit of $\mathbb{E}_\star$, since a single real block has no preferred integral lattice. The multiplicative version is the representation $\mathbb{W}$ in (400); the normalized state $\mathbf{U}(X)$ performs the common-denominator division in (401). The positive conjugate degrees in (419) explain why this construction reduces the problem to maximal slope but does not settle that estimate by lattice discreteness alone.

There is one necessary refinement. For a general endpoint P, the line $K\ell_P$ is horizontal on the universal cover but does not descend as a section of the triangle quotient. The natural algebraic base for directions is the projective flat bundle

$$\mathcal{Y} = \mathbf{P}(\mathcal{V}^\vee) \longrightarrow C,$$

with its tautological covector line. Its tautological affine cone, with the lattice in (405), is the base for degrees. The endpoint determines a horizontal real leaf in $\mathcal{Y}$; the parabolic invariant e determines a boundary incidence divisor; and pairing the universal tautological line with $\delta(e)$ gives the candidate framed normal function. A rigorous biextension proof must first construct this framed object over the affine cone and only then restrict its direction to the horizontal leaf. This

prevents both invalid simplifications: treating the endpoint covector as a global algebraic section of C, and treating the record ideal as a function of its projective direction.

The expected cusp/special alternatives have a simple lattice explanation. For an integral local system with an adapted metric, a nonzero integral vector in a strictly decreasing boundary subspace cannot follow a nonparabolic recurrent geodesic: at compact returns its norm would tend to zero, contradicting discreteness of the integral lattice. If the vector is only nonincreasing, its reduced transports lie in a finite set; two returns then agree. Provided the evaluated vector retains the boundary state, the intervening group element fixes the endpoint. It is parabolic, giving a cusp, or hyperbolic; a hyperbolic element fixing a point of $\mathbf{P}^1(K)$ has its eigenvalues in K.

This argument explains exactly why the positive Asai weight $\chi = -a_0 + a_1 + a_2$ is the obstruction. The desired Hodge–incidence twist must put the record section, including the finite ideal $(r_j)^{-1}$, into the nonpositive subspace. Once that is proved, lattice discreteness supplies the two conclusions; without it, a spectral gap for the operator does not control the evaluated record line.

There is a complementary representation-theoretic formulation. If ρ_i are the three real rank-two monodromy representations, consider

$$\rho_\star = \operatorname{Sym}^2 \rho_0 \oplus (\rho_1 \otimes \rho_2). \tag{495}$$

Writing $a_i(g)$ for the top singular exponent of $\rho_i(g)$, the gap between the two blocks is

$$\Delta_\star(g) = 2a_0(g) - a_1(g) - a_2(g).$$

For a parabolic shear T^n, all three exponents are $\log|n| + O_K(1)$, so $\Delta_\star(T^n) = O_K(1)$. The character is therefore neutral on long cusp motion, while the two Schwarz–Pick defects make it positive across compact returns after endpoint angles are controlled. This is the pattern of a relatively dominated, or cusped-Anosov, cocycle; Filip's work gives the relevant Hodge-theoretic model for this kind of statement [20]. A sufficient theorem would have to dominate the *evaluated logarithmic conormal line*, with transversality measured by (278) and equality only on parabolic axes or unit eigenlines. A singular-value gap for $\rho_\star$ alone would repeat the failed operator-norm argument and is not sufficient.

The same relative incidence data has a weight-three realization through tensor induction. On the rank-two Asai secant determined by P and c_j, the Cayley hyperdeterminant is

$$\prod_{i=0}^{2} \det(P_i, c_{j,i})^2 = |\,\mathrm{N}\, r_j|^2.$$

Along a variation of the i-th secant factor, its logarithmic conormal has coefficient $1/\sigma_i(r_j)$; the distinguished component is therefore the reciprocal ideal $(r_j)^{-1}$ found in Proposition 13.2. This identifies the pairwise tensor canonically as a logarithmic conormal component of the Asai discriminant. The remaining issue is still transport of the secant endpoint through the relative cocycle κ.

Precise unproved step. Problem 17.1 is the broad stopped-transport problem. Problem 17.15 is its active rank-two formulation, because the Wronskian supplies both the incidence ideal and the rank-drop locus. The finite ideal and pointwise real norm are proved, but their contraction along the changing prefix lines is not. Maximal slope synchronizes the finite places only after the new record is created; it does not repair the transport or index mismatch. The natural-extension calculation above sharpens the same point: using the correct orbit path restores the uncontrolled affine endpoint coordinate.

18 The fixed-point form and the nearest-wall method

There is a second group-wide way to see the same obstruction. It was developed in an earlier draft in terms of symmetric powers and a "nearest-wall" Hodge pencil. The useful part of that construction is not a calculation with the Rosen generators. It attaches one integral quadratic form to an arbitrary element of the full group.

For

$$g = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathcal{O}_K) \quad \text{and} \quad v = (p, q)^t \in \mathcal{O}_K^2$$

put

$$Q_g(v) = \det(v, gv) = cp^2 + (d - a)pq - bq^2. \tag{496}$$

Then

$$Q_g(v) = 0 \quad \Longleftrightarrow \quad g[v] = [v]. \tag{497}$$

If the left side is nonzero, it is an algebraic integer and hence

$$|N_{K/\mathbf{Q}} Q_g(v)| \geq 1. \tag{498}$$

There is an exact way to fuse two such forms. It turns the formerly unspecified global Bézout coefficient into the trace of one commutator.

Proposition 18.1 (Fixed-point resultant equals the commutator divisor)**.** *For $g, h \in \mathrm{SL}_2(\mathcal{O}_K)$, write*

$$Q_g(X, Y) = aX^2 + bXY + cY^2, \qquad Q_h(X, Y) = AX^2 + BXY + CY^2.$$

Then

$$\mathrm{Res}(Q_g, Q_h) = (aC - cA)^2 - (aB - bA)(bC - cB) = \mathrm{tr}([g, h]) - 2. \tag{499}$$

In particular, the resultant vanishes exactly when g and h have a common fixed point over $\overline{K}$.

Let $P \in \mathcal{O}_K^2$ be primitive and put $d_g = Q_g(P)$, $d_h = Q_h(P)$. If the resultant is nonzero, then at every finite prime $\mathfrak{p}$,

$$\min\{v_\mathfrak{p}(d_g), v_\mathfrak{p}(d_h)\} \leq v_\mathfrak{p}(\mathrm{tr}([g, h]) - 2). \tag{500}$$

Thus the common ideal of two fixed-point incidences is controlled as one global divisor by the commutator divisor.

Proof. For two binary quadratics the first equality in (499) is the Sylvester determinant, written in terms of the three 2-by-2 coefficient minors. If

$$g = \begin{pmatrix} u & v \\ w & z \end{pmatrix},$$

then its coefficient triple is $(w, z - u, -v)$. Substitution for g and h, followed by the two determinant-one identities, gives

$$\mathrm{Res}(Q_g, Q_h) = \mathrm{tr}(ghg^{-1}h^{-1}) - 2.$$

Both sides vanish precisely when the two quadratic fixed-point equations share a root.

For the valuation statement, complete P to a local unimodular basis and use it as the variables (X, Y), with $P = (1, 0)$. The leading coefficients of the two transformed forms are then d_g, d_h. If both are divisible by $\mathfrak{p}^e$, the two minors $aC - cA$ and $aB - bA$ are divisible by $\mathfrak{p}^e$. The displayed resultant is therefore divisible by $\mathfrak{p}^e$, which proves (500). □

The same identity separates spectral growth from geometric alignment. This is the form in which the real structure of the triangle group enters.

Corollary 18.2 (The commutator divisor is normalized axis separation)**.** *Assume that neither g nor h is parabolic, and let*

$$\alpha_g, \beta_g, \alpha_h, \beta_h \in \mathbf{P}^1(\overline{K})$$

be their two fixed points. Put

$$t(g,h) = \frac{(\alpha_g - \alpha_h)(\beta_g - \beta_h)}{(\alpha_g - \beta_h)(\beta_g - \alpha_h)}, \qquad D_g = \operatorname{tr}(g)^2 - 4, \quad D_h = \operatorname{tr}(h)^2 - 4.$$

Changing the order of either pair of fixed points replaces t by t^{-1}, and therefore leaves $t/(1-t)^2$ unchanged. One has the identity in K

$$\frac{\operatorname{tr}([g,h]) - 2}{D_g D_h} = \frac{t(g,h)}{(1 - t(g,h))^2}. \tag{501}$$

At a real embedding where both elements are hyperbolic, let ℓ_g, ℓ_h be their translation lengths and let θ be the complex angle between their axes: θ is the ordinary angle when the axes meet and is i times their distance when they are disjoint. Then

$$\operatorname{tr}([g,h]) - 2 = -4 \sinh^2(\ell_g/2) \sinh^2(\ell_h/2) \sin^2 \theta. \tag{502}$$

Thus the commutator is small only when spectral growth is paid for by alignment of the two fixed-point axes.

Proof. Write

$$Q_g(X,1) = a(X - \alpha_g)(X - \beta_g), \qquad Q_h(X,1) = A(X - \alpha_h)(X - \beta_h).$$

The resultant is a^2A^2 times the product of the four differences between a root of the first polynomial and a root of the second. The two discriminants are

$$D_g = a^2(\alpha_g - \beta_g)^2, \qquad D_h = A^2(\alpha_h - \beta_h)^2.$$

Dividing and using

$$\frac{(\alpha_g - \alpha_h)(\alpha_g - \beta_h)(\beta_g - \alpha_h)(\beta_g - \beta_h)}{(\alpha_g - \beta_g)^2(\alpha_h - \beta_h)^2} = \frac{t}{(1-t)^2}$$

proves (501). In the uniformizing real plane the last quotient is $-\sin^2\theta/4$, while $D_g = 4\sinh^2(\ell_g/2)$ and similarly for h. This gives (502); imaginary angle gives the disjoint-axis case with the correct positive sign. □

The outer embeddings of this divisor can in fact be synchronized by a theorem about the whole Hecke group. This is substantially stronger than estimating the three matrix products separately.

Proposition 18.3 (Trace dominance reduces the resultant height to one real axis)**.** *Let $M \in G_7$ and write $u = \operatorname{tr}(M)$. For the two nonuniformizing embeddings one has*

$$|\sigma_i(u)| \le \max\{2, |\sigma_0(u)|\}, \qquad i = 1, 2. \tag{503}$$

Consequently, if $g, h \in G_7$ and

$$R(g,h) = \operatorname{tr}([g,h]) - 2 \neq 0,$$

then

$$|\sigma_i(R(g,h))| \le |\sigma_0(R(g,h))| + 4, \tag{504}$$

$$\log|\,\mathrm{N}\,R(g,h)| \le 3\log\bigl(|\sigma_0(R(g,h))| + 4\bigr). \tag{505}$$

Combining this with (502) *gives*

$$\log|\,\mathrm{N}\,R(g,h)| \le 3\log\bigl(4 + 4\sinh^2(\ell_g/2)\sinh^2(\ell_h/2)|\sin\theta|^2\bigr)\,. \tag{506}$$

Thus the stopped resultant problem requires only a bound for the uniformizing commutator angle; no separate outer-axis estimate is needed.

Proof. For a hyperbolic element of the full Hecke reflection group, the distinguished absolute trace dominates every Galois-conjugate absolute trace. This is the trace-dominance theorem used by Bugeaud–Hubert–Schmidt in their height estimate for Rosen convergents [3, Lemma 3.1 and the preceding trace theorem]. It applies in particular to G_7. If M is not hyperbolic, discreteness of the uniformizing triangle group makes it elliptic or parabolic. Elliptic elements are conjugate to powers of the order-two or order-seven generators, and parabolic elements have trace ± 2; in either case all conjugate traces have absolute value at most two. This proves (503).

Apply the result to $M = [g,h]$. Since $|\sigma_0(u)| \le |\sigma_0(u-2)| + 2$, the triangle inequality gives (504). Multiplication over the three real embeddings gives (505), and Corollary 18.2 gives the last display. □

Trace dominance has a second consequence which is more directly adapted to the fixed-point equation. It is obtained by varying an arbitrary group element along the *entire parabolic stabilizer of an arbitrary group cusp.* Thus it is a statement about the full thin group, not a calculation with one preferred Rosen generator.

For a column $v = (p,q)^t$, put

$$N_v = \begin{pmatrix} -pq & p^2 \\ -q^2 & pq \end{pmatrix}. \tag{507}$$

Then $N_v^2 = 0$. If $v = A(1,0)^t$ is a cusp column with $A \in G_7$, its primitive parabolic is

$$U_v = ATA^{-1} = I + \lambda N_v.$$

Proposition 18.4 (Parabolic-ray domination of every cusp evaluation)**.** *Let $M \in G_7$, and let $v \in \mathcal{O}_K^2$ be the first column of an element of G_7. Then*

$$|\sigma_i(\lambda Q_M(v))| \le |\sigma_0(\lambda Q_M(v))|\,, \qquad i = 1,2. \tag{508}$$

In particular, write

$$Q_M(X,Y) = cX^2 + eXY - bY^2, \qquad e = d - a,$$

and set

$$H_0(Q_M) = \max\{|\sigma_0(c)|, |\sigma_0(e)|, |\sigma_0(b)|\}.$$

There are constants C_i, depending only on the three embeddings of K, such that

$$\max\{|\sigma_i(c)|, |\sigma_i(e)|, |\sigma_i(b)|\} \le C_i H_0(Q_M), \qquad i = 1,2. \tag{509}$$

Thus the complete outer coefficient height of the fixed-point form is bounded by its coefficient height in the uniformizing plane.

Proof. The rank-one identity

$$\operatorname{tr}(MN_v) = \det(v, Mv) = Q_M(v)$$

gives, for every integer n,

$$\operatorname{tr}(MU_v^n) = \operatorname{tr}(M) + n\lambda Q_M(v). \tag{510}$$

If $Q_M(v) = 0$, equation (508) is immediate. If it is nonzero, the distinguished trace in (510) has absolute value greater than two for all sufficiently large $|n|$. Apply the Hecke-group trace-dominance theorem to MU_v^n, divide by $|n|$, and let $|n|$ tend to infinity. This proves (508).

The three columns

$$(1,0)^t, \qquad (0,1)^t, \qquad (\lambda,1)^t$$

are cusp columns: they are the first columns of I, S, TS. Evaluation at the first two gives c and $-b$. Evaluation at the third gives

$$Q_M(\lambda,1) = \lambda^2 c + \lambda e - b.$$

Applying (508) to these three values and solving the last identity for e proves (509). All denominators which occur are the fixed nonzero numbers $\sigma_i(\lambda)$, so the constants depend only on the field. □

Corollary 18.5 (A one-place exactification threshold)**.** *For every primitive $P = (p,q)^t \in \mathcal{O}_K^2$, there is a constant $C(P) > 0$ such that every non-scalar $M \in G_7$ with $Q_M(P) \neq 0$ satisfies*

$$1 \le |\operatorname{N} Q_M(P)| \le C(P)\,|\sigma_0(Q_M(P))|\,H_0(Q_M)^2. \tag{511}$$

Consequently, if a sequence of non-scalar elements $M_n \in G_7$ satisfies

$$|\sigma_0(Q_{M_n}(P))| H_0(Q_{M_n})^2 \longrightarrow 0, \tag{512}$$

then $Q_{M_n}(P) = 0$ for all sufficiently large n, and the desired stabilizer has been found.

More generally, for the Hecke group G_q with totally real trace field of degree d, the same argument gives

$$1 \le |\operatorname{N} Q_M(P)| \le C_q(P)|\sigma_0(Q_M(P))|H_0(Q_M)^{d-1}. \tag{513}$$

Thus the exponent is 2 for $q = 7$ and 4 for $q = 11$. The constant can be made completely explicit. Put

$$\lambda_i = \sigma_i(\lambda), \qquad A_0 = \lambda_0^2 + |\lambda_0| + 1, \qquad r_i = \frac{|\lambda_0|}{|\lambda_i|},$$

$$C_i = \max\left\{ r_i,\ \frac{r_i}{|\lambda_i|}\left(A_0 + |\lambda_i|^2 + 1\right) \right\}, \qquad L_i(P) = |\sigma_i(p)|^2 + |\sigma_i(p)\sigma_i(q)| + |\sigma_i(q)|^2.$$

One may take

$$C(P) = C_1 C_2 L_1(P) L_2(P). \tag{514}$$

In particular, the strict finite criterion is

$$|\sigma_0(Q_M(P))| H_0(Q_M)^2 < C(P)^{-1}.$$

Proof. At an outer embedding, Proposition 18.4 gives

$$|\sigma_i(Q_M(P))| \le C_i H_0(Q_M) L_i(P).$$

Multiply the two outer inequalities by the distinguished absolute value. The left side is the absolute field norm of the nonzero algebraic integer $Q_M(P)$, and is therefore at least one. This proves (511);

its stated consequence is immediate. The proof for a trace field of degree d is identical, with one factor of $H_0(Q_M)$ for each of the $d-1$ outer embeddings.

For completeness, cusp evaluation at $(1,0)^t$ and $(0,1)^t$ gives $|\sigma_i(c)|, |\sigma_i(b)| \le r_i H_0(Q_M)$. If $w = Q_M(\lambda, 1)$, then

$$|\sigma_i(w)| \le r_i |\sigma_0(w)| \le r_i A_0 H_0(Q_M).$$

The identity

$$\sigma_i(e) = \frac{\sigma_i(w) - \lambda_i^2 \sigma_i(c) + \sigma_i(b)}{\lambda_i}$$

gives the displayed value of C_i. Expanding $Q_M(P)$ term by term then gives $C_i H_0(Q_M) L_i(P)$, proving (514). □

18.1 A thin-group stammering theorem

Let G_7^{ref} denote the full Hecke reflection group, in which G_7 has index two. The coefficient-dominance proposition and the exactification corollary remain valid for $M \in G_7^{\mathrm{ref}}$. Indeed, $MU_v^n \in G_7^{\mathrm{ref}}$, and the trace-dominance theorem used in their proofs is stated for this full reflection group. If such an M fixes a real point and reverses orientation, then $M^2 \in G_7$ fixes the same point. For an attracting periodic branch $M^2 \neq 1$.

We use the standard Rosen convergent matrices

$$C_n = \begin{pmatrix} p_{n-1} & p_n \\ q_{n-1} & q_n \end{pmatrix}, \qquad \xi = C_n \cdot T^n \xi,$$

with $q_n > 0$. If U, V are finite digit words, write $u = |U|$, $v = |V|$, and UV^∞ for the eventually periodic word. The notation UV^w, for real $w > 1$, includes the appropriate initial part of one last copy of V, as in the usual stammering condition.

Theorem 18.6 (Thin-group exactification of a repeated block)**.** *Let $\xi \in K \cap [-\lambda/2, \lambda/2)$ have an infinite Rosen expansion, and write $\xi = p/q$ with $P = (p,q)^t \in \mathcal{O}_K^2$ primitive. Suppose there are admissible words U_ν, V_ν, with $V_\nu \neq \varnothing$, and real numbers $w_\nu > 1$, such that $U_\nu V_\nu^{w_\nu}$ is a prefix of the expansion and $U_\nu V_\nu^\infty$ is admissible. Put*

$$u_\nu = |U_\nu|, \qquad v_\nu = |V_\nu|, \qquad N_\nu = |U_\nu V_\nu^{w_\nu}|.$$

If $N_\nu \to \infty$ and

$$\frac{(q_{u_\nu} q_{u_\nu + v_\nu})^3}{q_{N_\nu}^2} \longrightarrow 0, \tag{515}$$

then ξ is a cusp or is fixed by a hyperbolic element of G_7 whose eigenline is defined over K.

Proof. Let α_ν be the real number with Rosen word $U_\nu V_\nu^\infty$. The two numbers ξ and α_ν have their first N_ν digits in common, so the standard Rosen cylinder estimate gives

$$|\xi - \alpha_\nu| \le c_1 q_{N_\nu}^{-2}. \tag{516}$$

The return matrix

$$M_\nu = C_{u_\nu + v_\nu} C_{u_\nu}^{-1} \in G_7^{\mathrm{ref}}$$

fixes α_ν. Each entry of C_n is bounded at the distinguished real place by a fixed multiple of q_n, while C_n^{-1} has the same entries up to signs because $|\det C_n| = 1$. Consequently

$$H_0(Q_{M_\nu}) \le C_\lambda q_{u_\nu} q_{u_\nu + v_\nu}. \tag{517}$$

This is the central-place, coefficient-height part of the usual height estimate for a periodic Rosen surrogate [3, Lemma 3.4]; no field norm has yet been taken.

Write $M_\nu = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$. If β_ν is its other fixed point and $c \neq 0$, then

$$Q_{M_\nu}(z,1) = c(z-\alpha_\nu)(z-\beta_\nu).$$

Because both ξ and α_ν lie in the fixed Rosen interval, the identity

$$c(\xi-\beta_\nu) = c(\xi+\alpha_\nu) + (d-a)$$

shows that

$$|Q_{M_\nu}(\xi,1)| \le C_\xi H_0(Q_{M_\nu})|\xi-\alpha_\nu|.$$

The same estimate is immediate from the linear fixed-point equation when $c=0$. Since $P=(p,q)^t$ is fixed, homogeneity and (516) give

$$|\sigma_0 Q_{M_\nu}(P)| \le C(P) H_0(Q_{M_\nu}) q_{N_\nu}^{-2}. \tag{518}$$

Suppose $Q_{M_\nu}(P) \neq 0$. The one-place exactification inequality, valid in G_7^{ref} by the observation preceding the theorem, now yields

$$1 \le |\,\mathrm{N}\, Q_{M_\nu}(P)| \le C'(P)\frac{H_0(Q_{M_\nu})^3}{q_{N_\nu}^2} \le C''(P)\frac{(q_{u_\nu}q_{u_\nu+v_\nu})^3}{q_{N_\nu}^2}.$$

This contradicts (515) for all large ν. Hence some nontrivial return matrix M_ν fixes ξ. If it preserves orientation it is parabolic or hyperbolic. If it reverses orientation, its square lies in G_7, fixes ξ, and is nonidentity because a periodic inverse branch is attracting, whereas a reflection has multiplier of absolute value one. The parabolic case is a cusp and the hyperbolic case is the special alternative. □

Corollary 18.7 (A readable 3/2 criterion)**.** *Assume*

$$1 < b := \liminf q_n^{1/n} \le B := \limsup q_n^{1/n} < \infty.$$

Under the hypotheses of Theorem 18.6, its conclusion holds whenever

$$\limsup_{\nu\to\infty} \frac{N_\nu}{2u_\nu+v_\nu} > \frac{3}{2}\frac{\log B}{\log b}. \tag{519}$$

Proof. Pass to a subsequence on which the inequality has a fixed positive margin. For every small $\eta>0$, eventually

$$q_n \ge (b-\eta)^n, \qquad q_n \le (B+\eta)^n.$$

The logarithm of the quotient in (515) is therefore at most

$$3(2u_\nu+v_\nu)\log(B+\eta) - 2N_\nu\log(b-\eta),$$

which tends to $-\infty$ after η is chosen sufficiently small. Apply the theorem. □

What has improved, and what has not. Bugeaud–Hubert–Schmidt bound the ordinary integer height of the periodic surrogate by $(q_u q_{u+v})^{[K:\mathbf{Q}]}$, then use a degree-two algebraic-approximation inequality with exponent 3. For $K=\mathbf{Q}(2\cos(\pi/7))$, their stammering threshold is therefore $9/2$, and their conclusion *quadratic over K or transcendental* says nothing about

a point already in K. The argument above keeps the fixed-point equation over K, uses trace dominance to control its two outer coefficient vectors, and applies the product formula directly to its value at the given K-rational point. This removes the extra factor 3, giving $3/2$. It proves the classification in the dominant-repeated-block regime. It still does not treat the balanced many-block regime exhibited by $|a_j| = \lceil e^{\sqrt{j}} \rceil$, where no single repeated block is large enough.

For $q = 11$ the identical proof has four outer embeddings and gives the sufficient denominator condition

$$\frac{(q_u q_{u+v})^5}{q_N^2} \longrightarrow 0,$$

or the asymptotic threshold $\frac{5}{2}(\log B/\log b)$. This partial theorem is independent of the Hodge-defect sign. The $q = 11$ comparison therefore separates two mechanisms cleanly: trace dominance controls a sufficiently dominant literal repetition for every q, with a penalty equal to the trace field degree; only the Hodge–Arakelov argument could control the nonrepetitive accumulation of many cusp excursions, and its scalar sign changes at $q = 11$.

This result isolates exactly what an "optimal Roth estimate" would and would not do here. At $q = 7$, a bound

$$|\sigma_0(Q_M(P))| \ll H_0(Q_M)^{-2}$$

with an uncontrolled implied constant is critical: it need not contradict (511). One needs a constant below the explicit product-formula threshold, or a little-oh improvement as in (512). For $q = 11$, the same whole-group mechanism asks for four powers because there are four outer embeddings. This is a clean degree-dependent obstruction, separate from the scalar Hodge-sign calculation.

There is also a moving-frame version. Since $A^{-1}MA \in G_7$, Proposition 18.4 applies after every change of cusp frame $A \in G_7$. If $X = AP$ and $B \in G_7$, then

$$|\operatorname{N} Q_B(X)| \le C_7 |\sigma_0(Q_B(X))| H_0(Q_B)^2 \prod_{i=1}^{2} \big(|\sigma_i(X_1)| + |\sigma_i(X_2)|\big)^2, \tag{520}$$

where C_7 depends only on the field. Thus the dependence on the moving frame occurs only through the two outer sizes of X, not through outer coefficients of Q_B. For two reduced states $X_m = g_m P$, $X_n = g_n P$, take $B = g_n g_m^{-1}$; then

$$Q_B(X_m) = \det(X_m, X_n) = \Delta_{m,n}.$$

Thus trace dominance removes the outer fixed-form matrices from the stopped rank-two problem. It does not remove the outer height of the transported rational state. Controlling precisely that remaining factor is the pairwise Arakelov stopping problem of Section 12.

Proposition 18.8 (The target-dependent intermediate group)**.** *Let $P = (p, q)^t \in \mathcal{O}_K^2$ be primitive, put*

$$\mathcal{U}_P = \{n_P(t) = I + tN_P : t \in K\},$$

and choose a nonzero $t \in \mathcal{O}_K$ for which $n_P(t)$ is integral. For every cusp column v, one has

$$\operatorname{tr}\big(n_P(t)U_v^{-1}\big) - 2 = t\lambda \det(P, v)^2. \tag{521}$$

Moreover:

1. *$G_7 \cap \mathcal{U}_P \neq \{1\}$ if and only if $[P]$ is a cusp;*
2. *a nonparabolic element of G_7 normalizes $\mathcal{U}_P$ if and only if it fixes $[P]$, in which case it is one of the special hyperbolic elements required by the classification;*

3. *if $n_P(t) \notin G_7$, the distinguished real projection of*

$$\Gamma(P,t) = \langle G_7, n_P(t)\rangle$$

is dense in $\mathrm{PSL}_2(\mathbb{R})$*, although its diagonal embedding in* $\mathrm{SL}_2(\mathbb{R})^3$ *remains discrete.*

Proof. The rank-one identity

$$-\operatorname{tr}(N_P N_v) = \det(P,v)^2$$

and $U_v^{-1} = I - \lambda N_v$ give (521). A nonidentity element of $\mathcal{U}_P$ is parabolic with unique fixed line KP. Its membership in G_7 is therefore equivalent to $[P]$ being a group cusp. The normalizer of a one-dimensional unipotent group fixes its unique boundary line. Conversely, if $hP = \varepsilon P$, then

$$h n_P(s) h^{-1} = n_P(\varepsilon^2 s),$$

so h normalizes $\mathcal{U}_P$. A nonparabolic element of the Fuchsian group fixing a real boundary point is hyperbolic; because its eigenline is K-rational, its eigenvalue lies in K.

Finally, $\Gamma(P,t)$ is an integral intermediate group containing G_7. If its distinguished projection were discrete, it would be a discrete overgroup of the finite-covolume triangle group. Proposition 7.5 says that no proper such overgroup exists. Hence the projection is dense. Discreteness of the full diagonal embedding follows from the Minkowski discreteness of $\mathrm{SL}_2(\mathcal{O}_K)$. □

This proposition explains both the power and the exact boundary of the parabolic-ray argument. If w, v are two group cusps, then $U_w U_v^{-1} \in G_7$, so trace dominance controls $\lambda^2 \det(w,v)^2$. Replacing w by the rational target P produces the same trace formula (521), but the mixed product now belongs only to the dense target-dependent group $\Gamma(P,t)$. It is therefore invalid simply to apply the Hecke-group trace theorem to it.

The desired relative theorem can now be described without invoking an artificial fixed intermediate group. Along the cusp approximants selected by the ray to P, the failure of the mixed traces to satisfy the G_7 dominance rate must either be absorbed by the $1/5$ Hodge gap or become coherent under a normalizer of $\mathcal{U}_P$. The first alternative makes the integral right side of (521) tend to zero at all places and forces a cusp. The second is exactly the special hyperbolic alternative. Proving this rate dichotomy is another, group-theoretic formulation of the stopped rank-two Arakelov inequality; the elementary trace identity does not prove the dichotomy.

This factorization makes the relation with the $p = 11$ comparison precise. Hodge theory bounds the spectral factors D_g, D_h at all real embeddings. The thin-orbit input must bound the normalized axis separation $t/(1-t)^2$ for *consecutive stopped blocks*. Proposition 18.3 shows that, once the two factors are recombined into the commutator, this is a one-real-place problem. At $p = 7$, the scalar Hodge defect leaves a margin of $1/5$ if that alignment follows the geodesic-selected rate. At $p = 11$, the same scalar budget has the opposite sign, so an argument which ignores the axis-separation factor cannot possibly give a uniform prime-family theorem. Trace dominance therefore identifies a genuinely different uniform possibility: a commutator/resultant proof could overcome the scalar sign without estimating four outer axes separately. Its price is the sharp central two-sided alignment estimate in (506). The negative degrees of the abstract Hodge lines do not supply that estimate by themselves.

For the raw Rosen recurrence the commutator divisor has a much simpler form. This is the finite-place fusion statement which was missing from the earlier prime-support discussion.

Proposition 18.9 (Exact fusion of two consecutive Rosen incidences)**.** *Put*

$$B_a = \begin{pmatrix} 0 & -1 \\ 1 & -a\lambda \end{pmatrix}, \qquad v_{j+1} = B_{a_j} v_j, \qquad D_j = \det(v_j, v_{j+1}) = Q_{B_{a_j}}(v_j),$$

where $v_j \in \mathcal{O}_K^2$ is primitive. Pull the next form back to the same frame by

$$\widetilde{Q}_{j+1} = Q_{B_{a_j}^{-1} B_{a_{j+1}} B_{a_j}}.$$

Then

$$\widetilde{Q}_{j+1}(v_j) = D_{j+1}, \tag{522}$$

$$D_{j+1} - D_j = (a_{j+1} - a_j)\lambda\,(v_j)_2(v_{j+1})_2, \tag{523}$$

$$\mathrm{Res}(Q_{B_{a_j}}, \widetilde{Q}_{j+1}) = \lambda^2(a_{j+1} - a_j)^2. \tag{524}$$

Moreover, for every finite prime,

$$\min\{v_{\mathfrak{p}}(D_j), v_{\mathfrak{p}}(D_{j+1})\} \le v_{\mathfrak{p}}(a_{j+1} - a_j). \tag{525}$$

Equivalently, apart from the one principal divisor of the digit difference, the two consecutive incidence divisors are coprime.

Proof. Writing $v_j = (X, Y)^t$, direct conjugation gives

$$Q_{B_a}(X, Y) = X^2 - a\lambda XY + Y^2,$$
$$Q_{B_a^{-1} B_b B_a}(X, Y) = X^2 - (2a - b)\lambda XY + \{1 + a(a - b)\lambda^2\}Y^2.$$

Their difference is $(b - a)\lambda Y(X - a\lambda Y)$, proving the first two identities, and their Sylvester determinant is $\lambda^2(a - b)^2$.

If a prime divides D_j, it divides neither $Y = (v_j)_2$ nor $X - a_j\lambda Y = (v_{j+1})_2$: modulo these two coordinates,

$$D_j \equiv X^2 \pmod{Y}, \qquad D_j \equiv Y^2 \pmod{X - a_j\lambda Y},$$

and both adjacent state columns are primitive. Apply valuations to (523) to obtain (525). □

What this advances. The finite overlap of two incidence divisors is no longer an unknown collection of local primes. For arbitrary blocks it is bounded by the single algebraic integer $\mathrm{tr}([g, h]) - 2$; for adjacent Rosen steps it is bounded by the ordinary digit difference. Thus a global version of "new divisor" bookkeeping can be written as one Arakelov gcd degree, with the cusp digit paying the entire overlap.

What remains. The Hodge contractions used in the proof occur on long stopped blocks, not on isolated Rosen steps. For two such blocks the digit difference in (525) is replaced by the commutator divisor in (500). Controlling the height of that commutator along consecutive thin-orbit blocks is therefore the precise extra estimate needed to turn the proved local fusion into the stopped rank-two inequality. Proposition 18.3 reduces its all-place height to the one uniformizing commutator angle, but the exact record-block audit in `explore.tex` shows that this central quantity need not be bounded on finite closing prefixes. A static wedge does not provide this control: the $p = 11$ test below shows that $Q_{g_n} \wedge Q_{g_{n+1}}$ normally has the sum of the two coefficient growth rates rather than a new negative Hodge rate.

Independently of this two-form fusion, a sufficiently strong global upper bound for the three conjugates of $Q_g(v)$ proves exact closing, not merely approximate closing. This is the same product-formula mechanism as the transverse minor $\Delta_{m,n}$. In fact Proposition 8.15 proves the exact identity

$$\Delta_{m,n} doteq Q_{g_m^{-1} g_n}(P),$$

where $\doteq$ means equality up to a unit. Thus the two formulations have the same closing event: equality of two reduced rational states says that their interval quotient fixes the original line.

The direct estimate has a familiar limitation. At the distinguished embedding a long prefix has a periodic fixed point exponentially close to the endpoint of its cylinder, so the normalized value of $\sigma_0(Q_g(v))$ is small. Either outer embedding may compensate for that smallness. The integral Asai obstruction in Section 8 is exactly the statement that such compensation is unavoidable for a nonclosing rational ray. The nearest-wall construction tries to use the worse of the two outer embeddings, rather than their sum, and to amplify the remaining strict Hodge gap by taking symmetric powers.

Before discussing that arithmetic amplification, one can settle a loose geometric point from the earlier draft. Changing to an eigenline frame does have a cost, but the cost is elementary and is not independent of the eigenvalue gap. The next proposition records the exact calculation.

Let V be an oriented Euclidean plane. Give $\mathrm{Sym}^m(V^\vee)$ its Bombieri norm; in an orthonormal basis X, Y, this means

$$\|X^{m-j}Y^j\|^2 = \binom{m}{j}^{-1}.$$

For $d \geq 2r$, set $n = d - 2r$ and $k = n + 1$. Multiplication by $(XY)^r$ is the map

$$\mu_{d,r} : \mathrm{Sym}^n(V^\vee) \longrightarrow \mathrm{Sym}^d(V^\vee).$$

Its singular values are

$$s_j = \left(\frac{\binom{n}{j}}{\binom{d}{r+j}} \right)^{1/2}, \qquad 0 \leq j \leq n, \tag{526}$$

and therefore

$$\mathcal{E}(d,r) := -\log|\det \mu_{d,r}| = \frac{1}{2}\sum_{j=0}^{n} \left(\log\binom{d}{r+j} - \log\binom{n}{j} \right). \tag{527}$$

This is a universal metric term. It depends on d, r, but not on a return word.

Proposition 18.10 (Balanced eigenframes have a controlled determinant cost)**.** *Let $B \in \mathrm{SL}(V)$ have singular values $e^{\kappa(B)}, e^{-\kappa(B)}$. Use the standard action ρ_m of $\mathrm{SL}(V)$ on $\mathrm{Sym}^m(V^\vee)$, and put*

$$q_B = \rho_2(B)(XY), \qquad \mu_B(f) = q_B^r f.$$

Then, for the k-dimensional Jacobian computed with the Bombieri volumes,

$$|\log \mathrm{Jac}_k(\mu_B) - \log \mathrm{Jac}_k(\mu_{d,r})| \leq kd\,\kappa(B). \tag{528}$$

If $A \in \mathrm{SL}_2(\mathbb{R})$ is hyperbolic, has translation length ℓ, and B is a determinant-one eigenframe for A, then

$$Q_A = -2\sinh(\ell/2)\, q_B. \tag{529}$$

Consequently the possible ill-conditioning of the two eigenlines and the smallness of their eigenvalue separation occur in one formula; they may not be charged as two unrelated errors.

Proof. With the convention just chosen,

$$\mu_B = \rho_d(B)\,\mu_{d,r}\,\rho_n(B)^{-1}.$$

The determinant of $\rho_n(B)$ is $(\det B)^{n(n+1)/2} = 1$. Every singular value of $\rho_d(B)$ lies between $e^{-d\kappa(B)}$ and $e^{d\kappa(B)}$. Its volume distortion on any k-plane therefore lies between $e^{-kd\kappa(B)}$ and $e^{kd\kappa(B)}$, proving (528).

Choose eigenvectors e_+, e_- of determinant one, with $Ae_+ = e^{\ell/2}e_+$ and $Ae_- = e^{-\ell/2}e_-$. If $v = xe_+ + ye_-$, then

$$\det(v, Av) = (e^{-\ell/2} - e^{\ell/2})xy = -2\sinh(\ell/2)xy.$$

This is (529). □

There is an equally elementary hyperbolic interpretation. If the base Hodge metric is represented by $z \in \mathbb{H}$, the axis of A is at distance R from z, and $D = d_{\mathbb{H}}(z, Az)$, then

$$\sinh(D/2) = \cosh R \, \sinh(\ell/2). \tag{530}$$

A determinant-one eigenframe can be balanced by rescaling its two columns reciprocally; in that normalization its least possible Cartan size is $R/2$. Hence the right side of (528), together with the scalar in (529), is controlled by the endpoint displacement D. In particular, a small translation length does not create a separate uncontrolled eigenframe error: the factor $\sinh(\ell/2)$ becomes small at exactly the same time.

For a return segment let T be its length in the distinguished hyperbolic plane, let L_i be the length of its image under the i-th outer period map, let D_i be the distance between the two outer endpoints, and let ℓ_i be the translation length of the outer monodromy. Then

$$T - \ell_i = (T - L_i) + (L_i - D_i) + (D_i - \ell_i). \tag{531}$$

The first term is strict Schwarz–Pick contraction. The second is backtracking of the outer period path. The third is the axis/eigenframe term. Proposition 18.10 and (530) control the third term, including the case of nearly colliding eigenlines, once the unnormalized integral form Q_g is retained. Backtracking is favorable for an endpoint metric, since $D_i \le L_i$; it is not an additional loss which has to be paid by a new arithmetic estimate.

What this actually settles. The moving-eigenframe issue is a rank-two volume calculation. It is not the deep missing theorem. The calculation above supplies the axis-distance part of the old nearest-wall proposal and explains why its apparent singularity near the identity was an artifact of normalizing the two eigenline factors while discarding $2\sinh(\ell_i/2)$.

What it does not settle. The nearest-wall pencil changes the real weight from the sum of both outer expansions to the largest one by tensoring with a quotient of Hodge lines. That quotient has the desired negative curvature, but it is not an independent flat integral line. Treating it as one would incorrectly turn the proved real contraction into an arithmetic contraction of $Q_g(P)$. An integral realization must keep the affine framing of the rational Borel state. When that framing is written out, its missing coordinate is exactly s_n, and its transport is exactly

$$s_{n+1} = d_n^2 s_n + d_n \beta_n.$$

Thus the symmetric-power route and the Wronskian route meet at the same place: one must control an additive extension of the neutral unit line, not improve the estimate for an eigenbasis.

This gives a concrete synthesis. The fixed-point form (496) supplies the integral equation whose vanishing gives a stabilizer. The fundamental Hodge line $\mathcal{N}_2$, whose Higgs section has divisor P_7, supplies the strict nearest-wall curvature. The transverse Wronskian supplies an integral first

jet whose rank drops exactly when the rational Borel frame repeats. The affine cocycle records the extension data needed to put those three objects in one lattice. A proof of Problem 17.15 may therefore be viewed as the Arakelov realization of the nearest-wall twist: its finite part is the minor ideal, its curved part is P_7, and its boundary value is the affine endpoint cost. Once this realization is constructed, the product formula gives exact vanishing and Proposition 17.28 finishes the classification.

19 Numerical audit

The numerical work has one legitimate purpose here: falsification. It can reject an identity, a monotonicity assertion, or a proposed uniform bound by producing one exact counterexample. It cannot show that the changing prefix tensor is transported by the Hodge connection, because that is a statement about which geometric line is being followed, not about the values of a scalar on finitely many states. Earlier drafts gave the calculations too much evidentiary weight; the flat-line/record-line distinction was invisible to them.

Field elements are represented in the basis $1, \lambda, \lambda^2 - 1$ with rational coefficients. Addition, multiplication, inversion, field norm, the recurrence, and state repetition are exact. The Rosen digit is also certified exactly: a rational isolating interval for the distinguished root of (2) is refined until the relevant field element lies inside one nearest-integer cell. Floating-point logarithms are taken only after the algebraic element, its digit, and its record status have been fixed exactly.

The first grid calculation enumerated 36111 field states and found each terminating or exactly periodic within the chosen horizon. This is a check on the implementation and a source of examples. It supplies no evidence for the missing incidence transport, since every finite orbit segment necessarily has some finite maximum.

The program `rosen_cusp_transfer_tests.py` now records

$$\mathscr{E}(X) - 6h(P)$$

at every transition which creates a new norm record.

An exhaustive search with coefficient bound 5, rational denominators at most 20, and 120 steps produced 9862 such transitions. It found

$$\max\bigl(\mathscr{E}(X) - 6h(P)\bigr) = 10.7369152971\ldots.$$

The corresponding maximum for the smaller stopped envelope was $10.8197793730\ldots$. On every tested row,

$$3h(X) + \log(1 + \Theta) - \mathscr{E}(X) \le 0.640720 < \log 2,$$

which independently checks Lemma 12.2.

The pairwise energy is not monotone. On the balanced inverse-unit family, its increase from the initial state reaches

$$57.1719937053\ldots$$

at exponent 48. Nevertheless its excess over the initial sextic budget there is

$$-92.7052826767\ldots.$$

This separates the proposed bound from a false Lyapunov claim: initial arithmetic height pays for a large unit realignment, while a new norm record still remains within the global budget.

Several earlier scalar reductions fail on comparably explicit tests.

proposed simplification	diagnostic outcome
all-state monotonicity	the balanced inverse-unit family makes the candidate energy grow linearly over long initial segments;
state–prefix comparison	a weak-approximation family makes the signed comparison defect arbitrarily negative even when the second transition is a strict first record;
positivity of outer incidence	an exhaustive denominator-20 grid gives minimum $-8.886961\ldots$; only a bounded-below statement can survive;
operator norm of the cotangent	the full operator proxy reaches $180.0538\ldots$ while the evaluated stopped scalar remains within its budget;
separate Hodge drop and cusp charge	record samples show that the two terms can each be large with opposite signs and must be kept in one degree;
pairwise stopped degree	is not refuted on the tested families, but these values do not test whether the pointwise tensor is a flat or contracting section;

sample	largest increase of $\mathscr{E}$	largest $\mathscr{E} - 6h(P)$
500 random inputs	16.500605	negative in that run
coefficient 3, denominator 10	10.864047	5.835796
coefficient 5, denominator 20	16.825241	10.736915
balanced inverse, exponent 48	57.171994	-92.705283

The computations reject monotonicity and verify the elementary scalar identities, including Lemma 12.5. They do not validate the proposed Hodge mechanism. In particular, evaluating $\mathscr{E}$ at many states cannot detect the false inference from "these are the same pointwise weights" to "this is one transported contracting section." That inference is ruled out theoretically by Lemma 14.2, regardless of how many finite samples obey the stopped inequality.

The distinction among the samples matters. Random points explore typical short orbits. Coefficient/denominator grids expose low-height resonances. Unit rays test archimedean imbalance at fixed finite support. Weak-approximation families prescribe three different real directions and then approximate them by one exact element of K. These structured families remain useful for destroying overstrong scalar claims. The pairwise stopping inequality is retained only as a precisely stated sufficient conjecture, not because its survival in those samples supplies a reason to believe the missing transport theorem.

20 What has and has not been achieved

The coherent picture is now the following. The Hodge calculation on the actual geodesic is complete, and the dynamic-to-arithmetic comparison is an exact cross ratio, not an unspecified error. The analytic endpoint problem has also been reduced to a proved two-channel dichotomy. One rank-two Arakelov fusion theorem for consecutive quotients in a pointed thin orbit remains. It must select the larger channel without counting their common determinant divisor twice. The affine equality calculation then distinguishes the hyperbolic and parabolic alternatives.

1. *Proved reduction.* The classification is the rational-Borel intersection statement of Proposition 7.1. If a rational endpoint is not a cusp, its integral Asai vector forces outer Hodge growth

at least 1 along compact returns. The ordinary outer growth is $1/5 + 3/5 = 4/5$. Hence every hypothetical counterexample lies in a definite $1/5$-exceptional Lyapunov stratum. More precisely, Proposition 8.2 proves that the excess over central contraction is $6h([gP]) + O(1)$. Proposition 8.3 then closes the orbit whenever this excess is bounded along a subsequence, while Proposition 8.4 proves the sharp elementary lower barrier $\max_{n\le N} \mathcal{D}_P(t_n) \ge \log N - O(1)$ for any nonperiodic sequence of N compact returns.

2. *Exact endpoint calculation and missing theorem.* The dynamic future–cusp frame and the rational two-future Wronskian frame differ by $-A_{m,n}/\Delta_{m,n}$. Proposition 8.16 proves that the entire endpoint loss is the explicit quantity

$$2\log|\,\mathrm{CR}_0(m,n)| - \log|\,\mathrm{CR}_1(m,n)| - \log|\,\mathrm{CR}_2(m,n)|,$$

with algebraic endpoint lines $[E]$ and $[k_{m,n}P]$. Problem 17.16 records one sufficient scalar non-absorption estimate. The stronger structural advance is Proposition 8.20: if X_0 is the central cross-ratio gain, then

$$\mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_n) - \mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_m) \ge \mathfrak{C}(m,n) - 2X_0,$$

and hence the maximum of these two channels is at least $\mathfrak{C}(m,n)/3$. Moreover $X_0 = G_{m,n}(t_m) - G_{m,n}(t_n)$, so the first channel is literally a drop of the central two-frame Green function. Outer backtracking cannot erase both quantities. The remaining statement is Problem 17.17: retain the larger quantity in one stopped Galois-completed metric, with $(\Delta_{m,n})$ entered once. Proposition 17.18 constructs that integral extension explicitly, and Proposition 17.19 writes every one of its splitting lines as $Q_m \otimes (Q_n/\Delta_{m,n} + sQ_m)$, $s \in K$. Its real norms are the three distances from $\sigma_i(s)$ to moving Picard–Fuchs period ratios, and its finite content is $(s, \Delta_{m,n}^{-1})$. Corollary 17.24 proves that the dynamic splitting uses the entire determinant denominator. The subsequent audit is essential: Proposition 17.20 computes its ordinary degree and proves that it sees only the symmetric Green sum, so an ordinary Harder–Narasimhan argument cannot yield the cyclic gain. Proposition 17.23 constructs two explicit tensor lines whose norm ratio is exactly the controlled channel and whose commonly rescaled contents both have norm $|\,\mathrm{N}\,\Delta|^{-1}$. Their static global degree difference is zero. Thus the unproved part is the stopped Bott–Chern/relative holonomic Roth inequality for this explicit pair, not construction of an induced partition or an integral bundle. Proposition 17.26 proves that this stopped fusion gives the full classification.

3. *Relation to existing Hodge rigidity.* Filip's theorem proves that an integral vector in a strictly decreasing log-Anosov boundary subbundle comes from a cusp. Proposition 8.40 proves why it cannot be applied directly: rational descent restores the two outer Levi weights, and special unit eigenlines are neutral rather than decreasing. Conjecture 17.52 is the required equality-case extension for the explicit normalized Asai state.

4. *Affine rather than projective record data.* Proposition 17.48 proves that the normalized Asai direction forgets $(r_j)^{-1}$; the record is the lattice framing of its affine cone. Proposition 17.51 independently proves that an ordinary all-place two-form Roth product is identically at the critical power 2. Thus neither a theorem on the projective flat bundle nor a sharper standard version of Roth can provide the missing strict inequality. The required theorem is a stopped slope estimate for the framed affine tensor.

5. *One Hodge generator and its full arithmetic shape.* Proposition 17.30 and Corollary 17.31 prove $\mathcal{N}_1 \simeq \mathcal{N}_2^2$ and $\mathcal{N} \simeq \mathcal{N}_2^3$. Thus all Hodge curvature comes from the one fundamental

divisor P_7. Formula (442) separates the remaining transport into three copies of that curved coordinate and the flat arithmetic cube coordinate $\nu(m) = \mathrm{N}(m)/\tau^2(m)^3$. Fixed-support S-unit theory controls the latter. The new arithmetic synthesis is Proposition 17.59: the coordinate $\alpha_2(m) = \tau^2(m)/m$ paired with $\mathcal{N}_2$ is already a degree-three isogeny of the entire shape torus, and $\nu = \alpha_2^{-2}\tau(\alpha_2)^{-1}$. Thus the flat character is not a second independent arithmetic variable. Corollary 17.62 packages the norm and the full α_2-height into one degree-nine proper height, and Theorem 17.64 proves that boundedness of this height along infinitely many record states already closes the orbit. Establishing that stopped boundedness is the precise arithmetic component still required.

6. *What the geometric calculation now proves.* The cusp-normalized dynamic frame follows the original geodesic. Proposition 8.12 integrates the strict Schwarz defect on that path. The transverse Wronskian has finite divisor $(\Delta_{m,n})$, and Corollary 8.14 compares it with the dynamic frame by one rational shear. Hence neither the Hodge line, its first jet, its finite denominator, nor the relevant path remains to be guessed.

7. *Equality stage of the same attack.* Proposition 17.72 proves that a counterexample avoids both rational outer horoballs at compact returns. In the sublinear equality corridor, Proposition 17.73 turns the two remaining Hodge growth rates into an actual sequence of logarithmic units with normalized direction $(1, -\theta_1, -\theta_2)$, $\theta_1 + \theta_2 = 1$. Corollary 17.74 proves that every permitted special orbit gives a rational ray in exactly this cone. Proposition 8.8 proves that the infimal all-place height of the Borel coset is comparable with the Weil height of the complete state. Proposition 8.9 identifies the additional datum which cannot be omitted: the rational-state coset $C_n\mathscr{B}_P$. The gauge-dependent unipotent entry in a lifted transition shows that the Levi coordinate does not determine this coset, so unit locking alone does not repeat it. The active missing step is to show that a geodesic prefix sequence from the full thin group must either acquire positive transverse height or repeat the frame coset. The repeated frame gives the rational-Borel intersection; the unit cone then describes its neutral holonomy. Proposition 8.30 shows why the existing S-unit and moving-target theorems do not supply this recurrence: the former sees only the Levi entry, while the latter starts with a point sequence already contained in its moving hyperplane. Propositions 8.31 and 8.34 further prove that more coordinates or weighted copies of the same Asai state do not repair the defect. Proposition 17.4 now constructs the required transverse first jet: its Galois-cyclic ratio is the arithmetic shape of $\Delta_{m,n}$ multiplied by the Higgs section with divisor $3P_7$, and its rank drops exactly at frame recurrence or on that Hodge divisor. The viable synthesis is therefore the stopped integral slope estimate (333) for this explicit Wronskian. Proposition 8.12 proves the strict Schwarz bound directly on the original future–cusp geodesic, so there is no remaining path-transport problem. Corollary 8.14 gives the exact rational shear to the Wronskian frame, and Proposition 8.16 writes its full endpoint cost as one Galois-weighted cross ratio. Proposition 17.28 proves that this one estimate gives the full classification. Proposition 8.20 now proves the required positive fraction analytically, but in the maximum of two coordinates rather than in one signed scalar. The missing input is the rank-two stopped fusion inequality (335) for consecutive quotients in the pointed thin orbit. Arithmetic holonomy is a possible method only if it constructs this one framed rank-two lattice and does not count its determinant divisor twice. Abel's identity in Proposition 17.5 proves that the determinant lattice has no word-dependent denominator beyond $(\Delta_{m,n})$. The remaining verification is precisely (335); the scalar bound (334) remains an alternative sufficient route. Section 18 also reconciles this route with the symmetric-power proposal of the

earlier draft. The eigenframe contribution is now controlled by the exact Bombieri Jacobian estimate (528); the apparent singularity of nearly coincident eigenlines cancels against the factor $2\sinh(\ell/2)$ in the unnormalized integral fixed-point form. Hence the unresolved issue is not diagonalization. It is the integral realization of the curved nearest-wall Hodge quotient, which is the same affine-framing problem measured by (159).

Proposition A.6 prevents an important overclaim in the archived finite-quotient discussion: the pointed square-root estimate is not an off-the-shelf equidistribution conjecture. Under a trivial stabilizer its left side is identically zero for large R. A proof of that estimate would therefore already have to contain the frame-recurrence rigidity in inverse form. It supplies no independent route and is not part of the active proof strategy.

The longer ledger below records the exact Hodge and arithmetic reductions which lead to this picture and the alternative formulations tested along the way.

1. The global cusp-transfer inequality reduces the norm of a new remainder to one explicitly defined scalar at the preceding state.
2. That scalar is dominated, with loss at most $\log 2$, by the pairwise energy $\mathscr{E}$, whose initial value is at most $6h(P)$.
3. The energy is the degree of an explicit global metrized line; its finite fractional ideal is exactly the reciprocal remainder ideal.
4. The evaluated Hodge tensor has the same real weights and the same finite ideal at each fixed state. This is a pointwise identification, not a transport or contraction theorem.
5. McMullen's strict Schwarz–Pick contraction applies to the flat endpoint covector. Proposition 14.1 proves that this flat line is different from the moving line which contains the remainder divisor.
6. Lemma 14.2 computes the exact cost of changing between the two lines, and Proposition 14.3 identifies its product over the three embeddings with the desired norm.
7. Once the norms are bounded, the Galois four-term S-unit equation and the two-step glue give termination or repetition.
8. The order-seven monodromy splits the Hodge defect into the two divisors $2P_7$ and P_7. Equivalently, the three full triangle-group representations have relative Euler numbers $5/14, 1/14, 3/14$, whose defect $2e_0 - e_1 - e_2 = 3/7$ is the degree of $3P_7$. The scalar determinant loses one of the two unit directions.
9. The cyclic completion is still sharper: $\mathcal{Q}_0^2, \mathcal{Q}_1^2, \mathcal{Q}_2$ are $\mathcal{O}(-3P_7), \mathcal{O}(3P_7), \mathcal{O}$. Thus there is one curved elliptic direction and one flat arithmetic direction at the level of real Hodge characters. Proposition 17.59 proves that after Galois completion they belong to the one proper K-valued coordinate α_2: the flat character is a monomial in its conjugates. Proposition 17.32 identifies the full elliptic orbit with the CM norm $u^2 - \lambda uv + v^2$ and factors its rational incidence as cusp denominator times complete remainder state. This is a presentation-free bridge from the Hodge divisor to the S-unit problem. Proposition 17.35 strengthens this bridge: outside the fixed bad set, every prime in either CM factor is split and its divisibility is exactly one of the two incidences $g\zeta = \infty, x$ in the complete reduced projective orbit. Corollary 17.36 identifies both incidences with matrix coefficients in the single Steinberg constituent of the full finite action. Thus the unresolved non-recycling assertion is a coherent pointed-orbit inverse theorem for the whole thin group, not a condition on Rosen generators. Proposition 17.33

gives a concrete description of the word "thin": the ambient arithmetic group parametrizes all oriented $\mathcal{O}_K$-bases of $\mathcal{O}_L$, while G_7 selects one suborbit modulo common order-seven rotation. Proposition 17.40 makes the characteristic-zero state equally explicit: $u_g = c_\eta(g^{-1}x)$ is the complete target state, $v_g = c_\eta(g^{-1}\infty)$ is the cusp state, and $t_g = v_g/u_g$. Their heights are additive up to a fixed error. Consequently the exact whole-orbit closing target is $h_L(t_g) - h_L(v_g) = O_P(1)$ along infinitely many projectively distinct stopped returns. The deduction of a stabilizer from this bound is formal; proving the bound is the substantive relative-domination problem for the thin group.

10. The projectivized triple $\mathcal{Z}_0 \oplus \mathcal{Z}_1 \oplus \mathcal{Z}_2$ retains both integral Asai shape coordinates. Proposition 17.56 proves that bounded radial norm and bounded rank-two shape give finitely many complete states directly.

11. Proposition 7.3 proves that the two outer real images are jointly dense and remain so inside every finite-index congruence kernel. Thus a finite class and arbitrary open outer conditions can always be met simultaneously. The missing content is quantitative: the lift must have distinguished displacement at the information-theoretic scale while remaining in the Hodge-typical Cartan window.

12. Proposition A.19 proves that one entropy-optimal, Hodge-balanced lift of a fixed rational residue flag is an actual cusp. At the expected exponents the norm comparison has the strict margin $(1/10)$. The entropy-level local-limit version is a correct conditional implication, but (606) shows that its positivity already contains exact vanishing; prime averaging does not make it a standard equidistribution problem.

13. The fixed-flag Livšic route is circular: the rational flag descends only to the stabilizer cover, whose closed geodesics already are the desired stabilizers. The noncircular substitute is the tower of finite pointed congruence covers, whose inverse-limit stabilizer is $\bigcap_{\mathfrak{a}} \Gamma_x(\mathfrak{a})$.

14. The two-prime amplification replaces one level of size nearly e^R by two split-prime levels below $e^{R/2}$. Proposition A.23 shows that a Hodge-weighted second factorial-moment estimate for these two actions implies the full classification. This is a concrete bilinear large-sieve target for the whole thin group, not an estimate for individual letters.

15. A fixed incidence fibre is the intersection of the complete parabolic conjugacy class with an algebraic unipotent-conjugacy orbit. The connected algebraic equality subgroups are exactly the two expected unipotents. A midpoint point–line argument gives the genuine whole-block exponent $2/3$. A Hodge tilt does not improve this into rigidity: its larger fibre is exactly the conditional mean in the smaller incidence box. This failure explains why the successful congruence formulation must fix the rational target before counting.

16. Independently of stopping times, Proposition 8.1 proves that every noncuspidal trace-field endpoint has outer Hodge growth at least one along compact returns. Since the generic outer sum is $4/5$, the remaining counterexample set is confined to an explicitly exceptional Lyapunov stratum. Theorem 8.37 turns rigidity of that stratum into the full classification.

17. Thin-orbit local–global dispersion supplies a concrete whole-group architecture below the separate closing scales. The exact block identity (607) converts the fixed-point defect into a bilinear scalar coefficient. The congruence and real losses depend only on $r + \beta$, leaving the exact window $9/10 < r + \beta < 1$. The new input still required is a pointed exceptional-set theorem, with the two projective Hodge cocycles and a rational Borel as the only persistent resonance.

The rank-two stopped assertion is the active geometric formulation. More precisely, Proposition 8.20 proves that one third of every Hodge defect survives in either the Green or the controlled-Wronskian coordinate, and Problem 17.17 asks for their one-divisor Arakelov fusion. Proposition 17.34 realizes the Green coordinate as proximity to the identity in one norm-one CM torus. The pointed dispersion assertion is retained only as a record of an abandoned approach. Its missing resonance-to-Borel implication is not weaker than the theorem being sought.

For comparison, the geometric formulation asks first that (279) must bound the radial norm for the prefix cusps selected by first norm records. Second, the maximal slope of the projectivized three-component Asai lattice must bound both ratios in (422). The finite-state recurrence and the S-unit theorem imply neither statement. The scalar Hodge determinant does not imply the second without an additional rigidity statement, because its unit character is not proper. The two individual Green–Higgs factorizations provide the required archimedean curvatures, and the Asai components provide the matching integral lattices. The two-channel lemma now performs the analytic selection; the missing proof is its stopped global maximal-slope realization with the determinant lattice counted once. Equality must mean either a vanishing incidence or a rational unit eigenline. Those are precisely the cusp and special alternatives. In the equality corridor, Proposition 17.73 identifies the neutral Levi coordinate, while Proposition 8.9 identifies the complete frame which must recur. The stopped maximal-slope comparison must either force positive transverse height or repeat that frame coset. The latter gives the rational Borel intersection, hence the special alternative. This is the single path carried forward by the manuscript.

A Archived finite-quotient experiments

A.1 Discarded finite-quotient route: a presentation-free Fourier problem

This subsection and the subsequent finite-quotient variants are preserved to explain why they were abandoned. They are not used in the proof spine. All of them eventually require a pointed inverse theorem saying that a long rational cylinder with anomalous concentration comes from a rational Borel stabilizer. That is the theorem already isolated by the Asai and unit-cone argument, with additional finite-level machinery but no new source of rigidity. The reader interested only in the active route may stop before this appendix: the active argument and its exact remaining step have already been presented in Sections 8 through the concluding summary.

The projective reduction used above forgets the additive structure of the scalar defect. Retaining that scalar for the *full group shell*, including the right parabolic coordinate, replaces the family of separate congruence tests by one Fourier problem on the Minkowski torus of K.

Let $\widetilde{G}_7 \subset \mathrm{SL}_2(\mathcal{O}_K)$ be the standard lift. For a primitive $P \in \mathcal{O}_K^2$, put

$$D_g(P) = \det(P, gP) \in \mathcal{O}_K.$$

Proposition A.1 (Abundance of the full Hodge-typical shell)**.** *For every $\delta > 0$, there are annular sets $\widetilde{\Omega}_R(\delta) \subset \widetilde{G}_7$, with distinguished hyperbolic displacement in $[R, R+1]$, such that*

$$\#\widetilde{\Omega}_R(\delta) = M_R = e^{R-o(R)} \tag{532}$$

and

$$\log\|\sigma_0(g)\|_{\mathrm{op}} \le R/2 + O(1), \quad \log\|\sigma_1(g)\|_{\mathrm{op}} \le (1/10+\delta)R, \quad \log\|\sigma_2(g)\|_{\mathrm{op}} \le (3/10+\delta)R. \tag{533}$$

Proof. The number of elements of the cofinite Fuchsian group with displacement in $[R, R+1]$ is $e^{R-o(R)}$. Orbital equidistribution identifies the distribution of the two outer matrix cocycles

on these annuli with their Bowen–Margulis distribution. Oseledets' theorem applies because the logarithmic cocycles are integrable at the cusp. Their top exponents, in the normalization where the uniformizing exponent is one, are $1/5$ and $3/5$. Since the uniformizing operator norm grows as $e^{R/2+O(1)}$, the two outer operator norms grow typically as $e^{R/10+o(R)}$ and $e^{3R/10+o(R)}$. Removing the exceptional set leaves $e^{R-o(R)}$ elements and proves (533). The exponent values are the Hodge-degree calculation in (12). □

Fix $0 < \delta < 1/20$ and take the shell in Proposition A.1. The determinant estimate at the three embeddings gives, for g in this shell,

$$\begin{aligned} |\sigma_0 D_g(P)| &\ll_P e^{R/2}, \\ |\sigma_1 D_g(P)| &\ll_P e^{(1/10+\delta)R}, \\ |\sigma_2 D_g(P)| &\ll_P e^{(3/10+\delta)R}. \end{aligned} \tag{534}$$

These are the three column-growth exponents of the Hodge-balanced orbit. Their sum is $9/10 + 2\delta$. Thus $D_P(\widetilde{\Omega}_R(\delta))$ lies in a Minkowski box containing $e^{(9/10+2\delta)R+o(R)}$ algebraic integers, while the orbit supplies $e^{R-o(R)}$ group elements.

Let $K_\infty = K \otimes_{\mathbb{Q}} \mathbb{R}$, let $\mathfrak{d}_K^{-1}$ be the codifferent, and normalize Haar measure on

$$\mathbb{T}_K = K_\infty/\mathfrak{d}_K^{-1}$$

to have mass one. Define the complete-orbit exponential sum

$$S_R(P;\xi) = \sum_{g \in \widetilde{\Omega}_R(\delta)} e^{2\pi i \operatorname{Tr}_{K/\mathbb{Q}}(\xi D_g(P))}. \tag{535}$$

Proposition A.2 (Exact Fourier inversion and the Hodge major arc)**.** *One has*

$$\#\{g \in \widetilde{\Omega}_R(\delta) : D_g(P) = 0\} = \int_{\mathbb{T}_K} S_R(P;\xi)\, d\xi. \tag{536}$$

There are constants $c_P, c_P' > 0$ such that the box

$$\mathfrak{M}_R = \left\{ \xi : \begin{array}{l} |\sigma_0\xi| \le c_P e^{-R/2}, \\ |\sigma_1\xi| \le c_P e^{-(1/10+\delta)R}, \\ |\sigma_2\xi| \le c_P e^{-(3/10+\delta)R}, \end{array} \right\} \tag{537}$$

in a fixed fundamental domain satisfies

$$\operatorname{Re} \int_{\mathfrak{M}_R} S_R(P;\xi)\, d\xi \ge c_P' M_R e^{-(9/10+2\delta)R}. \tag{538}$$

Consequently, if $G_7 \cap B_P(K) = \{\pm 1\}$, the deliberately stronger minor-arc estimate

$$\int_{\mathbb{T}_K \setminus \mathfrak{M}_R} |S_R(P;\xi)|\, d\xi < c_P' M_R e^{-(9/10+2\delta)R} \tag{539}$$

for one sufficiently large R gives a contradiction. Thus (539) is a sufficient diagnostic, not the estimate pursued below. The rational arcs have necessary signed mass, so an absolute-value bound discards the local arithmetic.

Proof. The character group of $\mathbb{T}_K$ is $\mathcal{O}_K$. Character orthogonality therefore gives

$$\int_{\mathbb{T}_K} e^{2\pi i \operatorname{Tr}_{K/\mathbb{Q}}(\xi r)}\, d\xi = \mathbf{1}_{\{r=0\}},$$

which proves (536). By (534), c_P can be chosen so small that the real part of every summand in (535) is at least $3/4$ throughout $\mathfrak{M}_R$. The volume of that box is $\asymp_P e^{-(9/10+2\delta)R}$, proving (538). Estimate (539) then makes the right side of (536) positive. The shell is annular, so it does not contain $\pm I$. Thus it contains a projectively nonidentity g with $g[P] = [P]$, contradicting the assumed triviality of the Borel stabilizer. □

The last sentence is important. The zero arc is positive, but it is not the only structured arc. Rational torsion points contribute with alternating signs. The correct proof must evaluate their signed sum; it should not try to make their absolute contribution smaller than the zero arc.

The finite local transform in this formulation is unusually favorable and uses the complete thin group rather than a generating alphabet.

Lemma A.3 (Exact additive flattening of the complete group)**.** *Let $\mathfrak{p}$ be a good prime at which the standard lift of G_7 surjects onto $\mathrm{SL}_2(k_\mathfrak{p})$, and put $q = \#k_\mathfrak{p}$. If ψ is a nontrivial additive character of $k_\mathfrak{p}$ and $\bar{P} \neq 0$, then*

$$\frac{1}{|\mathrm{SL}_2(k_\mathfrak{p})|} \sum_{g \in \mathrm{SL}_2(k_\mathfrak{p})} \psi\big(\det(\bar{P}, g\bar{P})\big) = -\frac{1}{q^2-1}. \tag{540}$$

Equivalently,

$$\Pr(\det(\bar{P}, g\bar{P}) = 0) = \frac{1}{q+1}, \qquad \Pr(\det(\bar{P}, g\bar{P}) = a) = \frac{q}{q^2-1} \quad (a \neq 0). \tag{541}$$

More generally, over $\mathcal{O}_K/\mathfrak{p}^m$, the Fourier coefficient attached to a character of exact conductor $\mathfrak{p}^j$ is

$$\widehat{\nu}_{\mathfrak{p}^m}(\psi) = \begin{cases} 1, & j = 0, \\ -(q^2-1)^{-1}, & j = 1, \\ 0, & 2 \le j \le m, \end{cases} \tag{542}$$

and

$$\Pr\big(D_g(P) = 0 \pmod{\mathfrak{p}^m}\big) = \frac{q^{1-m}}{q+1}. \tag{543}$$

Proof. The action of $\mathrm{SL}_2(k_\mathfrak{p})$ on nonzero vectors is transitive, and every vector has an order-q stabilizer. Hence $g\bar{P}$ is uniform on $k_\mathfrak{p}^2 \setminus \{0\}$. The map $v \mapsto \det(\bar{P}, v)$ is a nonzero linear functional. Its additive character sum over all q^2 vectors is zero; deleting the zero vector leaves -1, which proves (540). Its zero fibre has $q-1$ nonzero vectors, and every nonzero fibre has q, proving (541).

For the prime-power assertion, a primitive vector modulo $\mathfrak{p}^m$ is a vector not divisible by $\mathfrak{p}$, so there are $q^{2m} - q^{2m-2}$ of them. The character sum over all vectors is zero. The sum over the deleted vectors divisible by $\mathfrak{p}$ is q^{2m-2} when the character has conductor $\mathfrak{p}$, and zero when its conductor has exponent at least two. This proves (542). In the zero fibre the free coordinate must be a unit, leaving $q^m - q^{m-1}$ vectors; division by $q^{2m} - q^{2m-2}$ gives (543). □

The coefficient $(q^2-1)^{-1}$ is substantially smaller than the projective Steinberg scale q^{-1}. There is no free power saving: Fourier inversion of the zero fibre sums q additive characters and recovers the projective density $(q+1)^{-1}$. The gain is structural. All finite places and all three archimedean coordinates can now be kept inside one global Fourier transform.

The global modulus and the signed target. Fix

$$9/10 + 2\delta < u < 1 \tag{544}$$

and form the single squarefree ideal

$$\mathfrak{A}_R = \prod_{\substack{\mathfrak{p}\ mgood \\ \mathrm{N}\,\mathfrak{p} \le uR}} \mathfrak{p}. \tag{545}$$

Strong approximation supplies a fixed finite set of excluded places for which the reduction map $widetildeG_7 \to \mathrm{SL}_2(mathcalO_K/\mathfrak{A}_R)$ is surjective for every R. Thus "good" in (545) is a simultaneous adelic condition, not merely a separate assertion at each factor. The prime-ideal theorem gives

$$\log \mathrm{N}\,\mathfrak{A}_R = uR + o(R). \tag{546}$$

This is an Arakelov construction: the ideal is chosen once, by its total degree, and is compared with the norm of the global scalar $D_g(P)$. There is no selection depending on the factorization of an individual defect.

For every $g \in \widetilde{\Omega}_R(\delta)$, (534) gives

$$|\,\mathrm{N}\,D_g(P)| \ll_P e^{(9/10+2\delta)R}. \tag{547}$$

It follows from (544) and (546) that

$$D_g(P) \equiv 0 \pmod{\mathfrak{A}_R} \quad \Longrightarrow \quad D_g(P) = 0 \tag{548}$$

for all sufficiently large R. Indeed, if the left-hand side held for a nonzero defect, divisibility would give $\mathrm{N}\,\mathfrak{A}_R \le |\,\mathrm{N}\,D_g(P)|$, contradicting (547).

The complete finite quotient has an exact density. By the Chinese remainder theorem and (541),

$$\begin{aligned}\rho(\mathfrak{A}_R) &:= \frac{1}{|\,\mathrm{SL}_2(\mathcal{O}_K/\mathfrak{A}_R)|}\#\{g : D_g(P) = 0 \pmod{\mathfrak{A}_R}\} \\ &= \prod_{\mathfrak{p}\mid\mathfrak{A}_R} \frac{1}{\mathrm{N}\,\mathfrak{p}+1} \\ &= \frac{1}{\mathrm{N}\,\mathfrak{A}_R}\prod_{\mathfrak{p}\mid\mathfrak{A}_R} \frac{\mathrm{N}\,\mathfrak{p}}{\mathrm{N}\,\mathfrak{p}+1}.\end{aligned} \tag{549}$$

Equivalently, the orbit of $[P]$ is the single product space

$$Y_R = \mathbb{P}^1(\mathcal{O}_K/\mathfrak{A}_R) = \prod_{\mathfrak{p}\mid\mathfrak{A}_R} \mathbb{P}^1(k_{\mathfrak{p}}), \qquad |Y_R| = \rho(\mathfrak{A}_R)^{-1}. \tag{550}$$

Thus the desired count is a return to one fixed point of one finite thin-group action. The product notation records its local structure; it does not split the proof into independent searches. The number-field Mertens formula gives

$$\prod_{\mathrm{N}\,\mathfrak{p}\le X} \frac{\mathrm{N}\,\mathfrak{p}}{\mathrm{N}\,\mathfrak{p}+1} \asymp_K \frac{1}{\log X}, \tag{551}$$

after deleting the finitely many bad primes. Hence

$$M_R\rho(\mathfrak{A}_R) = \frac{e^{(1-u)R-o(R)}}{\log R}. \tag{552}$$

The expected fibre is exponentially large even though the modulus is larger than every possible nonzero defect.

Theorem A.4 (Conditional completion from one adelic local limit)**.** *Suppose that for every primitive $P \in \mathcal{O}_K^2$ whose projective G_7-stabilizer is trivial, and every u satisfying (544), the Hodge-balanced shell satisfies*

$$\#\{g \in \widetilde{\Omega}_R(\delta) : D_g(P) \equiv 0 \pmod{\mathfrak{A}_R}\} = (1 + o_P(1))M_R\rho(\mathfrak{A}_R). \tag{553}$$

Then Conjecture 1.1 holds.

Proof. Suppose that a primitive P has trivial projective stabilizer. The right-hand side of (553) tends to infinity by (552); the fibre is therefore nonempty for all large R. Implication (548) gives an element with $D_g(P) = 0$. Since the shell is annular, this element is not projectively the identity. It belongs to $G_7 \cap B_P(K)$, contradicting triviality. Proposition 7.1 completes the classification. □

A deliberately strong pointed square-root hypothesis is sufficient for the local limit in Theorem A.4.

Conjecture A.5 (Pointed relative square-root bound for the Hodge shell)**.** *If the projective G_7-stabilizer of a primitive P is trivial, then for every $\epsilon > 0$ and the ideals $\mathfrak{A}_R$ above,*

$$\begin{aligned}\Big|\#\{g \in \widetilde{\Omega}_R(\delta) : D_g(P) \equiv 0 \pmod{\mathfrak{A}_R}\} - M_R\rho(\mathfrak{A}_R)\Big| \\ \ll_{P,\epsilon} \big(M_R\rho(\mathfrak{A}_R)\big)^{1/2}(\mathrm{N}\,\mathfrak{A}_R)^{\epsilon}.\end{aligned} \tag{554}$$

This is "relative" because the square root is taken after restricting to the expected fibre; a bare operator-norm error of size $M_R^{1/2}$ would forget the codimension-one local density and stop at half entropy. For fixed $u < 1$, choose $\epsilon < (1-u)/(2u)$. Equations (546) and (552) then show that the right-hand side of (554) is $o(M_R\rho(\mathfrak{A}_R))$. Thus Conjecture A.5 implies (553), and hence the classification.

The following elementary observation is essential for judging the force of this hypothesis.

Proposition A.6 (Why the pointed estimate already contains rigidity)**.** *If the projective stabilizer of $[P]$ is trivial, then for all sufficiently large R,*

$$\#\{g \in \widetilde{\Omega}_R(\delta) : D_g(P) \equiv 0 \pmod{\mathfrak{A}_R}\} = 0.$$

Consequently Conjecture A.5 cannot be a black-box equidistribution theorem valid for arbitrary thin-group actions. Any proof of it in the heptagonal setting must use an inverse alternative which turns persistent non-equidistribution into a rational Borel stabilizer.

Proof. Implication (548) turns every element counted on the left into an exact projective stabilizer. Under the stated assumption there is no such element. □

This proposition is the point at which earlier formulations were too optimistic. Calling (554) "Lindelöf" describes its square-root scale; it does not identify it with a standard automorphic L-function estimate. It is a pointed statement for the Hodge-tilted transfer operator at one coherent rational fibre. The semisimple torus theorem supplies part of the desired inverse mechanism: a coefficient which refuses to flatten must concentrate near rational torsion. The exact transform below gives the signed value at rational torsion. The remaining work has two logically distinct parts:

1. transfer that localization to the equilibrium Markov measure with the prescribed Hodge exponents and sum the rational torsion contributions with their signs;

2. prove that a torsion resonance persisting through the nested rational fibres comes from a nonidentity element of $G_7 \cap B_P(K)$, rather than from an unrelated high-injectivity phenomenon in the finite Schreier graphs.

The second part is the analytic form of the exceptional-ray rigidity in the proof spine.

The relative local limit (553) is the one estimate pursued by the present thin-group route. Its Fourier form explains both why absolute values fail and why the exact local transform matters. Finite Fourier inversion gives

$$\#\{g \in \widetilde{\Omega}_R(\delta) : D_g(P) \equiv 0 \pmod{\mathfrak{A}_R}\} = \frac{1}{\mathrm{N}\,\mathfrak{A}_R} \sum_{\xi \in \mathfrak{A}_R^{-1}\mathfrak{d}_K^{-1}/\mathfrak{d}_K^{-1}} S_R(P;\xi). \tag{555}$$

At one good prime, the signed sum of the complete local coefficients is

$$1 + (q-1)\left(-\frac{1}{q^2-1}\right) = \frac{q}{q+1}, \tag{556}$$

whereas the sum of their absolute values is

$$1 + (q-1)\frac{1}{q^2-1} = 1 + \frac{1}{q+1}. \tag{557}$$

Multiplying (556) over the factors of $\mathfrak{A}_R$ gives the factor $(\log R)^{-1}$ in (552); multiplying (557) gives a factor of order $\log R$. Thus an absolute minor-arc argument loses two logarithms relative to the signed answer and, more importantly, suppresses the mechanism which creates the correct local density. The signed torsion sum (555) is the correct endpoint.

The rational Borel is already visible in the Fourier representation. Write $P = (p,q)^t$ and $\ell_P = (-q,p)$, so that $D_g(P) = \ell_P g P$. The embedding

$$J_P : \mathbb{T}_K \longrightarrow (K_\infty^2)^*/(\mathcal{O}_K^2)^\vee, \qquad \xi \longmapsto \xi \ell_P, \tag{558}$$

turns $S_R(P;\xi)/M_R$ into the Fourier coefficient, at the fixed integral character P, of the contragredient torus orbit of $J_P(\xi)$. Moreover

$$B_P(K) = \{g \in \mathrm{SL}_2(K) : \ell_P g \in K^\times \ell_P\}. \tag{559}$$

Thus the desired inverse alternative has the correct algebraic endpoint. If the transfer analysis produces an invariant rational normal line, its stabilizer is exactly the Borel required by Proposition 7.1. The semisimple torus theorem used below first localizes large coefficients near rational torsion; it does not, by itself, turn torsion into a Borel intersection.

The relevant representation is canonical. Choose an integral basis of $\mathcal{O}_K$ and view $\mathcal{O}_K^2$ as $\mathbb{Z}^6$. The action of $\widetilde{G}_7$ is then an integral six-dimensional torus action. Its rational Zariski closure is $\mathrm{Res}_{K/\mathbb{Q}}\,\mathrm{SL}_2$, while over $\mathbb{R}$ it is the product of the three standard two-dimensional representations.

This real splitting is handled by the later semisimple torus theorem of He–de Saxcé [49]. We record its exact specialization because it is the first input here that uses the full thin group rather than a presentation of it. Let μ be a finite-exponential-moment probability measure whose support generates $\widetilde{G}_7$. If, for a fixed nonzero integral frequency P, the Fourier coefficient of the n-step contragredient walk starting at $J_P(\xi)$ has modulus at least t, then, for $n \geq C\log(\|P\|/t)$,

$$\widetilde{d}_\mu\left(J_P(\xi), \frac{a}{q}\right) \leq e^{-\vartheta n}, \qquad 0 < |q| \leq (\|P\|/t)^C, \tag{560}$$

for every fixed $0 < \vartheta < 1$, with $C = C(\mu, \vartheta)$. Here $\widetilde{d}_\mu$ is the anisotropic distance whose coordinate exponents are the three top Lyapunov exponents of the walk.

There is no hidden invariant subtorus in this application. If E is the real algebra generated by the identity component of the Zariski closure, then

$$E = M_2(\mathbb{R}) \times M_2(\mathbb{R}) \times M_2(\mathbb{R}).$$

All three factors are noncompact, so the compact summand called V_0 in the theorem of He–de Saxcé is zero. The image of the nonzero vector P in each Galois plane is nonzero; consequently $PE = (\mathbb{R}^6)^*$, and the exceptional space $W_0 = (PE)^\perp$ in that theorem is also zero. Formula (560) therefore says exactly that a large coefficient can occur only near a rational torsion point. This calculation, not merely the citation of a general equidistribution theorem, is where Zariski density of the complete group enters.

If this inverse theorem is transferred from an independent random walk to the equilibrium Markov measure defining the Hodge-balanced shell, its three Lyapunov widths, measured in central displacement R, are

$$e^{-\vartheta R/2}, \qquad e^{-\vartheta R/10}, \qquad e^{-3\vartheta R/10}.$$

On the three-dimensional subtorus $J_P(\mathbb{T}_K)$, the corresponding anisotropic ball has volume

$$e^{-9\vartheta R/10}. \tag{561}$$

Thus the existing theorem finds precisely the Hodge geometry required by (537); it does not lose the outer factors by taking a largest operator norm.

Two analytic operations are not contained in the published theorem. First, the measure in this paper is a Bowen–Margulis annular shell with specified Hodge exponents, not a convolution power of independent increments. Second, the hypothesis $n \geq C \log(\|P\|/t)$ does not by itself give a summable estimate at the exponentially small threshold $t = e^{-9R/10+o(R)}$. Li's Fourier-decay and renewal theorem for split semisimple groups [45] supplies a natural mechanism for the first transfer, but it does not supply the required lattice-scale local limit by itself. These two operations, combined with the signed torsion calculation above, are the precise analytic work still needed.

The threshold can be read directly from the single finite target Y_R in (550). A bare operator-norm error of size $M_R^{1/2}$ is smaller than the main term $M_R/|Y_R|$ only when $|Y_R| < M_R^{1/2}$, the half-entropy range. The relative error $(M_R/|Y_R|)^{1/2}$ in (554) is smaller whenever $|Y_R| < M_R$, which is precisely $u < 1$. The Hodge norm comparison requires $u > 9/10$. Hence the full usable interval is the visible one-tenth entropy gap

$$9/10 < u < 1.$$

There is a useful graph-theoretic way to see exactly what this asks for. The finite set Y_R is a Schreier graph for the full group action: its vertices are the finite orbit points and its labeled edges record the action of a fixed generating set. The point $[P]$ is a distinguished vertex compatible through the nested ideals $\mathfrak{A}_R$. The left side of (553) counts Hodge-weighted based loops at that vertex. Since $|Y_R| = e^{uR+o(R)}$ and the shell has $e^{R-o(R)}$ paths, the first random-model based loops appear when R is only slightly larger than $\log|Y_R|$. This is the window $u < 1$.

An ordinary spectral-gap or Ramanujan estimate controls the worst matrix coefficient by the square root of the total number of paths. It gives an error of size $e^{R/2+o(R)}$, which is smaller than the main term only for $u < 1/2$. The Hodge norm comparison needs $u > 9/10$. Thus the missing theorem is not merely expansion. It is a fixed-vertex return theorem at the information-theoretic time, together with an inverse statement for a coherent vertex that repeatedly has abnormally

large injectivity radius. In the present tower, the desired inverse statement is that such coherence forces a characteristic-zero Borel stabilizer. This formulation explains both why the full thin group matters and why generic finite-graph technology stops at exactly the wrong scale.

There is a concrete operator on which to prove this estimate. The Bowen–Series coding expresses a geodesic segment as a path in a Markov graph. Its transfer operator is simply the weighted adjacency operator which sums over all legal predecessor paths. Enlarge the state by Y_R, let each path act on that finite coordinate, and weight it by three real parameters: central displacement and the two outer Hodge cocycles. Inverse Laplace–Fourier transform in those parameters selects the annulus and the Hodge window. The constant functions on Y_R give $M_R\rho(\mathfrak{A}_R)$; the matrix coefficient between the two delta functions at $[P]$ is the left side of (553). Thus the remaining task is to prove (554) on the orthogonal complement of the constant finite state. The Markov alphabet is only a coordinate system for this operator. Its inverse theorem uses the semisimple Zariski closure of the entire group, and its main term uses the complete adelic quotient.

Proposition A.7 (The sparse-intersection form of the obstruction)**.** *Let $\mathcal{A}_R, \mathcal{B}_R \subset G_7$ be compatible half-shells such that*

$$\#\mathcal{A}_R, \#\mathcal{B}_R = e^{R/2-o(R)}, \qquad R(a), R(b) \le R/2 + O(1),$$

and suppose every block satisfies

$$a_1(a) + a_2(a),\ a_1(b) + a_2(b) \le (1/5 + \epsilon)R.$$

If $[P]$ has trivial projective stabilizer and $\mathrm{N}\,\mathfrak{p} > e^{(9/10+3\epsilon)R}$, then both reduction maps

$$a \longmapsto [a^{-1}P \bmod \mathfrak{p}], \qquad b \longmapsto [bP \bmod \mathfrak{p}]$$

are injective. Writing $\alpha_\mathfrak{p}$ and $\beta_\mathfrak{p}$ for the normalized endpoint measures, one has

$$\|\alpha_\mathfrak{p} - \mathbf{u}\|_2 \asymp e^{-R/4+o(R)}, \qquad \|\beta_\mathfrak{p} - \mathbf{u}\|_2 \asymp e^{-R/4+o(R)},$$

where $\mathbf{u}$ is uniform measure on the projective line. Consequently marginal L^2 information alone bounds their covariance only by $e^{-R/2+o(R)}$, whereas its constant term is $(\mathrm{N}\,\mathfrak{p})^{-1} = e^{-uR+o(R)}$ when $\mathrm{N}\,\mathfrak{p} = e^{uR+o(R)}$.

Proof. For distinct a, a', triviality of the stabilizer makes

$$\Delta_{a,a'} = \det(a^{-1}P, a'^{-1}P)$$

nonzero. The central operator bound contributes at most $R/2 + O_P(1)$ to $\log|\mathrm{N}\,\Delta_{a,a'}|$, and the two outer places contribute at most

$$a_1(a) + a_2(a) + a_1(a') + a_2(a') \le (2/5 + 2\epsilon)R.$$

Hence

$$\log|\mathrm{N}\,\Delta_{a,a'}| \le (9/10 + 2\epsilon)R + O_P(1).$$

Such a nonzero algebraic integer cannot be divisible by $\mathfrak{p}$ at the stated scale. This proves injectivity on the left; the proof for $\det(bP, b'P)$ is identical.

If $A = \#\mathcal{A}_R$, injectivity gives

$$\|\alpha_\mathfrak{p} - \mathbf{u}\|_2^2 = \frac{1}{A} - \frac{1}{\mathrm{N}\,\mathfrak{p} + 1} = e^{-R/2+o(R)}.$$

The same formula holds on the right. Cauchy–Schwarz now gives the claimed covariance scale. □

This proposition identifies the obstruction rather than estimating it away. At $u > 9/10$, each half-shell is an injective subset of size

$$e^{R/2} = q^{1/(2u)} = q^{5/9+o(1)}$$

at the endpoint $u = 9/10$. Two subsets of a q-point projective line of this size can be disjoint despite having the expected marginal energies. The desired collision count is $e^{(1-u)R+o(R)}$. Obtaining it requires improving the product of the two marginal L^2-norms by the additional factor $q^{-1/2}$. This is exactly the coherent Steinberg delocalization in the relative square-root estimate. Thus the place where the proof stops is a critical lower incidence theorem for two sparse complete block orbits, not a missing inequality for a generator.

There is a simple test which separates a calculation with the letters from a statement about the group. Put $a_i(g) = \log \|\sigma_i(g)\|_{\mathrm{op}}$. Since

$$\sigma_i(T^n) = \begin{pmatrix} 1 & n\lambda_i \\ 0 & 1 \end{pmatrix},$$

one has

$$a_i(T^n) = \log n + \log|\lambda_i| + o(1) \qquad (i = 0, 1, 2). \tag{562}$$

In particular,

$$\frac{a_i(T^n)}{a_0(T^n)} \longrightarrow 1, \qquad \frac{a_1(T^n) + a_2(T^n)}{a_0(T^n)} \longrightarrow 2.$$

Thus there is no worst-word estimate which bounds the outer Cartan sum by $\theta a_0 + O(1)$ with $\theta < 2$, let alone with the Hodge-typical coefficient $4/5$ in Cartan normalization. Accelerating a parabolic run into one Rosen digit does not alter this fact; it only hides the parameter n inside one letter.

This observation is elementary, but it rules out an entire proof style. The useful coefficient $4/5$, equivalently $2/5$ in hyperbolic-radius normalization, is an average for the Bowen–Margulis distribution of complete geodesic segments. It is not a pointwise inequality for the generators or for all words. The role of a Markov alphabet is merely to parametrize that distribution and its stopping times. Any argument which takes the largest loss over the six letters has discarded the thin-group information before the arithmetic problem begins.

This proposition explains why calculations with the six Markov letters cannot be the conceptual core of a proof. A polynomial identity on the ambient group which holds for every element of G_7 already holds on all of SL_2^3. Letter inequalities can still prove useful analytic domination on a selected positive semigroup, but they cannot manufacture the arithmetic intersection (78). That intersection is a statement about the integral orbit of the thin group inside the rational flag variety.

To state the global problem without a coding, embed the rational flag variety diagonally by Galois conjugation:

$$\Delta_K : \mathbf{P}^1(K) \longrightarrow \mathbf{P}^1(\mathbf{R})^3, \qquad x \longmapsto (\sigma_0 x, \sigma_1 x, \sigma_2 x).$$

Then Proposition 7.1 says exactly

$$\mathrm{Stab}_{G_7}\big(\Delta_K(x)\big) \neq \{\pm 1\} \qquad (x \in \mathbf{P}^1(K)). \tag{563}$$

Thus the target is an arithmetic rigidity theorem for rational flags under a Zariski-dense thin group. The genuinely group-wide structures available for such a theorem are:

1. the relatively hyperbolic boundary of the $(2, 7, \infty)$ group and its single orbit of parabolic points;

2. the three-factor limit flags of the Galois representation;

3. the strict domination supplied by the two modular embeddings; and

4. the integral $\mathcal{O}_K$-lattice, which turns recurrence of a complete reduced flag into an exact stabilizer.

The Rosen alphabet is only a finite-state coordinate system for the first of these structures. Its proper role is to recognize first returns and, at the end, to turn bounded integral data into repetition.

The missing quantitative theorem can now be described in group-wide terms. The cusp orbit consists of the flags $c_g = g^{-1}\infty$. For a rational flag x, the incidence

$$r_g = \det(x, c_g)$$

is its lower Bruhat coordinate relative to B_x. At the first prefixes for which $|\,\mathrm{N}\, r_g|$ reaches a prescribed level, the modular-embedding domination must force the global incidence lower bound (279). Equality, or zero domination, must force $g_2 g_1^{-1} \in B_x$, which is exactly (78). This is a thin-orbit separation theorem; it is stronger than a singular-value estimate for arbitrary words and more specific than an ordinary Roth theorem for arbitrary approximants.

There is an important correction to the phrase "use the thin group." Thinness itself is not a source of closing. Zariski density, strong approximation, and expansion are properties shared by many thin groups for which a prescribed rational Borel can have trivial intersection with the group. Those properties say that the orbit looks large after reduction modulo almost every ideal; they do not say that it meets the codimension-one subvariety B_x over K. In particular, the congruence condition

$$D_g(P) = \det(P, gP) \equiv 0 \pmod{\mathfrak{a}}$$

is a finite-quotient shadow of the desired equality, not evidence that the equality must occur.

What is special here is the combination of the whole group with its Hodge-correlated Galois representation. The uniformizing factor supplies entropy one; the two modular embeddings supply the strict outer contractions and the degrees $1/28$ and $3/28$; and the single $\mathcal{O}_K$-lattice couples all three real factors by the product formula. The actual missing assertion is therefore not a theorem about thin groups in general. It is the following rigidity statement for this particular integral Hodge orbit:

> If the Hodge-balanced orbit of a fixed rational flag persistently avoids its rational Borel, then its congruence and real boundary coordinates cannot remain locked together at every scale.

Equivalently, failure of the pointed dispersion estimate must produce a nonidentity element of $G_7 \cap B_x(K)$. The displayed matrices can parametrize this problem, but no inequality for those matrices proves that inverse implication.

A.2 Earlier congruence formulations and their limitation

The exact quantitative issue can be measured without choosing a word alphabet. Fix $x = [P] \in \mathbf{P}^1(K)$, with $P \in \mathcal{O}_K^2$ primitive, and put $e = (1, 0)^t$. For a good prime ideal $\mathfrak{p}$, a number $\theta > 0$, and a fixed constant C, define the *Hodge-constrained covering radius*

$$\mathcal{R}_{\theta,C}(x; \mathfrak{p}) = \inf \left\{ R(g) : \begin{array}{l} g \in G_7, \quad \det(P, ge) \in \mathfrak{p}, \\ a_1(g) + a_2(g) \le \theta R(g) + C \end{array} \right\}. \tag{564}$$

Here $R(g) = d_{\mathbb{H}}(i, \sigma_0(g)i)$. The determinant condition says that the cusp $g\infty$ has the prescribed reduction x mod $\mathfrak{p}$. Proposition 7.3 makes the set in (564) nonempty: first prescribe the finite class, and then use the dense outer image of its congruence fibre to make both outer matrices bounded. Distinct such lifts escape in the discrete uniformizing factor, so eventually the displayed Hodge inequality holds. This is a qualitative consequence for every residue class, not an average over classes.

Define the corresponding exponent, allowing a fixed error in the Hodge window, by

$$\kappa_\theta(x) = \inf_{C \geq 0} \liminf_{\mathrm{N}\,\mathfrak{p} \to \infty} \frac{\mathcal{R}_{\theta,C}(x;\mathfrak{p})}{\log \mathrm{N}\,\mathfrak{p}}, \tag{565}$$

where the limit runs through good primes. This invariant exposes the precise use still required from the thin group.

Proposition A.8 (The product-formula wall for the thin orbit)**.** *If x is not a cusp, then*

$$\kappa_\theta(x) \geq \frac{1}{1/2 + \theta}. \tag{566}$$

If x is a cusp, then $\kappa_\theta(x) = 0$. In particular, at the Hodge value $\theta = 2/5$, every noncuspidal rational flag has covering exponent at least $10/9$, whereas a cusp has exponent zero.

Proof. Put $d_g = \det(P, ge)$. If x is not a cusp, then $d_g \neq 0$ for every $g \in G_7$. Whenever g occurs in (564), integrality and $d_g \in \mathfrak{p}$ give

$$\log \mathrm{N}\,\mathfrak{p} \leq \log |\,\mathrm{N}\, d_g|.$$

At the three real places, the determinant estimate and $a_0(g) \leq R(g)/2 + O(1)$ give

$$\log |\,\mathrm{N}\, d_g| \leq a_0(g) + a_1(g) + a_2(g) + O_P(1) \leq (1/2 + \theta)R(g) + C + O_P(1).$$

Take the infimum over g, divide by $\log \mathrm{N}\,\mathfrak{p}$, and then take the two operations in (565). This proves (566). If $x = h\infty$ is a cusp, then $d_h = 0$. After choosing C large enough that the fixed element h satisfies the Hodge window, the covering radius is bounded by $R(h)$ for every $\mathfrak{p}$, so the exponent is zero. □

There is now a clean formulation of what it would mean to exploit the thin group at full strength. A ball of radius R contains $e^{R+o(R)}$ elements, while a projective congruence orbit has $\mathrm{N}\,\mathfrak{p} + 1$ points. The information-theoretic prediction is therefore $\kappa_{2/5}(x) = 1$ for a point with no coherent obstruction. Proposition A.8 says that this prediction is arithmetically impossible for a noncuspidal K-rational flag. Hence the following inverse statement would prove the theorem:

> If a coherent rational section of the projective congruence tower has $\kappa_{2/5}(x) > 1$, then its rational Borel already meets G_7 nontrivially.

The special fixed points are not counterexamples: their nontrivial Borel intersection is exactly the permitted obstruction, and the proposition shows that, being noncuspidal, they lie at or beyond the $10/9$ wall. Thus the problem is not to prove a uniform diameter theorem for all residue classes. It is to prove a pointwise inverse theorem for the particular thin orbit, with stabilizers as the exceptional case.

There is a second, more global way to use the same wall. It implements the Arakelov principle that all finite places should be charged at once. For $\delta > 0$, let

$$\Omega_R(\delta) = \{g \in G_7 : R \leq R(g) < R + 1,\ a_1(g) + a_2(g) \leq (2/5 + \delta)R\}, \qquad M_R = \#\Omega_R(\delta).$$

For a primitive P, put

$$D_g(P) = \det(P, gP)$$

and, at a good prime $\mathfrak{p}$, set

$$N_{\mathfrak{p}}(P;R) = \#\{g \in \Omega_R(\delta) : g[P] = [P] \pmod{\mathfrak{p}}\}.$$

Uniform distribution in the projective quotient predicts $N_{\mathfrak{p}} = M_R/(\mathrm{N}\,\mathfrak{p}+1)$.

Proposition A.9 (Signed global prime-sum closing)**.** *Fix u, δ with*

$$9/10 + \delta < u < 1$$

and suppose $M_R = e^{R-o(R)}$. If

$$\left| \sum_{\substack{\mathfrak{p} \text{ good} \\ \mathrm{N}\,\mathfrak{p} \le e^{uR}}} (\log \mathrm{N}\,\mathfrak{p}) \left(N_{\mathfrak{p}}(P;R) - \frac{M_R}{\mathrm{N}\,\mathfrak{p}+1} \right) \right| = o(RM_R), \tag{567}$$

then $[P]$ has a projectively nonidentity stabilizer in G_7.

Proof. Suppose otherwise. Then $D_g(P) \neq 0$ for every g in the shell. The Hodge window and the determinant estimate at all three real places give

$$\log|\mathrm{N}\,D_g(P)| \le (9/10+\delta)R + O_P(1). \tag{568}$$

On the other hand, interchange the order of summation to obtain

$$\begin{aligned} \sum_{\substack{\mathfrak{p} \text{ good} \\ \mathrm{N}\,\mathfrak{p} \le e^{uR}}} (\log \mathrm{N}\,\mathfrak{p}) N_{\mathfrak{p}}(P;R) &= \sum_{g \in \Omega_R(\delta)} \sum_{\substack{\mathfrak{p} \mid D_g(P) \\ \mathfrak{p} \text{ good},\ \mathrm{N}\,\mathfrak{p} \le e^{uR}}} \log \mathrm{N}\,\mathfrak{p} \\ &\le \sum_{g \in \Omega_R(\delta)} \log|\mathrm{N}\,D_g(P)| \\ &\le (9/10+\delta)RM_R + O_P(M_R). \end{aligned} \tag{569}$$

The prime-ideal Mertens formula, unchanged by deleting finitely many bad primes, says

$$\sum_{\substack{\mathfrak{p} \text{ good} \\ \mathrm{N}\,\mathfrak{p} \le X}} \frac{\log \mathrm{N}\,\mathfrak{p}}{\mathrm{N}\,\mathfrak{p}+1} = \log X + o(\log X).$$

Consequently (567) makes the left side of (569) equal to $(u+o(1))RM_R$, contradicting $u > 9/10+\delta$. □

This proposition is one divisor calculation in the Arakelov norm of $D_g(P)$, rather than a selection argument at individual finite places. Notice that (567) is a *signed* aggregate: errors at different primes may cancel. An absolute-value Elliott–Halberstam estimate is therefore unnecessarily strong. This is weaker still than finding a hit at each prime; only one weighted first moment of the complete Hodge-balanced shell is used.

There is a standard-shaped variance statement which supplies exactly this signed first moment. Put

$$E_{\mathfrak{p}}(P;R) = N_{\mathfrak{p}}(P;R) - \frac{M_R}{\mathrm{N}\,\mathfrak{p}+1}$$

and define

$$\mathcal{V}_R(P;u) = \sum_{\substack{\mathfrak{p}\ \text{good}\\ \mathrm{N}\,\mathfrak{p}\le e^{uR}}} (\log \mathrm{N}\,\mathfrak{p})|E_{\mathfrak{p}}(P;R)|^2. \tag{570}$$

Proposition A.10 (Steinberg large-sieve and common-divisor forms)**.** *Let $q_{\mathfrak{p}} = \mathrm{N}\,\mathfrak{p}$, let $\mathbf{u}_{\mathfrak{p}}$ be the uniform vector on $\mathbf{P}^1(\mathcal{O}_K/\mathfrak{p})$, and put*

$$e_{P,\mathfrak{p}} = \delta_{[P \bmod \mathfrak{p}]} - \mathbf{u}_{\mathfrak{p}}.$$

If $\mu_R = M_R^{-1}\sum_{g\in\Omega_R(\delta)}\delta_g$, then

$$E_{\mathfrak{p}}(P;R) = M_R\,\langle \mathrm{St}_{\mathfrak{p}}(\mu_R)e_{P,\mathfrak{p}}, e_{P,\mathfrak{p}}\rangle\,, \tag{571}$$

$$\mathcal{V}_R(P;u) = M_R^2 \sum_{\mathrm{N}\,\mathfrak{p}\le e^{uR}} (\log \mathrm{N}\,\mathfrak{p})\,|\langle \mathrm{St}_{\mathfrak{p}}(\mu_R)e_{P,\mathfrak{p}}, e_{P,\mathfrak{p}}\rangle|^2\,. \tag{572}$$

There is also an exact group-theoretic form of the uncentered correlation. Write

$$I_{\mathfrak{p}}(g) = \mathbf{1}_{\{\mathfrak{p}\mid D_g(P)\}}.$$

For all $g,h\in G_7$,

$$I_{\mathfrak{p}}(g)I_{\mathfrak{p}}(h) = I_{\mathfrak{p}}(h)I_{\mathfrak{p}}(h^{-1}g). \tag{573}$$

In particular, every prime dividing both $D_g(P)$ and $D_h(P)$ also divides $D_{h^{-1}g}(P)$. If

$$\omega_Q(\mathfrak{a}) = \sum_{\substack{\mathfrak{p}\mid\mathfrak{a}\\ \mathrm{N}\,\mathfrak{p}\le Q}} \log \mathrm{N}\,\mathfrak{p},$$

with each prime counted once, then

$$\sum_{\mathrm{N}\,\mathfrak{p}\le Q} (\log \mathrm{N}\,\mathfrak{p})I_{\mathfrak{p}}(g)I_{\mathfrak{p}}(h) = \omega_Q\big((D_g(P)) + (D_h(P))\big). \tag{574}$$

Proof. The projective permutation representation is $\mathbf{1}\oplus\mathrm{St}_{\mathfrak{p}}$, and

$$\langle \mathrm{St}_{\mathfrak{p}}(g)e_{P,\mathfrak{p}}, e_{P,\mathfrak{p}}\rangle = \mathbf{1}_{\{g[P]=[P] \bmod \mathfrak{p}\}} - \frac{1}{q_{\mathfrak{p}}+1}.$$

Sum over the shell and then over primes to obtain (571) and (572).

If $g[P]$ and $h[P]$ are both $[P]$ modulo $\mathfrak{p}$, then $h^{-1}g[P] = [P]$. Conversely, fixing by h and by $h^{-1}g$ implies fixing by g. This proves (573). Formula (574) merely sums this common-prime condition; the support inclusion follows from the difference law. □

Thus (575) is equivalently the family large-sieve estimate

$$\sum_{\mathrm{N}\,\mathfrak{p}\le e^{uR}} (\log \mathrm{N}\,\mathfrak{p})\,|\langle \mathrm{St}_{\mathfrak{p}}(\mu_R)e_{P,\mathfrak{p}}, e_{P,\mathfrak{p}}\rangle|^2 \le M_R^{-1}e^{o(R)}.$$

This formulation uses the complete shell and the full family of projective congruence representations. The continued-fraction generators do not occur. On expanding the square it asks for cancellation between the common-divisor energy (574) and its two centering terms. The identity also explains why ordinary expansion at each prime is too weak: it estimates the individual matrix coefficients, whereas the needed gain is orthogonality of these coherent coefficients across the prime family.

The conjectural status can now be placed on a familiar number-theoretic axis. Estimate (575) is the fixed-residue analogue of a level-one Barban–Davenport–Halberstam or Elliott–Halberstam assertion for an affine-sieve orbit. A Ramanujan bound for each congruence quotient does not imply it: that bound controls each summand in (572) separately and loses the prime-family orthogonality. Bilinear dispersion is precisely the technology which can go beyond expansion in thin orbits [69]; the two-prime formulation below exposes two variables below half entropy for that reason. The available affine-sieve estimates do not include the present fixed coherent Borel target, the Hodge cutoff, or level $u > 9/10$. Thus a standard-shaped conditional completion is now precise: *level-one fixed-section BDH for the Hodge-balanced G_7-orbit implies the theorem.* Ordinary Roth, GRH for scalar number-field L-functions, and quotient-by-quotient Ramanujan do not by themselves contain this family variance.

Conjecture A.11 (Fixed-section thin-orbit BDH)**.** *If $[P]$ has projectively trivial stabilizer, then for every fixed $u < 1$,*

$$\mathcal{V}_R(P;u) \le M_R e^{o(R)}. \tag{575}$$

Proposition A.12 (The fixed-section variance completes the proof)**.** *Assume $M_R = e^{R-o(R)}$. For any $9/10+\delta < u < 1$, the estimate (575) implies (567), and hence the heptagonal cusp classification.*

Proof. The prime-ideal theorem gives

$$\sum_{\mathrm{N}\,\mathfrak{p}\le e^{uR}} \log \mathrm{N}\,\mathfrak{p} = e^{uR+o(R)}.$$

Cauchy–Schwarz and (575) therefore give

$$\left|\sum_{\mathrm{N}\,\mathfrak{p}\le e^{uR}} (\log \mathrm{N}\,\mathfrak{p}) E_{\mathfrak{p}}(P;R)\right| \le e^{uR/2+o(R)} M_R^{1/2} e^{o(R)} = e^{(1+u)R/2+o(R)}.$$

Since $u < 1$, this is $o(RM_R)$. Proposition A.9 applies. □

The variance gap under the contrary hypothesis is exact. Choose

$$9/10 + \delta < v < u.$$

If $[P]$ has trivial stabilizer, then (568) shows that for all sufficiently large R,

$$N_{\mathfrak{p}}(P;R) = 0 \qquad (e^{vR} \le \mathrm{N}\,\mathfrak{p} \le 2e^{vR}).$$

Indeed, divisibility of a nonzero $D_g(P)$ by such a prime contradicts its norm bound. Hence on this dyadic prime interval

$$E_{\mathfrak{p}}(P;R) = -\frac{M_R}{\mathrm{N}\,\mathfrak{p}+1}.$$

The prime-ideal theorem now gives the lower bound

$$\mathcal{V}_R(P;u) \ge M_R^2 \sum_{e^{vR}\le \mathrm{N}\,\mathfrak{p}\le 2e^{vR}} \frac{\log \mathrm{N}\,\mathfrak{p}}{(\mathrm{N}\,\mathfrak{p}+1)^2} = M_R^2 e^{-vR+o(R)} = e^{(2-v)R-o(R)}. \tag{576}$$

Thus the conjectural BDH scale $e^{R+o(R)}$ must save precisely the factor $e^{(1-v)R}$. This is the same coherent-hole obstruction seen in the two-prime dispersion formulation below, but (575) is its weakest signed prime-average form.

A.3 Reduction by the entire parabolic orbit

There is a second group-wide reformulation which makes precise what is lost when one checks only the six Rosen moves. Let

$$e = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \qquad f = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

For every $g \in G_7$, put $c = ge$, $d = gf$, and

$$U_{c,n} = gT^n g^{-1} \qquad (n \in \mathbf{Z}). \tag{577}$$

This is the full family of parabolic transvections in the cusp orbit. If an integral vector is written in the cusp frame as

$$P = uc + rd,$$

then the action is the elementary shear

$$U_{c,n}P = (u + n\lambda r)c + rd. \tag{578}$$

The point is not the formula, which is trivial, but the simultaneous requirement that follows when it is imposed for every $c \in G_7 e$.

Let h denote the absolute projective Weil height and define the global reduced set

$$\mathscr{R}_{G_7} = \{x \in \mathbf{P}^1(K) : h(\gamma x) \geq h(x) \text{ for every } \gamma \in G_7\}. \tag{579}$$

Proposition A.13 (Global thin-group reduction)**.** *Every G_7-orbit in $\mathbf{P}^1(K)$ meets $\mathscr{R}_{G_7}$. Conjecture 1.1 is equivalent to the assertion that every point of $\mathscr{R}_{G_7}$ has a projectively nontrivial G_7-stabilizer. Moreover, if $x = [P] \in \mathscr{R}_{G_7}$, then for every cusp frame and every $n \in \mathbf{Z}$,*

$$h\big([(u + n\lambda r)c + rd]\big) \geq h([uc + rd]). \tag{580}$$

Proof. Northcott finiteness says that $\mathbf{P}^1(K)$ contains only finitely many points of bounded height. Intersect the orbit of x with the nonempty finite set of points of height at most $h(x)$, and choose one of least height. It belongs to $\mathscr{R}_{G_7}$. A stabilizer of this minimal representative conjugates to a stabilizer of the original point, so Proposition 7.1 proves the equivalence. Finally (580) is the defining minimality applied to (577), using (578). □

This proposition turns the proposed classification into a Voronoi problem for the whole thin group. A calculation with the standard generators checks only a few faces of this infinite reduction domain. In contrast, (580) is indexed by the complete cusp orbit and is independent of a continued-fraction alphabet. If $r = 0$ in one frame, then $x = c$ is already a cusp. Otherwise one must show that the simultaneous minima in all cusp frames retain only finitely many complete Borel frames; repetition then gives a stabilizer. The rank-two unit lattice is precisely the reason that finiteness of the ordinary projective height does not prove this last assertion by itself. The scalar character of $\mathcal{N}$ detects only the sum of the two unit directions and is not proper on their rank-two lattice. The proper Hodge object must retain the two pairwise characters separately; this refinement is made in §17.1.

What is new and what is elementary. The matrix equivalence in Proposition 7.1 and the coordinate identity (81) are elementary. Their value is diagnostic: they identify the desired conclusion before any height is chosen. The genuinely geometric input begins with the embedding of the Teichmüller curve in the Hilbert modular threefold and the Hodge character $\mathcal{N}$. The missing theorem is a properness statement for their combination at first norm-hitting times, not a further estimate for word matrices in isolation.

A.4 The joint compression functor of the full cusp orbit

McMullen's compression of cusps gives a second global structure which is especially well adapted to the cubic field. It is defined on every pair of cusps, and not by checking the effect of a preferred generator.

For a cusp ξ, let u_ξ be its positive primitive parabolic generator. For distinct cusps ξ, η, set

$$A(\xi,\eta) = -\operatorname{Tr}\big((u_\xi - I)(u_\eta - I)\big) \in \mathcal{O}_K.$$

The positively oriented outer representations give two compression ratios

$$r_i(\xi,\eta) = \frac{\sigma_i A(\xi,\eta)}{\sigma_0 A(\xi,\eta)} \in (0,1], \qquad i = 1,2. \tag{581}$$

Proposition A.14 (Joint compression is the integral shape coordinate)**.** *Let v_ξ, v_η be primitive cusp columns and put $d = \det(v_\xi, v_\eta)$. Then*

$$A(\xi,\eta) = \lambda^2 d^2, \tag{582}$$

$$\log r_i(\xi,\eta) = 2\left(\log|\sigma_i d| - \log|\sigma_0 d|\right) + 2\log\frac{|\lambda_i|}{\lambda_0}. \tag{583}$$

On the modular-symbol category of G_7, the map

$$\mathbf{r} = (r_1, r_2) : \mathscr{S}(G_7) \longrightarrow [0,1]^2$$

is continuous and multiplicative under concatenation. If a sequence of symbols has thick-part length tending to infinity, then $\mathbf{r} \to (0,0)$.

Proof. Write $T = I + \lambda E_{12}$. If $v = (a,c)^t$ is the first column of g, then

$$gTg^{-1} - I = \lambda \begin{pmatrix} -ac & a^2 \\ -c^2 & ac \end{pmatrix} = \lambda N_v.$$

The rank-one identity $-\operatorname{Tr}(N_v N_w) = \det(v,w)^2$ proves (582); applying the embeddings and taking logarithms proves (583). Continuity and multiplicativity are McMullen's compression theorem applied to the two modular embeddings separately. His thick-part estimate gives

$$r_i(\xi,\eta) \le \exp(-\delta_i L_{\text{thick}}(\xi,\eta))$$

for constants $\delta_i > 0$, which proves the last assertion. □

Corollary A.15 (Radial norm and joint compression are proper)**.** *Fix $B > 0$ and $\delta > 0$. Up to the diagonal action of G_7, there are only finitely many oriented cusp pairs (ξ,η) satisfying*

$$|\,\mathrm{N}\,A(\xi,\eta)| \le B, \qquad r_1(\xi,\eta), r_2(\xi,\eta) \ge \delta. \tag{584}$$

Consequently an infinite sequence of distinct modular symbols with bounded $|\,\mathrm{N}\,A|$ must have $r_1 r_2 \to 0$ along a subsequence.

Proof. All three numbers $A_i = \sigma_i A$ are positive and

$$|\,\mathrm{N}\,A| = A_0 A_1 A_2 = A_0^3 r_1 r_2.$$

Thus (584) gives $A_0^3 \le B/\delta^2$. McMullen's parabolic-pairing formula identifies $\log A_0$ with the renormalized length of the cusp-to-cusp geodesic. On a finite-area hyperbolic orbifold there are only finitely many parabolic double cosets of bounded renormalized length. Equivalently, the relative word metric on $\langle T\rangle\backslash G_7/\langle T\rangle$ is proper. This proves finiteness. The last assertion is its contrapositive. □

The preceding construction has a consequence which is easy to miss if one works only with a continued-fraction alphabet. It puts every matrix entry of the whole group in one compact set of Galois shapes.

Theorem A.16 (The compact joint ratio set of the thin group)**.** *Let $\mathcal{E}_7$ be the set of nonzero entries of all matrices in G_7, and put*

$$\mathcal{R}_7 = \left\{ \left(\frac{|\sigma_1 m|}{|\sigma_0 m|}, \frac{|\sigma_2 m|}{|\sigma_0 m|} \right) : m \in \mathcal{E}_7 \right\}.$$

The closure of $\mathcal{R}_7$ is a compact countable subset of

$$\left[0, \frac{\lambda_0}{|\lambda_1|}\right] \times \left[0, \frac{\lambda_0}{|\lambda_2|}\right].$$

More precisely, every point of $\mathcal{R}_7$, up to finitely many sign choices, is the coordinatewise square root of a value of McMullen's joint compression functor, followed by the fixed rescaling

$$(s_1, s_2) \longmapsto \left(\frac{\lambda_0}{|\lambda_1|}\sqrt{s_1}, \frac{\lambda_0}{|\lambda_2|}\sqrt{s_2} \right).$$

If the modular symbols producing a sequence of entries escape through unbounded thick-part length, both coordinates tend to zero.

Proof. Let $e = (1,0)^t$ and $f = (0,1)^t$. If $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, then

$$c = \det(e, ge), \quad -a = \det(f, ge), \quad d = \det(e, gf), \quad -b = \det(f, gf).$$

The four vectors are primitive cusp columns. For any nonzero one of these determinants, equations (582) and (581) give

$$\frac{|\sigma_i m|}{|\sigma_0 m|} = \frac{\lambda_0}{|\lambda_i|}\sqrt{r_i(\xi,\eta)} \qquad (i = 1, 2).$$

McMullen's modular-symbol space, with its one point at infinity, is a compact countable space. Proposition A.14 extends $\mathbf{r}$ continuously to that space by sending the point at infinity to $(0,0)$. Its image is therefore compact and countable. Equation (A.4) proves all the assertions. □

This is where the word *thin* has mathematical content. In the distinguished copy of $\mathrm{SL}_2(\mathbf{R})$, the group is a lattice, and ordinary cusp equidistribution sees only that copy. In the full Hilbert modular group, units make the two ratios in (A.16) unbounded. The diagonal Galois image of G_7, by contrast, confines them to the compact countable joint ratio set. This is a global restriction on every group element, not an inequality for a preferred generator.

There is also a useful way to apply this cusp–cusp theorem to the incidence with a fixed rational point. It does not canonically normalize that rational point, but it avoids needing to do so.

Proposition A.17 (Two-anchor control of a rational incidence)**.** *Fix a primitive $P = (p, q)^t \in \mathcal{O}_K^2$. Let $v = (a, c)^t$ be a primitive cusp column representing a cusp $\xi \neq 0, \infty$, and define*

$$R_i(v) = \max\{r_i(\infty, \xi), r_i(0, \xi)\}, \qquad H_0(v) = \max\{|\sigma_0 a|, |\sigma_0 c|\}.$$

There are constants $C_{P,i}$, depending only on P, such that

$$\max\{|\sigma_i a|, |\sigma_i c|\} \leq \frac{\lambda_0}{|\lambda_i|} H_0(v)\sqrt{R_i(v)}, \tag{585}$$

$$|\,\mathrm{N}\det(P, v)| \leq C_P\,|\sigma_0 \det(P, v)| H_0(v)^2 \sqrt{R_1(v)R_2(v)}. \tag{586}$$

Consequently, if the right side of (586) is less than one, then $[P] = [v]$ and the rational point is a cusp.

Proof. The determinants of v with the primitive columns at ∞ and at 0 are respectively c and $-a$. Apply (A.4) to these two cusp pairs and take their maximum; this proves (585). Next,

$$|\sigma_i(pc - qa)| \leq (|\sigma_i p| + |\sigma_i q|) \max\{|\sigma_i a|, |\sigma_i c|\}.$$

Multiplying this estimate for $i = 1, 2$ by the central absolute value gives (586). A nonzero algebraic integer has absolute norm at least one, proving the last assertion. □

The maximum in Proposition A.17 cannot be replaced by the norm of one of the two entries. This failure already occurs for one element of the full group. Put $\alpha = \lambda^2 - 1$. The first column of

$$A_3 = \begin{pmatrix} \alpha & \alpha \\ \lambda & \alpha \end{pmatrix}$$

has

$$\frac{|\sigma_1\alpha|}{|\sigma_0\alpha|} > \frac{|\sigma_1\lambda|}{|\sigma_0\lambda|}, \qquad \frac{|\sigma_2\alpha|}{|\sigma_0\alpha|} < \frac{|\sigma_2\lambda|}{|\sigma_0\lambda|}.$$

Thus one entry supplies the first outer maximum and the other supplies the second. This is the smallest possible example of the rank-two obstruction; it is not caused by a long or badly chosen word.

There is nevertheless a canonical integral object which retains all four dominance chambers. Let τ generate $\mathrm{Gal}(K/\mathbf{Q})$, with $\sigma_i = \sigma_0\tau^i$, and for $v = (a, c)^t$ put

$$Z_{12}(v) = \tau(v) \otimes \tau^2(v) = \big(\tau(a)\tau^2(a), \tau(a)\tau^2(c), \tau(c)\tau^2(a), \tau(c)\tau^2(c)\big). \tag{587}$$

Proposition A.18 (The split maximum is one integral tensor norm)**.** *If $(a, c) = \mathcal{O}_K$, then the four coordinates of $Z_{12}(v)$ generate the unit ideal. If $H_i(v) = \max\{|\sigma_i a|, |\sigma_i c|\}$, then*

$$\|\sigma_0 Z_{12}(v)\|_\infty = H_1(v)H_2(v), \qquad h([Z_{12}(v)]) = 2h([v]). \tag{588}$$

For fixed $P = (p, q)^t$, the fixed linear functional

$$L_P = (\tau(p)Y - \tau(q)X) \otimes (\tau^2(p)Y - \tau^2(q)X)$$

on the tensor space satisfies

$$L_P(Z_{12}(v)) = \tau(\det(P, v))\tau^2(\det(P, v)), \qquad \det(P, v)L_P(Z_{12}(v)) = \mathrm{N}\det(P, v). \tag{589}$$

Proof. The content ideal of the tensor is $\tau((a,c))\tau^2((a,c)) = \mathcal{O}_K$. Its largest coordinate at σ_0 is the product of the largest coordinates of v at σ_1 and σ_2, proving the first equality in (588). At the other two embeddings the logarithmic maxima are $\log H_2 + \log H_0$ and $\log H_0 + \log H_1$. Their sum is twice the sum for v, which proves the height identity. Finally, evaluate the two tensor factors and use $r\tau(r)\tau^2(r) = \mathrm{N}\, r$. □

Proposition A.18 is the precise repair of the failed scalar ordering. It turns the product of the two outer incidences into a fixed linear evaluation on a primitive integral rank-four object. It does not yet give descent: its projective height is exactly twice the original height, so the ordinary Subspace Theorem sees the baseline unit scaling with no surplus. The Hodge input must control the *evaluated* tensor at a synchronized stopping time. This is the rank-four version of the stopped incidence transport problem below.

The pair in (583) is exactly the proper rank-two coordinate

$$(\log|\sigma_1 d| - \log|\sigma_0 d|, \log|\sigma_2 d| - \log|\sigma_0 d|)$$

on the unit lattice. In fact, its signed version determines d up to a rational scalar: if two nonzero elements $d, d' \in K$ have the same two signed ratios, then all three embeddings of d/d' are equal, so $d/d' \in \mathbf{Q}$. Squaring in (581) introduces only finitely many sign choices. Thus one outer representation gives one projection of the unit shape, while the two together retain the complete shape.

This observation also identifies the exact boundary of the method. For two cusps, the primitive parabolic generators canonically remove all unit scaling, and the joint compression functor supplies both shape coordinates. For the determinant $\det(P, v_\xi)$ with a general rational endpoint $[P]$, there is no parabolic generator at $[P]$ unless the conclusion is already true. Proposition A.17 shows that this does *not* prevent a canonical outer upper bound: the two fixed cusp anchors control the two entries of v_ξ separately. What it does prevent is an exact multiplicative compression coordinate for $\det(P, v_\xi)$. Cancellation between pc and qa, and the radial norm paid at successive stopping times, are not recorded by the two upper bounds. The remaining stopped Arakelov problem is to prove that these local payments do not compound forever. It is the same place where the unit frame entered above, now expressed without a Rosen alphabet.

There is a simple model showing that this loss of radial information is essential. Let $\varepsilon = \lambda^2 - 1$. It is a unit whose distinguished absolute value is greater than one and whose other two absolute values are less than one. Both the unit progression ε^n and, for sufficiently large fixed k, the nonunit progression

$$\alpha^n, \qquad \alpha = 2\varepsilon^k, \qquad |\mathrm{N}\,\alpha| = 8,$$

have compact countable sets of pairwise Galois shapes tending to $(0,0)$. Indeed

$$\alpha^n - \alpha^m = \alpha^m(\alpha^{n-m} - 1).$$

The first progression is motion in the unit lattice and is the local model for a special hyperbolic orbit. The second has the same archimedean shape behavior but carries the growing finite degree

$$\log|\mathrm{N}(\alpha^n)| = n \log 8.$$

Consequently no theorem stated only in terms of compression ratios can distinguish the desired special equality case from a dangerous nonunit escape. The full proof must combine McMullen's two shape coordinates with an irreversible bound for this radial Arakelov degree.

Deep input versus elementary identity. The trace calculation (582) is elementary. The continuity, multiplicativity, and thick-part decay of $\mathbf{r}$ use McMullen's global modular-symbol compactification and the two contractive modular embeddings. Their synthesis is substantive: it identifies those global Hodge compressions with the two integral Asai shape coordinates required for arithmetic finiteness.

A.5 A whole-orbit pigeonhole which does not give rigidity

One can use the complete cusp orbit, rather than a preferred coding, to form a tempting additive argument. In a central displacement ball of radius R, Hodge-typical cusp columns have sizes

$$H_0 \le e^{R/2}, \qquad H_1 \le e^{R/10+o(R)}, \qquad H_2 \le e^{3R/10+o(R)}.$$

There are $e^{R-o(R)}$ such columns. For fixed primitive P, their incidences $r(v) = \det(P, v)$ lie in a Minkowski box containing $e^{9R/10+o(R)}$ algebraic integers. Pigeonholing therefore produces one incidence value assumed by $e^{R/10-o(R)}$ cusp columns.

This does *not* produce an exceptional fibre. The quotient

$$\frac{e^{R-o(R)}}{e^{9R/10+o(R)}} = e^{R/10-o(R)}$$

is exactly the mean occupancy predicted by uniform distribution. The incidence value also changes with R. Thus cardinality at this scale does not by itself force a stabilizer; any such conclusion must use an additional structural property of the thin parabolic conjugacy class.

The trace set does not add a hidden restriction. If $U_v = I + \lambda N_v$ is the positive primitive parabolic belonging to a cusp column v, then the elementary rank-one calculation gives

$$\operatorname{Tr}(U_v U_w^{-1}) - 2 = \lambda^2 \det(v, w)^2.$$

If $r(v) = r(w) = r$, write $v = v_0 + tP$ and $w = v_0 + sP$; the formula becomes

$$\operatorname{Tr}(U_v U_w^{-1}) - 2 = \lambda^2 r^2 (t - s)^2.$$

But the first identity holds for every pair of cusps. The second merely rewrites it in fibre coordinates; it is not an independent constraint on the trace set.

This failed route clarifies why the optimal-lifting criterion above keeps the rational flag fixed modulo an externally chosen ideal. The 1/10 margin can annihilate one prescribed incidence. If it is first spent on choosing a popular moving incidence, nothing remains with which to prove that the incidence is zero.

There is nevertheless more group structure in a fibre than the trace identity records. For $P = (p, q)^t$, set

$$N_P = \begin{pmatrix} -pq & p^2 \\ -q^2 & pq \end{pmatrix}, \qquad n_P(s) = I + sN_P.$$

Thus $N_P^2 = 0$, and $n_P(s)$ is the one-parameter unipotent subgroup fixing $[P]$. If $r = \det(P, v) \neq 0$, direct multiplication gives

$$n_P(t/r)v = v + tP, \qquad U_{v+tP} = n_P(t/r) U_v n_P(-t/r). \tag{590}$$

The second identity uses $U_w = I + \lambda N_w$ and $N_{gw} = g N_w g^{-1}$. Consequently a fixed incidence fibre is exactly the intersection of the primitive parabolics of G_7 with one algebraic conjugacy orbit of

n_P. This is a statement about the complete parabolic conjugacy class of the thin group, not about a chosen generator.

There are no further connected algebraic equality groups hidden here. If Q, e are independent and $F_Q(g) = \det(Q, ge)$, then every positive-dimensional connected subgroup on which F_Q is a nonzero constant is unipotent and fixes either e or Q. Indeed, a torus gives a nonconstant expression $Aa + A^{-1}B$, a Borel contains such a torus, and SL_2 moves e through a Zariski-open set. A unipotent with fixed line L sends e to $e + s\ell$; constancy forces either $e \in L$, or $\det(Q, \ell) = 0$, hence $L = KQ$. The first case is the right cusp stabilizer and merely repeats the same column. The second is n_Q. This elementary classification is useful because it shows what a quantitative inverse theorem must produce; it is not itself that inverse theorem.

The two alternatives in the theorem have an equally simple expression in this language. A nonidentity $n_P(s) \in G_7$ makes $[P]$ a cusp. If instead $hP = \varepsilon P$, then

$$hn_P(s)h^{-1} = n_P(\varepsilon^2 s), \tag{591}$$

so a special hyperbolic element normalizes the same unipotent subgroup and acts on its parameter by a nontrivial unit dilation. Conversely, a projectively nonidentity element of G_7 which normalizes n_P fixes its unique boundary point $[P]$. Proposition 7.1 then gives the cusp or special alternative.

The most direct group-level conjecture would say that a fibre has $O_P(e^{(1/10-\delta_P)R})$ points unless G_7 normalizes $n_P(K)$. That would prove the theorem, but it asks for a saving below the mean fibre size and has no supporting mechanism.

A tempting attempt to repair this defect is to tilt the outer Hodge window. Suppose, for small $t > 0$, that

$$H_0(v) \le e^{R/2}, \qquad H_1(v)H_2(v) \le e^{(2/5-t)R},$$

and that pressure gives $e^{(1-I(t)+o(1))R}$ such columns, with $I(t) = O(t^2)$. The ideal-plus-unit count is global and correct: their nonzero incidences occupy at most

$$e^{(9/10-t+o(1))R}$$

values. It therefore gives a fibre of size

$$e^{(1/10+t-I(t)-o(1))R}.$$

This fibre is not exceptional. It is exactly the conditional mean for the tilted shell. Indeed, adelic equidistribution in a shell of incidence volume $e^{(9/10-t)R}$ predicts probability $e^{-(9/10-t)R}$ for one incidence cell; multiplying by the tilted population gives the same exponent $1/10 + t - I(t)$. The linear gain in the pigeonhole count is simply the linear shrinkage of the output space. It supplies no entropy surplus relative to the selected population.

Equivalently, let $\mathfrak{a} = (r)$, and set

$$\epsilon_i \asymp \frac{|\sigma_i r|}{H_i(v)}.$$

The natural adelic Frostman law is

$$\Pr\{\det(P, v) \in \mathfrak{a},\ d_i([v_i], [P_i]) \ll \epsilon_i\} \asymp \frac{\epsilon_0\epsilon_1\epsilon_2}{\mathrm{N}\,\mathfrak{a}}.$$

The product formula cancels the numerator $\prod_i |\sigma_i r|$ against $\mathrm{N}(r)$, leaving $(H_0H_1H_2)^{-1}$. On the tilted shell this is precisely $e^{-(9/10-t)R}$, not $e^{-9R/10}$. Thus a uniform bound by $e^{R/10}$ in the larger untilted window would rule out ordinary large-deviation conditioning and is not a legitimate thin-group conjecture.

The conclusion is that a moving popular incidence value cannot prove the classification, even after a Hodge tilt. A valid argument must keep the target fixed before counting: either the rational residue flag in the optimal-lifting theorem, or one fixed incidence value through infinitely many scales. The unipotent description (590) remains exact, but pigeonholing does not make its intersection exceptional.

There is nevertheless one rigorous whole-block estimate worth retaining. For a collection admitting midpoint factorizations $g = ab$ with both blocks of length $R/2 + o(R)$, the symplectic form gives

$$\det(P, abe) = \det(a^{-1}P, be). \tag{592}$$

Put $X = \{a^{-1}P\}$ and $Y = \{be\}$, and fix $0 \neq r \in \mathcal{O}_K$. Then the fibre is contained, with bounded multiplicity, in the real point–line incidence set

$$\{(x, y) \in X \times Y : \det(x, y) = \sigma_0 r\}.$$

If $[P]$ has trivial stabilizer, the lines indexed by x are distinct. Since $|X|, |Y| \leq e^{R/2+o(R)}$, the Szemerédi–Trotter theorem gives the actual bound

$$\#\{v : \det(P, v) = r\} \leq e^{(2/3+o(1))R}$$

on that midpoint-regular collection. This is a genuine full-group incidence estimate, rather than a generator inequality. It is far too weak to prove the theorem, and it does not turn the moving fibre into a fixed target.

The distinction is now sharp. A joint adelic Frostman theorem at the *shell-dependent* volume exponent would describe ordinary distribution and reproduce the conditional mean above. What proves the classification is instead the horizontal local limit later in this paper: it fixes $[P]$, chooses an external modulus independently of the word, and produces an optimally short lift. No incidence value is selected after the orbit has been counted.

There is a related homogeneous-dynamics analogy which also has to be used with care. Tomanov's split-torus orbit rigidity on $\mathrm{SL}_2(\mathbb{R})^3/\mathrm{SL}_2(\mathcal{O}_K)$ concerns one fixed torus acting on the ambient arithmetic lattice [54]. The lifted Teichmüller curve

$$z \longmapsto (z, F_1(z), F_2(z))$$

is not such an orbit: the two outer coordinates are nonlinear modular embeddings. Each deep cusp segment is approximately diagonal only in its own cusp frame, and successive frames differ by arbitrary elements of the thin group. Putting all segments in one torus would require those frame changes to preserve a rational Borel, which is already the desired cusp or special conclusion. Ambient Hilbert-modular torus rigidity is therefore an analogy for the missing stopped shape–degree theorem, not a proof of it.

A.6 Optimal congruence lifting of a cusp

The preceding formulations still leave open how thinness can produce an arithmetic inequality. There is a particularly direct answer. Reduce the whole cusp orbit modulo an ideal and ask for a short lift of one prescribed residue flag. This uses neither a Rosen digit nor a prime already occurring in a remainder.

Fix $x = [P] \in \mathbf{P}^1(K)$, where $P \in \mathcal{O}_K^2$ is primitive, and retain $e = (1, 0)^t$. For $h \in G_7$, write

$$r_h = \det(P, he), \qquad R(h) = d_{\mathbb{H}}(i, \sigma_0(h)i), \qquad a_i(h) = \log \|\sigma_i(h)\|_{\mathrm{op}}.$$

The congruence $h\infty \equiv x \pmod{\mathfrak{a}}$ means precisely that the two primitive columns he and P are projectively equal modulo $\mathfrak{a}$. In particular,

$$h\infty \equiv x \pmod{\mathfrak{a}} \quad \Longrightarrow \quad r_h \in \mathfrak{a}. \tag{593}$$

Proposition A.19 (A Hodge-balanced short lift is an actual cusp)**.** *Suppose there are integral ideals $\mathfrak{a}_n$, with $\mathrm{N}\,\mathfrak{a}_n \to \infty$, and elements $h_n \in G_7$ such that*

$$h_n\infty \equiv x \pmod{\mathfrak{a}_n}, \tag{594}$$

$$R(h_n) \le A \log \mathrm{N}\,\mathfrak{a}_n + O(1), \tag{595}$$

$$a_1(h_n) + a_2(h_n) \le \theta R(h_n) + O(1). \tag{596}$$

If

$$A\left(\frac{1}{2} + \theta\right) < 1, \tag{597}$$

then $x \in G_7\infty$.

Proof. At the distinguished embedding, the standard comparison between matrix norm and displacement gives

$$\log|\sigma_0 r_{h_n}| \le \frac{1}{2} R(h_n) + O_P(1).$$

At the two other embeddings,

$$\log|\sigma_i r_{h_n}| \le a_i(h_n) + O_P(1) \qquad (i = 1, 2).$$

Consequently

$$\log|\mathrm{N}\, r_{h_n}| \le \left(\frac{1}{2} + \theta\right) R(h_n) + O_P(1) \le A\left(\frac{1}{2} + \theta\right) \log \mathrm{N}\,\mathfrak{a}_n + O_P(1). \tag{598}$$

If $r_{h_n} \neq 0$, divisibility (593) gives the reverse inequality $|\mathrm{N}\, r_{h_n}| \ge \mathrm{N}\,\mathfrak{a}_n$. This contradicts (598) for large n because of (597). Hence $r_{h_n} = 0$. The columns P and $h_n e$ are then parallel over K, so $x = h_n\infty$. □

This elementary closing argument identifies the exact global input which would suffice. The projective congruence orbit has

$$\#\mathbf{P}^1(\mathcal{O}_K/\mathfrak{a}) = (\mathrm{N}\,\mathfrak{a})^{1+o(1)},$$

whereas the number of central group elements of displacement at most R is $e^{R+o(R)}$. Thus the entropy-optimal scale for reaching a prescribed residue flag is

$$R = (1 + o(1)) \log \mathrm{N}\,\mathfrak{a}. \tag{599}$$

The two conjugate Hodge representations have mean total exponent $\theta = 2/5$. At these two values, the left side of (597) is

$$1 \cdot \left(\frac{1}{2} + \frac{2}{5}\right) = \frac{9}{10} < 1.$$

The available margin is therefore exactly $1/10$. No shrinking real target is needed.

The following is the clean thin-group statement suggested by the calculation.

Conjecture A.20 (Hodge-balanced optimal lifting dichotomy)**.** *For every $x \in \mathbf{P}^1(K)$, one of the following holds.*

1. *The stabilizer of x in G_7 is projectively nontrivial.*
2. *For every $\epsilon > 0$, there are ideals $\mathfrak{a}_n$ of unbounded norm and $h_n \in G_7$ satisfying*

$$\begin{aligned} h_n\infty &\equiv x \pmod{\mathfrak{a}_n}, \\ R(h_n) &\le (1+\epsilon)\log \mathrm{N}\,\mathfrak{a}_n, \\ a_1(h_n) + a_2(h_n) &\le (2/5+\epsilon)R(h_n). \end{aligned}$$

Theorem A.21 (Conditional completion by optimal lifting)**.** *Conjecture A.20 implies the heptagonal cusp classification.*

Proof. In the first alternative, Proposition 7.1 gives a cusp or a special point. In the second, choose $\epsilon > 0$ so small that

$$(1+\epsilon)(9/10+\epsilon) < 1.$$

Proposition A.19 then shows that x is in fact a cusp. □

This conjecture is not a disguised generator estimate. Zariski density and strong approximation produce the large projective congruence orbits; expansion concerns the corresponding family of Schreier graphs; and the two Galois representations supply the Hodge constraint. On the other hand, expansion by itself gives only $R \le C \log \mathrm{N}\,\mathfrak{a}$ for an unspecified constant C. The norm argument requires a nearly optimal constant, less than $10/9$, and also requires a Hodge-balanced lift. The mathematical issue is therefore concentration: if all entropy-short lifts of one rational residue flag have atypically large conjugate norm, one must show that the flag has a hyperbolic stabilizer. That is the precise point at which the special alternative enters.

The phrase *optimal lifting* is standard in the arithmetic theory of congruence quotients. Sarnak's formulation, and the general framework of Golubev–Kamber [37], predict lifting at the information-theoretic exponent and relate it to density bounds for nontempered spectrum. There are two important differences here. Their basic optimal-lifting conclusion concerns almost all pairs of congruence classes in arithmetic quotients. Conjecture A.20 concerns the one class obtained by reducing a fixed K-rational flag in a thin orbit. It also imposes a bound for a nonunitary cocycle from the two conjugate Hodge representations. Thus the conjecture above is a prescribed-target, Hodge-weighted version of a standard number-theoretic lifting principle; the periodic alternative records precisely the exception which an almost-all theorem is allowed to discard.

The distinction between "almost all" and the present fixed target is not cosmetic. Kamber–Varjú analyze the projective action of $\mathrm{SL}_2(\mathbf{Z})$ itself [38]. In their normalization the information-theoretic distance on $\mathbf{P}^1(\mathbf{Z}/q\mathbf{Z})$ is $(1/2 + o(1))\log q$, while the worst projective distance has limsup $(1 + o(1))\log q$; their explicit lower bound uses special composite moduli. Since hyperbolic displacement is twice the logarithm of the matrix norm, these become respectively $\log q$ and $2\log q$ in the normalization used here. Our closing inequality requires displacement below $(10/9 - o(1))\log q$. Thus even in the full arithmetic group a theorem about almost every projective target cannot simply be evaluated at one prescribed section. Restricting to split prime ideals removes the known composite hard example, but does not turn an almost-all assertion into a pointwise one.

This comparison identifies the exact role of the thin group. One must prove that the particular targets $[P \bmod \mathfrak{p}]$, as $\mathfrak{p}$ varies, cannot remain in the entropy-short exceptional set unless the rational Borel B_P already meets G_7. The Hodge condition asks in addition that a short lift lie in the typical two-cocycle window. Neither conclusion follows from the cardinality of the exceptional set. A proof needs an inverse theorem for this coherent algebraic section across the congruence tower.

There is a useful amplification which makes the required whole-group estimate more concrete. Instead of one modulus of norm nearly e^R, use a product of two split primes, each of norm below $e^{R/2}$. For a primitive P, put

$$D_g(P) = \det(P, gP),$$

and let $\Omega_R(\delta)$ be the elements satisfying

$$R(g) \le R, \qquad a_1(g) + a_2(g) \le (2/5 + \delta)R.$$

Choose

$$9/10 + \delta < s < 1, \qquad u = s/2 < 1/2, \tag{600}$$

and let $\mathcal{P}_R(u)$ be the good split primes with $e^{uR} \le \mathrm{N}\,\mathfrak{p} \le 2e^{uR}$.

Conjecture A.22 (Two-prime Hodge dispersion)**.** *If $[P]$ has projectively trivial stabilizer, then*

$$\sum_{\substack{\mathfrak{p},\mathfrak{q}\in\mathcal{P}_R(u)\\ \mathfrak{p}\neq\mathfrak{q}}} \#\{g \in \Omega_R(\delta) : g[P] = [P] \pmod{\mathfrak{p}\mathfrak{q}}\}$$
$$\gg \#\Omega_R(\delta) \sum_{\substack{\mathfrak{p},\mathfrak{q}\in\mathcal{P}_R(u)\\ \mathfrak{p}\neq\mathfrak{q}}} \frac{1}{(\mathrm{N}\,\mathfrak{p}+1)(\mathrm{N}\,\mathfrak{q}+1)}. \tag{601}$$

A recognizable sufficient form is a Hodge-weighted Barban–Davenport–Halberstam variance estimate. If

$$E_{\mathfrak{p},\mathfrak{q}}(P;R) = \#\{g \in \Omega_R(\delta) : g[P] = [P] \pmod{\mathfrak{p}\mathfrak{q}}\} - \frac{\#\Omega_R(\delta)}{(Nm\mathfrak{p}+1)(\mathrm{N}\,\mathfrak{q}+1)},$$

assume

$$\sum_{\substack{\mathfrak{p},\mathfrak{q}\in\mathcal{P}_R(u)\\ \mathfrak{p}\neq\mathfrak{q}}} |E_{\mathfrak{p},\mathfrak{q}}(P;R)|^2 \le e^{R+o(R)}. \tag{602}$$

There are

$$e^{sR+o(R)}/R^2$$

prime pairs. Cauchy–Schwarz bounds the sum of their errors by

$$e^{(1+s)R/2+o(R)},$$

whereas the total main term in (601) is $e^{R-o(R)}/R^2$. Since $s < 1$, (602) implies the preceding conjecture. This is a standard variance scale, but its application here is not the classical theorem: it concerns a thin orbital shell, one coherent rational residue section, exponential level, and two nonunitary Hodge weights. The obstruction can be measured exactly. If $[P]$ had no nonidentity stabilizer, (604) would make every count in the definition of $E_{\mathfrak{p},\mathfrak{q}}$ equal to the identity contribution. Since the predicted count per pair is $e^{(1-s-o(1))R}$, the variance would then have size

$$e^{(2-s-o(1))R}, \tag{603}$$

not $e^{R+o(R)}$. The missing inverse theorem must rule out precisely this excess factor $e^{(1-s)R}$, or identify it with concentration in B_P. This is the quantitative wall encountered by all the mixing arguments above.

Proposition A.23 (Two-prime dispersion is sufficient)**.** *Conjecture A.22, together with the full-growth estimate $\#\Omega_R(\delta) = e^{R-o(R)}$, implies the heptagonal cusp classification.*

Proof. At each real place, $|\sigma_i D_g(P)| \ll_P \|\sigma_i(g)\|$. Hence

$$\log|\mathrm{N}\, D_g(P)| \le (9/10+\delta)R + O_P(1) \tag{604}$$

for $g \in \Omega_R(\delta)$. A counted element has $D_g(P) \in \mathfrak{pq}$. If this algebraic integer were nonzero, its norm would be at least $\mathrm{N}(\mathfrak{pq}) \ge e^{sR}$, contrary to (604). Thus every counted element fixes $[P]$ exactly.

The right side of (601) is $e^{R-o(R)}/R^2$. The identity contributes only the number of prime pairs, $e^{sR+o(R)}/R^2$. Since $s<1$, a projectively nonidentity element is counted for large R, contradicting triviality of the stabilizer. Proposition 7.1 finishes the proof. □

This is strictly more structured than asking for an unspecified spectral gap. Each prime action lies below the half-entropy scale, while their product crosses the Hodge norm threshold 9/10. The analytic target is a second factorial moment of prime divisibility of the fixed matrix coefficient $D_g(P)$, averaged over two prime variables. A prospective proof would use a macroscopic factorization of arbitrary elements of the thin group, a bilinear large sieve for the two congruence actions, and the two Galois Cartan weights. What remains nonstandard is the inverse alternative: persistent failure along the coherent reductions of P must force concentration in its rational Borel. This reformulation does not reduce the logical strength of the missing statement: because $\mathrm{N}(\mathfrak{pq})$ is larger than the entire Hodge defect range, positivity of the moment already contains exact closing. Its gain is analytic, not logical: both prime variables now lie below half entropy, the range in which a bilinear dispersion argument can treat them separately before recombination.

For the application one can weaken the conjecture further by varying the modulus. It is enough that, for every rational flag without a stabilizer and infinitely many Q, there be one split prime ideal $\mathfrak{p}$, with $Q \le \mathrm{N}\,\mathfrak{p} \le 2Q$, and one $h \in G_7$ such that

$$h\infty \equiv x \pmod{\mathfrak{p}}, \qquad R(h) \le (1+\epsilon)\log Q, \qquad a_1(h)+a_2(h) \le (2/5+\epsilon)R(h).$$

This is a horizontal optimal-lifting statement: the rational flag is fixed while the finite field varies. The proof is still Proposition A.19. Its advantage is conceptual. A ball of this radius contains $Q^{1+\epsilon+o(1)}$ elements, while each projective congruence orbit has $Q^{1+o(1)}$ points, so there should be $Q^{\epsilon+o(1)}$ admissible lifts per prime. Its difficulty is equally precise: proving this for a prescribed rational section requires distribution essentially up to the entropy limit. Ordinary expander mixing, or an almost-all-target optimal-lifting theorem for one fixed quotient, does not imply it.

Even a conjectural Ramanujan bound does not remove this pointwise obstruction. In the regular nonbacktracking model, if a Schreier graph has N vertices and branching rate β, the Ramanujan estimate for two prescribed vertices has the form

$$K_n(u,v) = \frac{\beta^n}{N} + O(n\beta^{n/2}).$$

Comparison of the main term with the absolute error proves $K_n(u,v) > 0$ only for $n > (2+o(1))\log_\beta N$. The closing proposition needs $n < (10/9+o(1))\log_\beta N$. At $n = (1+\epsilon)\log_\beta N$, an L^2 argument controls almost all targets but can discard the one rational target under consideration. The missing assertion is therefore not "expansion with a better constant." It is a target-sensitive no-holes theorem across the thin congruence tower, weighted by the two outer Cartan cocycles, with a rational Borel orbit as its only exceptional family.

Here is a precise analytic form. For a canonical cusp column $v = (a, c)^t$, set $H_i(v) = \max\{|\sigma_i a|, |\sigma_i c|\}$, and let $\mathcal{C}_R(\delta)$ be the set satisfying

$$H_0(v) \le e^{R/2}, \qquad H_1(v) \le e^{(1/10+\delta)R}, \qquad H_2(v) \le e^{(3/10+\delta)R}.$$

For a good split prime $\mathfrak{p}$, write $X_\mathfrak{p}$ for the projective congruence orbit and

$$K_{\mathfrak{p},R,\delta}(P) = \#\{v \in \mathcal{C}_R(\delta) : [v] = [P] \bmod \mathfrak{p}\}.$$

The natural pointwise local limit prediction is

$$K_{\mathfrak{p},R,\delta}(P) = M_\mathfrak{p} + O\big(M_\mathfrak{p}^{1/2}(\mathrm{N}\,\mathfrak{p})^{o(1)}\big), \qquad M_\mathfrak{p} = \frac{\#\mathcal{C}_R(\delta)}{\#X_\mathfrak{p}}, \tag{605}$$

provided $[P]$ has no projectively nontrivial stabilizer. It is enough to assume the corresponding asymptotic after summing over split primes in one dyadic interval. At $R = (1+\eta)\log \mathrm{N}\,\mathfrak{p}$, the Hodge volume law gives $M_\mathfrak{p} = (\mathrm{N}\,\mathfrak{p})^{\eta - o(1)}$, so the main term in (605) dominates its error.

Theorem A.24 (Conditional completion by a Lindelöf-sized local limit)**.** *Assume the Hodge volume law $\#\mathcal{C}_R(\delta) = e^{R-o(R)}$ and either the pointwise estimate (605) or its prime-averaged asymptotic, for every rational flag with trivial projective stabilizer. Then the heptagonal cusp classification holds.*

Proof. Choose $\eta, \delta > 0$ with $(1+\eta)(9/10 + 2\delta) < 1$. The local limit gives a prime $\mathfrak{p}$ and a column $v \in \mathcal{C}_R(\delta)$ congruent to P. Thus $r = \det(P, v) \in \mathfrak{p}$, whereas

$$\log|\mathrm{N}\,r| \le (9/10 + 2\delta)R + O_P(1) < \log \mathrm{N}\,\mathfrak{p}.$$

Therefore $r = 0$, making $[P]$ an actual cusp with a parabolic stabilizer. This contradicts triviality. Every rational flag therefore has a nontrivial stabilizer, and Proposition 7.1 gives the cusp or special alternatives. □

This conditional implication is exact, but it is not yet an independent thin-group mechanism. If $\omega_Q(\mathfrak{b})$ counts good split prime divisors of $\mathfrak{b}$ with norm in $[Q, 2Q]$, then

$$\sum_{\substack{Q \le \mathrm{N}\,\mathfrak{p} \le 2Q \\ \mathfrak{p} \text{ good split}}} K_{\mathfrak{p},R,\delta}(P) = \sum_{v \in \mathcal{C}_R(\delta)} \omega_Q\big((\det(P, v))\big), \tag{606}$$

with a zero incidence counted for every prime on the left. Under the margin used in the proof, every nonzero incidence has norm less than Q, and hence contributes zero to the right side. Positivity of the prime average is therefore already equivalent to an actual zero incidence. Prime averaging has not softened the closing problem.

This is an important limit on the interpretation of (605). The prescribed congruence class lies outside the complete norm support of every nonzero incidence. A pointwise theorem saying that the orbit hits it has already supplied the desired exact vanishing; it is not a routine extrapolation of equidistribution. The local-limit theorem remains a concise conditional reduction, but it should be read as a no-holes form of the classification, not as the proof of that no-holes property.

The size of the error is the substantive conjecture. It is square-root in the *local expected count* $M_\mathfrak{p}$. Ramanujan gives an error of square-root size in the *whole ball*, which is the factor-two barrier above. Calling the desired estimate "Lindelöf-sized" describes its scale; it does not mean that it is a formal consequence of a standard automorphic Lindelöf conjecture.

There is also a structural point about the Hodge weights. They are not scalar potentials on the Bowen–Series shift. For an outer embedding σ_i, the increment is the matrix-cocycle quantity

$$\psi_i(\gamma, [u]) = \log \frac{\|\sigma_i(\gamma)u\|}{\|u\|}.$$

It becomes additive only after the state space is enlarged by the moving outer direction $[u] \in \mathbf{P}^1(\mathbb{R})$. The required transfer operator must therefore act on the accelerated symbolic coordinate, two outer projective directions, and the finite congruence orbit at once. The acceleration has unbounded cusp branches. These are real additional features, not bookkeeping details.

This formulation explains exactly where the thin group is used. Its complete nonbacktracking orbit supplies exponentially many candidates; strong approximation supplies the large finite projective quotients; and the two non-geometric representations supply the Hodge balance which creates the 1/10 norm margin. None of these three facts is visible in a calculation with a single generator. A rational Borel stabilizer certainly gives a closed loop with exceptional cocycle holonomy. The converse—every persistent target-specific failure of the local limit comes from such a loop—is a separate spectral-rigidity assertion, and is part of the conditional input rather than a proved consequence of the coding.

Published transfer estimates do not already contain this assertion. Mayer–Mühlenbruch–Strömberg construct transfer operators for cofinite Hecke triangle groups [50], but without the congruence-uniform, two-cocycle local limit above. Uniform congruence mixing of Oh–Winter is convex cocompact [52], hence does not include this cusp. The Schottky congruence count of Magee–Oh–Winter has a polynomial loss in the level and a fixed power saving in radius [53]; it does not reach congruence level $\mathrm{N}\,\mathfrak{p} \asymp e^R$ with error square-root in the local expected count. At that entropy-size level, however, the preceding divisor identity shows that the proposed theorem already contains exact closing. A less taut analytic formulation is the joint congruence/real shrinking-target problem below: its two restrictions are imposed below their separate closing scales, their total entropy loss lies in the available 1/10 window, and persistent failure along the coherent rational section must produce a rational Borel.

A.7 The relevant thin-orbit method: a pointed local–global theorem

There is an established analytic architecture which is closer to this joint problem than ordinary expander mixing. Bourgain and Kontorovich study integral values of a scalar matrix coefficient on a thin SL_2-orbit by cutting a long group element into macroscopic blocks and applying the circle method, bilinear estimates, and dispersion [68, 69]. Zhang treats the corresponding local–global problem for thin groups containing parabolics [70]. The point of these methods is precisely that they use convolution of large pieces of the group to go beyond the level of distribution supplied by a spectral gap.

The same algebra is present here. For

$$D_g(P) = \det(P, gP)$$

and a block factorization $g = ab$, invariance of the symplectic form gives

$$D_{ab}(P) = \det(a^{-1}P, bP). \tag{607}$$

Thus a congruence character of $D_g(P)$ becomes a bilinear exponential sum in two complete macroscopic group blocks. A smooth real shrinking-target cutoff can be Fourier expanded at the

distinguished embedding. The two outer Hodge restrictions are imposed on the same ensemble by retaining the moving projective directions of the matrix cocycles. None of these operations singles out a Rosen generator.

The complete finite quotient has an unusually clean scalar distribution. Let $k = \mathbb{F}_q$, let $0 \neq P \in k^2$, and let ψ be a nontrivial additive character of k. Then

$$\#\{g \in \mathrm{SL}_2(k) : D_g(P) = 0\} = q^2 - q, \tag{608}$$

$$\#\{g \in \mathrm{SL}_2(k) : D_g(P) = t\} = q^2 \quad (t \neq 0), \tag{609}$$

$$\sum_{g \in \mathrm{SL}_2(k)} \psi(\xi D_g(P)) = -q \quad (\xi \neq 0). \tag{610}$$

Indeed, conjugate P to e_1. If $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, then $D_g(e_1) = c$. For $c \neq 0$, the choices of a, d determine $b = (ad-1)/c$, giving q^2 matrices. For $c = 0$, there are $q-1$ choices for a, then $d = a^{-1}$, and q choices for b. Fourier summation gives the last identity.

There is also no hidden prime-power obstruction. Let

$$R_m = \mathfrak{o}/\varpi^m\mathfrak{o}, \qquad q = \#(\mathfrak{o}/\varpi), \qquad Q = q^m,$$

and let $P \in R_m^2$ be unimodular. The same first-column count gives

$$\#\{g \in \mathrm{SL}_2(R_m) : D_g(P) = t\} = \begin{cases} Q^2, & t \in R_m^\times, \\ Q^2(1 - q^{-1}), & t \notin R_m^\times. \end{cases} \tag{611}$$

Indeed, after sending P to e_1, fix the lower-left entry $c = t$. For each primitive first column $(a, c)^t$, the possible second columns form an affine copy of R_m. When c is a unit there are Q choices for a; when it is a nonunit, a must be a unit.

Consequently, for an additive character ψ of R_m,

$$\sum_{g \in \mathrm{SL}_2(R_m)} \psi(D_g(P)) = \begin{cases} Q^3(1 - q^{-2}), & \psi = 1, \\ -Q^3 q^{-2}, & \psi \neq 1 \text{ and } \psi|_{\varpi R_m} = 1, \\ 0, & \psi|_{\varpi R_m} \neq 1. \end{cases} \tag{612}$$

To see this, sum (611) over c. If $T = \sum_{R_m} \psi$ and $U = \sum_{\varpi R_m} \psi$, the result is $Q^2(T - q^{-1}U)$.

The Chinese remainder theorem now gives, for every good composite ideal $\mathfrak{a}$,

$$\frac{\#\{g \in \mathrm{SL}_2(\mathcal{O}/\mathfrak{a}) : D_g(P) = 0\}}{\#\,\mathrm{SL}_2(\mathcal{O}/\mathfrak{a})} = \frac{1}{\#\mathbf{P}^1(\mathcal{O}/\mathfrak{a})}. \tag{613}$$

Moreover every Fourier character with conductor genuinely exceeding the residue field at some prime has zero complete local coefficient; the only nonzero local modes factor through a residue field and have normalized coefficient $-1/(q^2-1)$. Hence prime powers introduce no additional major-arc resonance.

The same calculation has a more pointed interpretation for the three-block argument below. If $u, v \in R_m^2$ are unimodular, then the distribution of

$$c \longmapsto \det(u, cv), \qquad c \in \mathrm{SL}_2(R_m), \tag{614}$$

is exactly the distribution in (611). Indeed, choose $h, k \in \mathrm{SL}_2(R_m)$ with $hu = kv = e_1$. The change of variable $c' = hck^{-1}$ is a bijection and

$$\det(u, cv) = \det(e_1, c'e_1).$$

Thus a complete middle bridge erases every dependence on the two outer endpoints at the finite places. At a residue field its nontrivial normalized Fourier coefficient is exactly $-1/(q^2-1)$, and at higher conductor it vanishes. The difficulty is therefore to recover enough of this complete-quotient cancellation from a short thin-group block while simultaneously retaining the real and Hodge cutoffs.

There is also an exact description of an incomplete bridge. On $V = \mathbb{F}_q^2$ put

$$(\mathcal{F}_\xi f)(u) = q^{-1} \sum_{v \in V} \psi(\xi \det(u, v)) f(v)$$

and let π be the permutation representation of $\mathrm{SL}_2(\mathbb{F}_q)$ on V. For a probability measure ν on the finite group, the bridge operator is

$$\begin{aligned} (\mathcal{T}_{\nu,\xi} f)(u) &= q^{-1} \sum_c \nu(c) \sum_v \psi(\xi \det(u, cv)) f(v) \\ &= \mathcal{F}_\xi \pi(\nu) f. \end{aligned} \tag{615}$$

The first factor is unitary by additive orthogonality. Hence, on the mean-zero functions on $V \setminus \{0\}$,

$$\|\mathcal{T}_{\nu,\xi}\| = \|\pi(\nu)\|.$$

This identity locates the analytic bottleneck. The favorable complete coefficient $O(q^{-2})$ is the result of full projection onto the invariant functions; a short thin block has exactly its congruence spectral norm and no additional automatic Fourier gain.

After removal of the diagonal, the complete bridge has more cancellation than this operator statement displays. For nonzero $u, u', v, v' \in \mathbb{F}_q^2$, put

$$K = vu^t J - v'u'^t J, \qquad J = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

Then

$$\det K = -\det(u, u') \det(v, v'). \tag{616}$$

Additive Fourier inversion for $\det c = 1$ gives

$$\sum_{c \in \mathrm{SL}_2(\mathbb{F}_q)} \psi\big(\det(u, cv) - \det(u', cv')\big) = q \,\mathrm{Kl}_\psi(\det K) \tag{617}$$

when $K \neq 0$. It equals $-q$ when K has rank one, and has absolute value at most $2q^{3/2}$ when K is invertible. For $K = 0$ it is $q(q^2-1)$.

Thus the singular dual variety is exactly

$$\det(u, u') \det(v, v') = 0.$$

Under the assumption that $[P]$ has trivial stabilizer, a projective coincidence between two left endpoints or between two right endpoints is possible only on the literal diagonal of the corresponding block. Once the diagonal and two semi-diagonals are subtracted, every complete finite correlation is in the Weil-bound case. At the distinguished real place, stationary phase for $\int_{\mathrm{SL}_2(\mathbb{R})} e^{i\xi \operatorname{tr}(Kc)} w(c)\, dc$ has the same dual cone: away from $\det K = 0$, its nondegenerate three-dimensional decay is $O_w(|\xi|^{-3/2})$.

This is an exact piece of the desired inverse theorem. The bad finite and real correlations have the same algebraic support, and that support is endpoint repetition. The remaining step is no longer to guess the structured obstruction; it is to transfer the complete Kloosterman and real stationary-phase cancellation to a short Hodge-weighted thin bridge by means of the two long blocks.

The singular cone also disappears from a horizontal prime average for a genuinely off-diagonal pair. Indeed, for four block elements put

$$Z_{a,a';b,b'}(P) = \det(a^{-1}P, a'^{-1}P)\det(bP, b'P) \in \mathcal{O}_K. \tag{618}$$

If $[P]$ has trivial projective stabilizer and $a \neq a'$, $b \neq b'$, then $Z_{a,a';b,b'}(P) \neq 0$. If all four blocks have displacement at most R, the modular embeddings give

$$\log|\,\mathrm{N}\, Z_{a,a';b,b'}(P)| \leq C_P R + O_P(1). \tag{619}$$

It follows that, for every fixed $r > 0$, only $O_P(1/r)$ prime ideals of norm at least e^{rR} can divide Z. The dyadic interval $[e^{rR}, 2e^{rR}]$ contains $e^{rR+o(R)}/R$ good split primes. Therefore an exponentially vanishing proportion of those primes sees an off-diagonal correlation on the singular cone; at all the others the complete correlation is in the nonsingular Weil case.

This is a whole-orbit statement, not a generator estimate. All finite places are controlled simultaneously by the norm of the one global discriminant (618). It removes the finite algebraic inverse obstruction. The remaining difficulty is analytic: the middle block is far shorter than a complete congruence quotient, so the two long blocks must turn this almost-every-prime complete cancellation into an incomplete bilinear estimate.

One can quantify exactly what the long blocks already accomplish. If $\mathcal{A}$ is a set of A distinct block elements and

$$\alpha_{\mathfrak{p}} = A^{-1}\sum_{a\in\mathcal{A}} \delta_{[a^{-1}P \bmod \mathfrak{p}]},$$

then triviality of the stabilizer and (619) give

$$\frac{1}{\#\mathcal{P}}\sum_{\mathfrak{p}\in\mathcal{P}} \|\alpha_{\mathfrak{p}}\|_2^2 \leq \frac{1}{A} + \frac{O_P(1)}{\#\mathcal{P}} \tag{620}$$

for a dyadic family $\mathcal{P}$ of exponential-size split primes. Indeed, the off-diagonal part of the square counts primes dividing $\det(a^{-1}P, a'^{-1}P)$, and every nonzero such determinant has only boundedly many prime divisors in the interval. Thus a long thin block has the optimal endpoint collision exponent on average over primes.

There is a sharper representation-theoretic description of the finite variable. At every good split prime, strong approximation and double transitivity give

$$\ell^2\big(\mathbf{P}^1(\mathbf{F}_q)\big) = \mathbf{1} \oplus \mathrm{St}_q, \tag{621}$$

where St_q is the q-dimensional Steinberg representation. Put $e_y = \delta_y - \mathbf{1}/(q+1)$. For every finite shell $\Omega \subset G_7$, normalized counting measure μ_Ω, and primitive P, one therefore has the exact formula

$$\frac{\#\{g\in\Omega : g\infty = [P \bmod \mathfrak{p}]\}}{\#\Omega} = \frac{1}{q+1} + \left\langle \mathrm{St}_q(\mu_\Omega)e_\infty, e_{[P \bmod \mathfrak{p}]}\right\rangle. \tag{622}$$

Thus the prescribed-target problem contains exactly one nonconstant finite representation, not an unspecified family of expander modes.

For a Hodge-balanced shell $\Omega_R(\delta)$ of size $e^{R-o(R)}$ and $q = e^{sR+o(R)}$, a relative square-root estimate would read

$$\#\{g \in \Omega_R(\delta) : g\infty = [P] \bmod \mathfrak{p}\} = \frac{\#\Omega_R(\delta)}{q+1} + O_\epsilon\Big((\#\Omega_R(\delta)/q)^{1/2} q^\epsilon\Big). \tag{623}$$

It is sufficient whenever $s < 1$, and in normalized form gives a Steinberg coefficient $o(q^{-1})$. On the other hand, if $s > 9/10 + \delta$ and $[P]$ has trivial stabilizer, the product formula forbids every hit in the left side: a hit would make $\det(P, ge) = 0$. Formula (622) then forces the extremal value $-1/(q+1)$. Hence the exact inverse theorem needed from the congruence tower is:

> Persistent extremal negative Steinberg coefficient for the coherent vectors $e_{[P \bmod \mathfrak{p}]}$, after the two Hodge cutoffs, forces $G_7 \cap B_P(K) \neq \{1\}$.

This uses the complete thin group in four places: its geodesic shells, its strong-approximation quotients, the Steinberg representation of its projective action, and its two correlated modular-embedding cocycles. The candidate exceptional subgroup is already visible: $e_{[P]}$ is Borel-fixed, and the characteristic-zero Borel is exactly the stabilizer whose nontrivial intersection gives the cusp or special alternative.

The corresponding relative Lindelöf hypothesis has an exact normalization. Form the accelerated transfer operator on the symbolic coordinate, the two moving outer projective directions, and St_q, with complex variables z, v_1, v_2 dual to central time and the two Hodge cocycles. If Φ_∞, Φ_P have finite components e_∞, e_P, the pointed counting series is

$$\big\langle (I - \mathscr{L}_{z,\mathbf{v},q})^{-1} \Phi_\infty, \Phi_P \big\rangle .$$

The existence of fast and slow Hecke-triangle transfer operators with an arbitrary finite-dimensional unitary twist is established by Adam–Pohl [51]; the new demand here is uniformity for the growing Steinberg twist and the two additional Hodge variables. These Hodge variables are genuine representation parameters. The cocycle identity for

$$\psi_i(g, [u]) = \log \frac{\|\sigma_i(g)u\|}{\|u\|}$$

turns the weighted projective action into a principal-series representation Π_{i,v_i}. After the usual half-density normalization, the group part of the operator acts in

$$\mathrm{St}_q \otimes \Pi_{0,z} \otimes \Pi_{1,v_1} \otimes \Pi_{2,v_2}.$$

The common algebraic stabilizer of the Galois endpoint

$$([\sigma_0 P], [\sigma_1 P], [\sigma_2 P])$$

is $\mathrm{Res}_{K/\mathbb{Q}} B_P$, and hence its intersection with G_7 is exactly $G_7 \cap B_P(K)$. The full SL_2^3 Galois Zariski closure excludes a hidden diagonal factor between two embeddings. Thus the proposed exceptional term is forced by the tensor representation itself.

The sufficient critical-line estimate is

$$\big\langle (I - \mathscr{L}_{z,\mathbf{v},q})^{-1} \Phi_\infty, \Phi_P \big\rangle - \mathcal{E}_{P,q}(z, \mathbf{v}) \ll_{P,\epsilon,A} q^{-1/2+\epsilon} (1 + |\Im z|)^{-A}, \qquad \Re z = \frac{1}{2} + \epsilon, \tag{624}$$

uniformly in the Fourier range imposing the Hodge-balanced shell. Here $\mathcal{E}_{P,q}$ is the explicit sum of cusp or primitive hyperbolic polar terms indexed by nonidentity elements of $G_7 \cap B_P(K)$, and is zero when that stabilizer is trivial. Smooth inversion gives

$$q^{-1/2+\epsilon} e^{(1/2+\epsilon)R} = (e^R/q)^{1/2} q^{O(\epsilon)},$$

which is exactly the error in (623). Thus (624) conditionally proves the classification: either the remainder produces a cusp hit, or $\mathcal{E}_{P,q} \neq 0$ already gives the rational-Borel stabilizer.

The factor $q^{-1/2}$ is decisive. Published uniform congruence transfer estimates have schematic error $q^C e^{(1-\eta)R}$. Even a conjectural Ramanujan operator norm gives $e^{R/2+o(R)}$, still too large at $q = e^{sR}$, $s > 9/10$. The additional $q^{-1/2}$ is the local Borel-period factor for the two Steinberg newvectors. Averaging over all projective targets can recover it by Plancherel; the theorem requires it for the coherent target $[P \bmod \mathfrak{p}]$, after the actual B_P-orbital terms have been removed.

The local $q^{-1/2}$ factor itself can be proved. Realize $\mathrm{St}_q = \ell_0^2(\mathbf{P}^1(\mathbf{F}_q))$, let ϕ_0 be the normalized centered delta function at ∞, and for $\xi \neq 0$ put

$$\phi_\xi(x) = q^{-1/2}\psi(\xi x) \quad (x \in \mathbf{F}_q), \qquad \phi_\xi(\infty) = 0.$$

This is an orthonormal basis. The complete cusp unipotent is diagonal:

$$u_t \phi_\xi = \psi(-\xi t)\phi_\xi, \qquad u_t \phi_0 = \phi_0.$$

The Weyl element $w : x \mapsto -1/x$, which changes between opposite cusp frames, has matrix

$$\langle w\phi_\xi, \phi_\eta \rangle = \frac{1}{q} \sum_{x \neq 0} \psi(-\xi/x - \eta x),$$

and hence every entry is at most $2q^{-1/2}$ by the Weil bound. Rows involving ϕ_0 have entries $-1/q$ or $\sqrt{q+1}/q$. Thus the square-root local factor is an unconditional consequence of the Bruhat structure of the full finite quotient. For arbitrary distinct cusp frames the transition lies in $B_1 w B_2$; the Borel factors add phases and permute Fourier frequencies, so the same flatness bound holds.

What remains is not this complete Kloosterman estimate. Accelerated cusp blocks insert incomplete diagonal multipliers

$$D_\nu(\xi) = \sum_n \nu(n)\psi(-\xi n\lambda),$$

where the same measure ν is cut out by central time and both Hodge cocycles. A product of flat Weyl matrices and correlated diagonal matrices can still have norm one. The unique line fixed by the complete cusp unipotent is $\mathbf{C}\phi_0$, the Borel-fixed line. Therefore the required inverse theorem can be stated concretely: persistent failure of Kloosterman flatness through the stopped product must concentrate on successive Borel-fixed lines, and the global height must turn that concentration into an exact repeated rational Borel.

This still does not close by ordinary L^2 composition. At the one-prime scale write $q = e^{sR}$, $9/10 < s < 1$, and split the word into blocks of entropy $s/2, 1-s, s/2$. The two endpoint sets then have size $q^{1/2+o(1)}$, while the bridge has size $q^{(1-s)/s+o(1)}$. Formula (620) and Cauchy–Schwarz would require the bridge spectral norm to be at most $e^{-sR/2}$. Its entire length is less than $R/10$, so even a formal entropy-rate contraction cannot supply that bound. Equivalently, the missing finite statement is the lower-incidence estimate

$$\sum_{c \in C} |U \cap cV| \asymp |C|, \qquad |U| = |V| = q^{1/2+o(1)},$$

for these three correlated thin-orbit blocks. General finite-field Möbius-incidence theorems give upper-tail estimates; they do not force this critical lower incidence. This isolates a genuinely off-diagonal whole-group variance theorem as the required next input.

More generally, for a unitary boundary representation ρ, the joint bridge operator factors as

$$\mathcal{T}_{\nu,\xi,\rho} = (\mathcal{F}_\xi \otimes I) \sum_c \nu(c)\big(\pi(c) \otimes \rho(c)\big). \tag{625}$$

Hence its norm is exactly the norm of the diagonal group average on the tensor product. Fourier decomposition of the real shrinking target provides the boundary factor, while Laplace transforms of the Hodge cutoffs attach the two outer projective cocycles. Separate congruence and real spectral gaps do not bound this tensor norm: a correlated graph can be invisible in both marginals. The needed estimate is therefore a matrix-coefficient or energy estimate for this diagonal action of the whole Galois representation.

Strong approximation transfers this calculation to every good split prime quotient of G_7. It proves that there is no finite local bias against the rational Borel: its density is exactly $1/(q+1)$, and the nontrivial scalar Fourier coefficients have normalized size $O(q^{-2})$. Thus any persistent obstruction is created by coupling the finite variable to the real and Hodge restrictions, not by a defect in the finite quotient itself.

Write

$$r = \frac{\log \#X_{\mathfrak{p}}}{R}$$

for the congruence entropy and let β be the real shrinking rate. The expected number of elements per prime is

$$\exp\big((1 - r - \beta + o(1))R\big). \tag{626}$$

The product-formula closing inequality requires

$$r + \beta > \frac{9}{10} + 2\delta, \tag{627}$$

while a positive expected count requires $r + \beta < 1$. The allocation between congruence and real approximation can vary, but the total budget is always the same interval of width $1/10 - 2\delta$.

This suggests the following precise full-group input. For a rational flag $[P]$, average over good split primes satisfying $\log \mathrm{N}\,\mathfrak{p} = rR + O(1)$, and count elements in a central shell of radius R which fix $[P]$ modulo $\mathfrak{p}$, move $[P]$ by at most $e^{-\beta R}$ at the distinguished real place, and have outer Cartan sum at most $(2/5 + \delta)R$. If $[P]$ has trivial projective stabilizer, the desired pointed local–global theorem gives the lower bound (626) for the average over primes whenever $r + \beta < 1$. Choosing $9/10 + 2\delta < r + \beta < 1$ and applying the closing proposition below then gives a contradiction.

The published local–global theorems do not imply this assertion. They represent a density-one set of admissible scalar values. Here the value zero is fixed in advance and is singular: it means that a rational line has been preserved. Moreover a special hyperbolic line is congruence admissible at every level but is not a cusp, so congruence admissibility alone cannot identify the exceptional set. The required upgrade is not "all admissible values are represented." It is an inverse theorem: failure of the pointed bilinear estimate along one coherent rational section forces the block measures into its rational Borel. The exact Borel calculation then gives the cusp or special alternative.

> **What is genuinely new in this route.** The bilinear identity (607) is elementary. The deep proposed input is a dispersion/minor-arc theorem uniform in the two projective Hodge cocycles, followed by a classification of a persistent major-arc resonance as a K-rational Borel. Existing thin-orbit local–global work supplies the correct whole-group factorization and shows how one can beat the spectral square-root barrier; it does not supply this pointed exceptional-set classification.

A.8 Why a black-box spectral gap still does not use enough of the group

The preceding proposal should not be confused with the assertion that some unspecified spectral gap will suffice. The loss can be read off before choosing an analytic method.

Proposition A.25 (The decay budget)**.** *Let Ω_R be a family of $e^{R+o(R)}$ Hodge-balanced elements, and suppose an event has the local density*

$$m_R = \exp\bigl(-(r+\beta+o(1))R\bigr).$$

Assume that a normalized counting theorem gives only

$$\left|\frac{\#\{g\in\Omega_R : g \text{ satisfies the event}\}}{\#\Omega_R} - m_R\right| \le \exp\bigl(-(\kappa-o(1))R\bigr). \tag{628}$$

Then this estimate forces a positive count only when

$$\kappa > r+\beta. \tag{629}$$

In the closing window $r+\beta > 9/10+2\delta$, it therefore requires $\kappa > 9/10+2\delta$.

Proof. The main term in normalized form is $e^{-(r+\beta+o(1))R}$. The error in (628) is smaller than this main term exactly under (629). The last assertion is (627). □

This proposition explains why dividing a word into more pieces does not, by itself, solve the problem. If the only input for their convolution is a fixed spectral radius, the final error is still $e^{-\kappa R}$ with the same decay per unit hyperbolic length. The required exponent is not merely positive; it is almost the full orbital entropy.

There is nevertheless a genuine full-group theorem available on the archimedean side. Jialun Li proves Fourier decay, a uniform high-frequency spectral gap for the Iwasawa cocycle, and an exponential renewal error for a finite-exponential-moment random walk whose support is Zariski dense in a connected split semisimple real group [45]. Proposition 7.2 makes his theorem applicable to any finitely supported generating measure on $\rho_{\mathrm{Gal}}(G_7) \subset \mathrm{SL}_2(\mathbb{R})^3$. This is an actual use of the whole diagonal thin group: the conclusion would degenerate if the three factors lay in a proper algebraic subproduct.

What that theorem supplies is the archimedean minor-arc mechanism for the three flag directions and the two Hodge cocycles. What it does not supply is a congruence twist whose level grows like e^{rR}, a Bowen–Margulis shell conditioned on one prescribed rational cylinder, or a decay exponent satisfying (629). Its positive exponent is measure-dependent and is not asserted to be close to one. It is therefore a component of a proof, not the missing proof.

The part of thin-orbit theory which can cross this budget is dispersion: one keeps the two macroscopic variables in (607) separate and estimates their additive energy, rather than replacing the entire group ball by one mixing operator. The remaining inverse statement is correspondingly precise. If dispersion fails along the reductions of one fixed K-rational line, the concentration must lie in the rational Borel which fixes that line. Establishing this implication, while retaining the Hodge cutoffs, is the point at which the multiplication law of the whole thin group must do work.

There is a particularly useful way to divide the available entropy which makes this last sentence concrete. Choose

$$\frac{9}{10}+2\delta < s < 1, \qquad r=\beta=\frac{s}{2}, \tag{630}$$

and write $g = ab$, with each block having central radius $R/2 + O(1)$. Each block contains $e^{R/2+o(R)}$ elements. The congruence orbit has $e^{rR+o(R)}$ points and the real line has $e^{\beta R+o(R)}$ intervals of radius $e^{-\beta R}$. Since

$$r < \frac{1}{2}, \qquad \beta < \frac{1}{2},$$

each of these two marginal partitions is smaller than one block. Their joint partition, however, has $e^{sR+o(R)}$ cells, more than one block but fewer than the $e^{R+o(R)}$ ordered pairs (a, b). Independent block distributions therefore predict

$$e^{(1-s+o(1))R} \tag{631}$$

pairs for which $[a^{-1}P]$ and $[bP]$ lie in the same congruence class and the same real interval.

The block languages here are chosen from compatible one-sided nonbacktracking cylinders. Thus every concatenation ab is a reduced positive word of radius $R + O(1)$. This removes the spurious pairs $b = a^{-1}$, which would produce many copies of the identity but no projectively nontrivial stabilizer.

Such a pair gives

$$D_{ab}(P) = \det(a^{-1}P, bP) \in \mathfrak{p}$$

and the real interval supplies the factor $e^{-\beta R}$. If both blocks are Hodge balanced, their outer costs add to the required $(2/5 + \delta)R$. Equations (630) and (627) then turn the collision into exact vanishing.

This symmetric split isolates the remaining difficulty more sharply than a pointwise mixing statement. A Ramanujan-sized L^2 congruence estimate is now asked to flatten either block only below the half-entropy scale, and the analogous boundary estimate is asked to do the same in the real marginal. Separate flattening does not imply a joint collision: two sparse subsets of the product can have uniform marginals and still be disjoint, like two different graphs of permutations. The needed inverse-dispersion theorem must rule out precisely this congruence–real graph correlation. In the present problem its algebraic persistent form is concentration in the rational Borel of $[P]$, which is already one of the desired outcomes.

A short middle block makes the proposed inverse mechanism still more explicit. Choose $0 < \tau < 1 - s$ and factor

$$g = acb, \qquad R(a) = R(b) = \frac{1-\tau}{2}R + O(1), \qquad R(c) = \tau R + O(1). \tag{632}$$

Then

$$D_{acb}(P) = \det(a^{-1}P, cbP). \tag{633}$$

The two outside blocks still have more entropy than either marginal, because $s/2 < (1-\tau)/2$. The middle family contributes $e^{\tau R+o(R)}$ different translates of the second sparse joint support without changing the total predicted collision count $e^{(1-s+o(1))R}$.

The bridge is an inverse-theorem device. Equation (614) proves that the complete finite group destroys a graph correlation uniformly in its two endpoints. A congruence–real graph which prevents collisions for every middle block must remain coherent under exponentially many elements of the thin group. An adelic product/flattening theorem should turn that persistence into concentration near a proper algebraic stabilizer. For the incidence (633), the target-dependent stabilizer is the rational Borel B_P. More precisely, the required inverse alternative must produce either a compatible approximate incidence to which the product formula applies, or a projectively nonidentity element of $G_7 \cap B_P$ directly. Merely saying that a measure is metrically concentrated near B_P would not be enough. This is the precise three-block inverse-dispersion statement now required.

There is an exact reason why the bridge cannot be handled by a black-box spectral gap. The joint congruence–real partition has $M_R = e^{(s+o(1))R}$ cells, while either outside block has only

$$N_R = e^{((1-\tau)/2+o(1))R}$$

points. If f is its endpoint density relative to uniform measure on the joint cells, then

$$\|f-1\|_2^2 \geq \frac{M_R}{N_R} - 1 = e^{(s-(1-\tau)/2+o(1))R}. \tag{634}$$

This lower bound already assumes the most favorable possible placement, with essentially no repeated joint cells. A bridge norm $e^{-\kappa\tau R}$, inserted into the usual Cauchy–Schwarz correlation estimate, can dominate these endpoint norms only if

$$\kappa\tau > s - \frac{1-\tau}{2}. \tag{635}$$

Since $s > 9/10$ and $\tau < 1-s$, the required contraction rate is greater than four per unit bridge length; a Ramanujan-sized rate is $1/2$.

Equivalently, the self-correlation of one outside block has N_R diagonal pairs, while the desired number of distinct birthday collisions is only

$$\frac{N_R^2}{M_R} = e^{(1-\tau-s+o(1))R}.$$

The diagonal is larger by the factor in (634), and every such pair produces the identity. Hence positivity, a squared bridge measure, or an ordinary second moment does not prove closing. The analytic input must be a genuinely off-diagonal trilinear estimate for the two long blocks and the bridge. This is a narrower and more demanding statement than the earlier phrase "joint spectral gap."

The bridge cannot do this pointwise by itself. Even granting the Ramanujan-sized estimate

$$\left\| \pi(\nu_{\mathcal{C}_R})\big|_{\ell_0^2} \right\| \leq e^{-(\tau/2)R+o(R)},$$

hitting one projective class of mass e^{-rR} from a point would require $\tau/2 > r$. But $\tau < 1-s$ and $r = s/2 > 9/20$. Thus the required gain must be genuinely three-linear: the two long outer blocks must flatten the endpoint data before the short congruence operator is applied, while also averaging the real frequency.

One qualitative point is already forced by the central lattice. Let X be any finite transitive congruence orbit of G_7. A nonempty closed G_7-invariant subset of

$$X \times \mathbf{P}^1(\mathbb{R})$$

is the whole product. Indeed, the stabilizer H of a point $x \in X$ has finite index in G_7, hence is a cofinite non-elementary Fuchsian group. Its limit set is the whole boundary circle and its boundary action is minimal. Invariance and transitivity make the projection of the closed set all of X, so the fibre is nonempty. That fibre is a nonempty closed H-invariant set, so it is the whole circle; transitivity does the rest.

This elementary fact sharpens the inverse problem. An exactly invariant congruence–real graph cannot be the obstruction. The quantitative argument must produce a near-return before an invariant limit is reached; the finite divisibility and the three archimedean bounds then make that

return exact. Only at that point does the rational Borel appear, as the stabilizer of the exact incidence.

The joint representation has a standard geometric interpretation. If $X = G_7/H$ is the projective congruence orbit and ρ is the boundary representation, the tensor-induction identity gives

$$\ell^2(X) \otimes \rho \simeq \mathrm{Ind}_H^{G_7}(\rho|_H). \tag{636}$$

Thus the operator in (625) is the spectral operator on the pointed congruence cover of the triangle orbifold, with boundary coefficients. Under the map

$$G_7\backslash\mathbb{H} \longrightarrow \mathrm{SL}_2(\mathcal{O}_K)\backslash\mathbb{H}^3, \qquad z \longmapsto (z, F_1(z), F_2(z)),$$

the Hodge twists are the two outer automorphic line bundles restricted to the triangle curve.

There is a useful exact stabilizer statement behind this picture. Let $\widetilde{C}_7 = \{(z, F_1(z), F_2(z))\} \subset \mathbb{H}^3$. Then

$$\mathrm{Stab}_{\mathrm{PSL}_2(\mathcal{O}_K)}(\widetilde{C}_7) = G_7. \tag{637}$$

To prove this, project the ambient stabilizer to the first real factor. The projection is faithful because each F_i is nonconstant, and it is discrete because convergence in the first factor forces convergence in the other two by equivariance. It is therefore a cofinite Fuchsian supergroup of G_7. If its index were $n > 1$, its cusped quotient would have

$$-\chi = \frac{5}{14n} \le \frac{5}{28} < \frac{1}{4}.$$

The only orientable cusped hyperbolic orbifold with $-\chi < 1/4$ has signature $(2, 3, \infty)$ and $-\chi = 1/6$, which is not $5/(14n)$. Hence $n = 1$.

At every good ideal $\mathfrak{a}$, the pointed congruence cover has an equally concrete ambient description. Let $\Lambda_{B,\mathfrak{a}} \subset \mathrm{PSL}_2(\mathcal{O}_K)$ be the inverse image of a residue Borel and put $G_{7,B,\mathfrak{a}} = G_7 \cap \Lambda_{B,\mathfrak{a}}$. Then

$$G_{7,B,\mathfrak{a}}\backslash\mathbb{H}$$

is the connected pullback of the Borel-level Hilbert modular cover. Indeed, its components are indexed by

$$G_7\backslash \mathrm{PSL}_2(\mathcal{O}_K)/\Lambda_{B,\mathfrak{a}},$$

and strong approximation makes this double-coset set a singleton. In particular the degree is exactly

$$[G_7 : G_{7,B,\mathfrak{a}}] = \#\mathbf{P}^1(\mathcal{O}_K/\mathfrak{a}).$$

There is a subtle but important reason to work with this finite tower rather than invoke a closed-orbit theorem directly. Fix $x \in \mathbf{P}^1(K)$, put $\Gamma = G_7$, and let $\Gamma_x = \mathrm{Stab}_\Gamma(x)$. The denominator identity

$$j_{gh}(x) = j_g(hx)j_h(x)$$

shows that $j_g(x)$ is a cocycle on the action groupoid, not a scalar cocycle on the compact triangle quotient with x held fixed. The space on which that fixed flag does descend is

$$\Gamma\backslash\big(T^1\mathbb{H} \times \Gamma x\big) \simeq \Gamma_x\backslash T^1\mathbb{H}. \tag{638}$$

Its closed geodesics are represented by hyperbolic elements of Γ_x. Thus, if the stabilizer is projectively trivial, an ordinary Livšic argument has no periodic orbits on which to test its cocycle. Using such an argument to produce a closed orbit would be circular.

The pointed congruence covers avoid that circle. Define

$$\Gamma_x(\mathfrak{a}) = \{g \in \Gamma : g[x] = [x] \pmod{\mathfrak{a}}\}.$$

They have finite index before anything is known about Γ_x, and

$$\Gamma_x = \bigcap_{\mathfrak{a}} \Gamma_x(\mathfrak{a}). \tag{639}$$

Indeed, for a primitive representative P, membership in the right-hand side says that $\det(P, gP)$ is divisible by every ideal, and hence vanishes. The genuinely group-theoretic target is now exact: construct Hodge-balanced closed loops on these finite covers, uniformly in their level, at a radius for which the product formula forces one loop into the intersection (639). The optimal-lifting formulation below is precisely this inverse-limit problem.

The logarithmic normal bundle of the modular graph has an exact Hodge splitting:

$$N^{\log}_{C_7/X_K} \simeq \mathcal{L}_1^{-2} \oplus \mathcal{L}_2^{-2}, \qquad (\deg \mathcal{L}_1^{-2}, \deg \mathcal{L}_2^{-2}) = \left(-\frac{1}{14}, -\frac{3}{14}\right). \tag{640}$$

Indeed, the logarithmic tangent of the threefold restricts as $\mathcal{L}_0^{-2} \oplus \mathcal{L}_1^{-2} \oplus \mathcal{L}_2^{-2}$; because the first coordinate of the graph is the identity, quotienting its tangent line leaves the last two factors. Both have negative degree, so

$$H^0(C_7, N^{\log}_{C_7/X_K}) = 0.$$

Thus the curve admits no infinitesimal logarithmic deformation in the Hilbert modular threefold. This is a genuine Hodge-theoretic rigidity of the whole thin embedding.

The $(1/10)$ closing window is itself a Hodge degree. Define

$$\mathcal{S} = \mathcal{L}_0 \otimes \mathcal{L}_1^{-1} \otimes \mathcal{L}_2^{-1}.$$

Then

$$\deg \mathcal{S} = \frac{1}{28}, \qquad \frac{\deg \mathcal{L}_0 + \deg \mathcal{L}_1 + \deg \mathcal{L}_2}{2 \deg \mathcal{L}_0} = \frac{9}{10}, \qquad \frac{\deg \mathcal{S}}{2 \deg \mathcal{L}_0} = \frac{1}{10}. \tag{641}$$

The first ratio is the sum of the three column-growth exponents in hyperbolic-radius normalization; the second is the remaining entropy-versus-height surplus. Thus the one-tenth margin is a fixed positive line-bundle degree. The missing Arakelov construction is a canonical map carrying the finite incidence divisor into this line at a stopped return; changing the Borel frame separately at each return loses that global map.

This makes an ambient relative trace formula a concrete alternative to the intrinsic transfer operator. Its stabilizer term is exactly the thin-group orbital kernel; the other terms come from translates of the curve. Since the curve has complex codimension two, its two normal Hodge directions are the two outer line bundles with exponents $1/5$ and $3/5$. A global Arakelov intersection bound for the non-stabilizing translates would therefore replace separate finite-place dispersion and retain precisely the two missing Hodge coordinates.

Here "relative trace" does not mean that the modular graph is a homogeneous subgroup orbit. It cannot be one. A complex one-dimensional symmetric subgroup orbit in $\mathbb{H}^3$, nonconstant in all three factors, has Möbius projections and hence equal tangent degrees in the three factors. The modular graph has degree ratio

$$\deg \mathcal{L}_0 : \deg \mathcal{L}_1 : \deg \mathcal{L}_2 = 5 : 1 : 3,$$

so its outer period maps are strict contractions. The correct identity is instead a nonhomogeneous restriction pretrace formula.

It has an exact elementary geometric form. Write $j(z) = (z, F_1(z), F_2(z))$, let $\mathcal{F}$ be a fundamental domain for G_7, and let k be a compactly supported point-pair kernel on $\mathbb{H}^3$. Then

$$\begin{aligned}&\int_{\mathcal{F}^2} a(z)\overline{b(w)} \sum_{\lambda\in \mathrm{PSL}_2(\mathcal{O}_K)} k(j(z), \lambda j(w))\, d\mu(z)d\mu(w)\\ &\quad= \int_{\mathcal{F}^2} a(z)\overline{b(w)} \sum_{\gamma\in G_7} k(j(z), j(\gamma w))\, d\mu(z)d\mu(w) + \mathcal{I}_k^{\mathrm{tr}}(a,b),\end{aligned} \tag{642}$$

where $\mathcal{I}_k^{\mathrm{tr}}$ is the same sum over ambient elements outside G_7. The full-stabilizer theorem proves that the middle term is the kernel of the *entire* thin group and that every term in $\mathcal{I}_k^{\mathrm{tr}}$ moves the curve transversely.

At Borel level $\mathfrak{p}$, the connected-pullback theorem identifies the restriction of (642) with the pointed degree-$q+1$ cover. Removing the constant fibre leaves exactly the rank-one finite test

$$e_{[P \bmod \mathfrak{p}]} \otimes e_\infty^* \quad \text{on} \quad \mathrm{St}_{\mathfrak{p}}.$$

When a transverse term changes residue cusp frames, it passes through an open Bruhat cell. The Weyl matrix computed above is Kloosterman-flat, so every such cusp-Fourier coefficient already has the required $O(q^{-1/2})$ local factor. There is one important closed-Bruhat channel: a translate may move the curve in characteristic zero while remaining in the same residue Borel modulo $\mathfrak{p}$. Its finite vector lies in the unflattened Borel-fixed line. At one prime this says only congruence incidence, not that the translate lies in $B_P(K)$. Thus the remaining problem has two exact parts: retain the local $q^{-1/2}$ while summing the open-Bruhat transverse terms, and prove that coherent closed-Bruhat concentration across growing levels and Hodge stopping scales forces $G_7 \cap B_P(K) \neq \{1\}$. The latter is where the product formula and the negative normal degrees must enter; neither part is a generator inequality.

The embedding alone does not identify the intrinsic spectrum of a curve cover with the ambient Hilbert modular automorphic spectrum. A relative trace comparison and a bound for all non-stabilizing translates are additional requirements. Only after that comparison would ordinary ambient Ramanujan and generalized Lindelöf period conjectures imply the square-root estimate below.

The missing geometric estimate can now be stated without referring to generators: construct the arithmetic normal displacement of a distinct level translate and show that its finite congruence vanishing together with its archimedean shadowing contradicts the two negative degrees in (640). That would be a global Arakelov gap principle controlling every finite place at once.

This gives a natural conditional theorem. Assume tempered new spectrum for these covers and a generalized Lindelöf bound for the resulting packet of restriction periods, uniformly in the level, real frequency, and the bounded Mellin parameters imposing the two Hodge windows. Precisely, assume that the three-block count has main term

$$\mathcal{M}_R = e^{(1-s+o(1))R}$$

and spectral error

$$|\mathcal{E}_R| \le \mathcal{M}_R^{1/2} e^{o(R)}. \tag{643}$$

Then $s < 1$ makes the count positive. Any counted triple has incidence defect divisible by the chosen prime and bounded at the three real places by the real and Hodge cutoffs. Since $s > 9/10 + 2\delta$,

the product formula makes the defect zero, and the compatible positive block language makes the resulting stabilizer projectively nonidentity. Hence this standard-shaped Ramanujan–Lindelöf package implies the classification.

Ramanujan alone does not imply (643): it gives square-root cancellation relative to the entire orbital ball. The Lindelöf assertion is square-root in the much smaller local expected count. The remaining work in this route is therefore a relative-period estimate on the whole congruence tower, not a better inequality for the generators.

A.9 Comparison: a congruence shrinking-target criterion

There is a less efficient variant which replaces a chosen continued-fraction prefix by a short element in a congruence Borel. The key point is that the Hodge exponent gap is just large enough to turn a simultaneous real and congruence return into an exact stabilizer.

Fix $x = [P] \in \mathbf{P}^1(K)$, with $P = (p, q)^t \in \mathcal{O}_K^2$ primitive. For $g \in G_7$, put

$$D_g(P) = \det(P, gP) \in \mathcal{O}_K.$$

Thus $D_g(P) = 0$ precisely when g fixes x. Let

$$R(g) = d_{\mathbb{H}}(i, \sigma_0(g)i), \qquad a_i(g) = \log \|\sigma_i(g)\|_{\mathrm{op}}.$$

Changing the basepoint norms changes every formula below by $O_P(1)$, and one has

$$a_0(g) \le \frac{1}{2} R(g) + O(1). \tag{644}$$

Proposition A.26 (Congruence–Hodge closing)**.** *Suppose there are integral ideals $\mathfrak{a}_n$, with $\mathrm{N}\,\mathfrak{a}_n \to \infty$, and projectively nonidentity elements $g_n \in G_7$ such that, for constants $A, \beta, \theta > 0$,*

$$g_n x \equiv x \pmod{\mathfrak{a}_n}, \tag{645}$$

$$R(g_n) \le A \log \mathrm{N}\,\mathfrak{a}_n + O(1), \tag{646}$$

$$|\sigma_0(g_n) x_0^\dagger - x_0^\dagger| \le \exp(-\beta R(g_n)), \tag{647}$$

$$a_1(g_n) + a_2(g_n) \le \theta R(g_n) + O(1). \tag{648}$$

Assume that the affine expression in (647) stays in one fixed chart and that

$$A\left(\frac{1}{2} - \beta + \theta\right) < 1. \tag{649}$$

Then x is fixed by a projectively nonidentity element of G_7. Consequently x is a cusp or is special.

Proof. Write $g_n = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$. In the affine chart $x = p/q$,

$$\sigma_0(D_{g_n}(P)) = q(cp + dq)\big(x - \sigma_0(g_n)x\big).$$

Equations (644) and (647) therefore give

$$\log |\sigma_0 D_{g_n}(P)| \le \left(\frac{1}{2} - \beta\right) R(g_n) + O_P(1).$$

At the other two real places,

$$|\sigma_i D_{g_n}(P)| \le \|\sigma_i(P)\|^2 \, \|\sigma_i(g_n)\|_{\mathrm{op}},$$

so (648) yields

$$\log|\mathrm{N}\, D_{g_n}(P)| \le \left(\frac{1}{2} - \beta + \theta\right) R(g_n) + O_P(1). \tag{650}$$

The congruence (645) says exactly that $D_{g_n}(P) \in \mathfrak{a}_n$. If this determinant were nonzero, then

$$|\mathrm{N}\, D_{g_n}(P)| \ge \mathrm{N}\, \mathfrak{a}_n.$$

Put $\kappa = 1/2 - \beta + \theta$. If $\kappa \le 0$, then (650) is bounded independently of n and is eventually smaller than $\mathrm{N}\, \mathfrak{a}_n$. If $\kappa > 0$, then (646) and $A\kappa < 1$ give the same strict inequality. Hence $D_{g_n}(P) = 0$. Proposition 7.1 gives the two stated alternatives. □

The constants reveal the point. Lattice points in a hyperbolic ball of radius R number $e^{R+o(R)}$. Fixing one projective point modulo $\mathfrak{a}$ has expected cost $\mathrm{N}\, \mathfrak{a}$, and asking for a real return of radius $e^{-\beta R}$ has expected cost $e^{\beta R}$. Thus one expects such elements whenever

$$A(1-\beta) > 1, \qquad R = A \log \mathrm{N}\, \mathfrak{a}. \tag{651}$$

For Hodge-typical elements,

$$\theta = \frac{1}{2}\left(\frac{1}{5} + \frac{3}{5}\right) = \frac{2}{5}.$$

The counting condition and the closing condition become

$$\frac{9}{10} - \frac{1}{A} < \beta < 1 - \frac{1}{A}.$$

This interval has width $1/10$. For example, $A = 2$ permits any $2/5 < \beta < 1/2$. The width is exactly the integral Asai gap

$$\frac{1}{2} - \left(\frac{1}{10} + \frac{3}{10}\right) = \frac{1}{10}$$

in hyperbolic-distance normalization.

This suggests a precise full-group conjecture.

Conjecture A.27 (Borel congruence shrinking-target dichotomy)**.** *For every $x \in \mathbf{P}^1(K)$, either* $\mathrm{Stab}_{G_\tau}(x) \ne \{\pm 1\}$*, or there are $A > 1$, a number*

$$\frac{9}{10} - \frac{1}{A} < \beta < 1 - \frac{1}{A},$$

ideals $\mathfrak{a}_n$ of unbounded norm, and elements g_n satisfying (645)–(648) with $\theta = 2/5 + o(1)$.

Theorem A.28 (Conditional congruence completion)**.** *Conjecture A.27 implies the heptagonal cusp classification.*

Proof. The first alternative of the conjecture is the desired conclusion. In the second, choose the displayed β with a small margin. Then (649) holds, and Proposition A.26 supplies the same conclusion. □

Strong approximation supplies large congruence images of the Zariski-dense thin group, and expansion predicts the factor $(\mathrm{N}\, \mathfrak{a})^{-1}$ in (651). Hyperbolic mixing predicts the shrinking-target factor $e^{-\beta R}$, while the two Hodge Lyapunov exponents predict (648). What must be proved is their joint version at this entropy scale, with a periodic orbit as the only possible concentration obstruction. This does not track a prime which happens to appear in a remainder. The modulus is chosen externally, and the proof compares its entire ideal norm with the product over all three real places in one step.

A.10 What a full-group transfer theorem would have to say

It is important not to hide the remaining arithmetic problem inside the word "mixing." The presentation

$$G_7 \simeq C_2 * C_7$$

gives a global Markov coding of the whole group. After accelerating across the parabolic cusp, let $\mathscr{H}$ denote its inverse branches. The central length increment is scalar, for example $\ell(h,z) = -\log|h'(z)|$. The two outer increments are matrix-cocycle quantities. For a moving direction $[u_i]$, put

$$\psi_i(h,[u_i]) = \log \frac{\|\sigma_i(h)u_i\|}{\|u_i\|}.$$

The relevant operator therefore acts on functions of $(z,[u_1],[u_2],\xi)$. Schematically, one inverse branch changes the state by

$$(z,[u_1],[u_2],\xi) \longmapsto \big(hz,[\sigma_1(h)u_1],[\sigma_2(h)u_2],h\xi\big), \tag{652}$$

and is weighted by $e^{-s\ell+u_1\psi_1+u_2\psi_2}$. Here ξ belongs to the finite projective orbit

$$X_{\mathfrak{a}}(x) = G_7 \cdot [P \bmod \mathfrak{a}].$$

Thus the symbolic coordinate, two moving projective directions, the congruence action, and both Galois representations occur in one operator. This is a use of the full thin group; estimating six generators one at a time is not.

For fixed $\epsilon > 0$, let $N_x(R,\mathfrak{a},\beta,\epsilon)$ count the elements in a hyperbolic shell of radius R which satisfy

$$gx \equiv x \pmod{\mathfrak{a}}, \qquad |\sigma_0(g)x_0^\dagger - x_0^\dagger| \le e^{-\beta R},$$
$$a_1(g)+a_2(g) \le (2/5+\epsilon)R.$$

The transfer statement sufficient for the closing argument is a uniform lower bound of the form

$$N_x(R,\mathfrak{a},\beta,\epsilon) \gg_{x,\epsilon} \frac{e^{(1-\beta)R}}{\#X_{\mathfrak{a}}(x)} \tag{653}$$

through a sequence for which the right side tends to infinity. A useful spectral formulation would prove this with a relative error $O(e^{-\eta R})$, uniformly while

$$\log \#X_{\mathfrak{a}}(x) \le (1-\beta-2\eta)R.$$

The periodic concentration alternative must be retained: if the boundary ray closes, its cyclic stabilizer is the desired special case.

There are two reasons why separate expansion and mixing theorems do not yet give (653). First, an interval of radius $e^{-\beta R}$ around a prescribed boundary point is a cylinder which fixes roughly βR units of the symbolic future. An unrestricted large-deviation theorem for the Hodge cocycle does not imply the same theorem after that cylinder has been fixed. For an arbitrary K-rational endpoint, this conditional estimate is precisely the arithmetic boundary rigidity under investigation. Second, there is very little entropy available for an error term. Put

$$r = \frac{\log \#X_{\mathfrak{a}}(x)}{R}.$$

The expected count requires $r + \beta < 1$, whereas the closing lemma, with the typical outer sum $2/5$, requires $r > 9/10 - \beta$. Consequently

$$0 < 1 - r - \beta < \frac{1}{10}. \tag{654}$$

This bound is independent of how the congruence scale is chosen. If A is close to one, the real target is short but congruence mixing must be nearly entropy-optimal. If A is large, congruence mixing is easier but the target fixes almost nine tenths of the boundary word. The two regimes are the same obstruction viewed from opposite sides.

> **What the thin group has and has not supplied.** The free-product coding, the projective congruence orbits, and the two Hodge weights in (652) are genuine global structures of G_7. They produce the unconditional closing Proposition A.26 and the exact budget (654). They do not by themselves prove the conditional Hodge estimate on every rational cylinder. Treating that estimate as ordinary mixing would simply rename the central difficulty.

References

[1] W. Ljunggren, *Zur Theorie der Gleichung* $x^2 + 1 = Dy^4$, Avh. Norske Vid. Akad. Oslo I, no. 5 (1942).

[2] J. H. E. Cohn, *Perfect Pell powers*, Glasgow Math. J. **38** (1996), 19–20, https://doi.org/10.1017/S0017089500031207.

[3] Y. Bugeaud, P. Hubert, and T. A. Schmidt, *Transcendence with Rosen continued fractions*, J. Eur. Math. Soc. **15** (2013), 39–51, https://arxiv.org/abs/1007.2050.

[4] B. Adamczewski and Y. Bugeaud, *On the complexity of algebraic numbers II. Continued fractions*, Acta Math. **195** (2005), 1–20.

[5] P. Arnoux and T. A. Schmidt, *Cross sections for geodesic flows and* α*-continued fractions*, Nonlinearity **26** (2013), 711–726.

[6] P. Arnoux and T. A. Schmidt, *Veech surfaces with nonperiodic directions in the trace field*, J. Mod. Dyn. **3** (2009), 611–629, https://arxiv.org/abs/0906.4571.

[7] I. Bouw and M. Möller, *Teichmüller curves, triangle groups, and Lyapunov exponents*, Ann. of Math. (2) **172** (2010), 139–185.

[8] P. Brosnan and G. Pearlstein, *Jumps in the Archimedean height*, Duke Math. J. **168** (2019), 1737–1842.

[9] J. I. Burgos Gil, D. Holmes, and R. de Jong, *Positivity of the height jump divisor*, Int. Math. Res. Not. IMRN 2019, no. 7, 2044–2068.

[10] J. Bourgain and A. Kontorovich, *On representations of integers in thin subgroups of* $\mathrm{SL}_2(\mathbb{Z})$, Geom. Funct. Anal. **20** (2010), 1144–1174, https://arxiv.org/abs/1001.4534.

[11] J. Bourgain and A. Kontorovich, *The affine sieve beyond expansion I: thin hypotenuses*, Int. Math. Res. Not. IMRN 2015, no. 19, 9175–9205, https://arxiv.org/abs/1307.3535.

[12] Y. Benoist and S. Miquel, *Arithmeticity of discrete subgroups containing horospherical lattices*, Duke Math. J. **169** (2020), 1485–1539, https://arxiv.org/abs/1805.00045.

[13] A. Baker and G. Wüstholz, *Logarithmic forms and group varieties*, J. Reine Angew. Math. **442** (1993), 19–62, https://doi.org/10.1515/crll.1993.442.19.

[14] R. Canary, T. Zhang, and A. Zimmer, *Cusped Hitchin representations and Anosov representations of geometrically finite Fuchsian groups*, Adv. Math. **404** (2022), Paper No. 108439, https://arxiv.org/abs/2103.06588.

[15] F. Zhu and A. Zimmer, *Relatively Anosov representations via flows I: theory*, preprint, https://arxiv.org/abs/2207.14737.

[16] S. Cavallar and F. Lemmermeyer, *The Euclidean algorithm in cubic number fields*, Proceedings of the International Conference, Eger, 1996, 123–146 (1998), https://arxiv.org/abs/1202.6019.

[17] P. Lezowski and K. J. McGown, *The Euclidean algorithm in quintic and septic cyclic fields*, Math. Comp. **86** (2017), 2535–2549, https://arxiv.org/abs/1601.03433.

[18] J.-H. Evertse, H. P. Schlickewei, and W. M. Schmidt, *Linear equations in variables which lie in a multiplicative group*, Ann. of Math. (2) **155** (2002), 807–836.

[19] F. Calegari, V. Dimitrov, and Y. Tang, *Arithmetic holonomy bounds and effective Diophantine approximation*, preprint (2026), https://arxiv.org/abs/2510.04156.

[20] F. Calegari, V. Dimitrov, and Y. Tang, *The linear independence of* 1, $\zeta(2)$, *and* $L(2, \chi_{-3})$, preprint (2024), https://arxiv.org/abs/2408.15403.

[21] F. Calegari, V. Dimitrov, and Y. Tang, *The unbounded denominators conjecture*, J. Amer. Math. Soc. **38** (2025), 627–702, https://arxiv.org/abs/2109.09040.

[22] F. Calegari, *The adelic closure of triangle groups*, to appear in Groups Geom. Dyn., https://www.math.uchicago.edu/~fcale/papers/Chaos.pdf.

[23] M. Ru and P. Vojta, *Schmidt's subspace theorem with moving targets*, Invent. Math. **127** (1997), 51–65.

[24] G. Zhao and Y. Li, *On a theorem of Schmidt's subspace theory with moving targets*, preprint (2026), https://arxiv.org/abs/2603.01627.

[25] S. Filip, *Uniformization of some weight* 3 *variations of Hodge structure, Anosov representations, and Lyapunov exponents*, arXiv:2110.07533v2.

[26] K. Golubev and A. Kamber, *On Sarnak's density conjecture and its applications*, Forum Math. Sigma **11** (2023), e48.

[27] E. Hanson, A. Merberg, C. Towse, and E. Yudovina, *Generalized continued fractions and orbits under the action of Hecke triangle groups*, Acta Arith. **134** (2008), 337–348, https://doi.org/10.4064/aa134-4-4.

[28] Y. Hao, *Groups generated by root unipotents: higher-rank and rank-one*, preprint (2026), https://arxiv.org/abs/2607.04580.

[29] A. Kamber and P. P. Varjú, *Lifting all elements in* $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$, preprint, https://arxiv.org/abs/2310.10269.

[30] D. Mayer, T. Mühlenbruch, and F. Strömberg, *The transfer operator for the Hecke triangle groups*, Discrete Contin. Dyn. Syst. **32** (2012), 2453–2484, https://arxiv.org/abs/0912.2236.

[31] A. Adam and A. Pohl, *A transfer-operator-based relation between Laplace eigenfunctions and zeros of Selberg zeta functions*, https://arxiv.org/abs/1606.09109.

[32] H. Oh and D. Winter, *Uniform exponential mixing and resonance free regions for convex cocompact congruence subgroups of* $\mathrm{SL}_2(\mathbb{Z})$, https://arxiv.org/abs/1410.4401.

[33] M. Magee, H. Oh, and D. Winter, *Uniform congruence counting for Schottky semigroups in* $\mathrm{SL}_2(\mathbb{Z})$, https://arxiv.org/abs/1601.03705.

[34] C. Kraaikamp, T. A. Schmidt, and I. Smeets, *Natural extensions for* α*-Rosen continued fractions*, J. Math. Soc. Japan **62** (2010), 649–671.

[35] J. Li, *Fourier decay, renewal theorem and spectral gaps for random walks on split semisimple Lie groups*, Ann. Sci. Éc. Norm. Supér. (4) **55** (2022), 1613–1686, https://arxiv.org/abs/1811.06484.

[36] W. He and N. de Saxcé, *Semisimple random walks on the torus*, Ergodic Theory Dynam. Systems **45** (2025), 3090–3147, https://arxiv.org/abs/2204.11453.

[37] C. McMullen, *Billiards, heights, and the arithmetic of non-arithmetic groups*, Invent. Math. **228** (2022), 1309–1351.

[38] D. Rosen, *An arithmetic characterization of the parabolic points of* $G(2\cos(\pi/5))$, Glasgow Math. J. **6** (1963), 88–96, https://doi.org/10.1017/S204061850003478X.

[39] G. A. Margulis, *Discrete subgroups of semisimple Lie groups*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3), vol. 17, Springer-Verlag, Berlin, 1991.

[40] M. Möller and D. Zagier, *Modular embeddings of Teichmüller curves*, Compos. Math. **152** (2016), 2269–2349.

[41] D. Rosen and T. A. Schmidt, *Hecke groups and continued fractions*, Bull. Austral. Math. Soc. **46** (1992), 459–474.

[42] R. Schoof, *Computing Arakelov class groups*, in *Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography*, Math. Sci. Res. Inst. Publ. **44**, Cambridge Univ. Press, 2008, 447–495, https://arxiv.org/abs/0801.3835.

[43] G. Tomanov, *Locally divergent orbits on Hilbert modular spaces and Margulis conjectures*, J. Reine Angew. Math. **689** (2014), 1–42.

[44] L. N. Vaserstein, *On the group* SL_2 *over Dedekind rings of arithmetic type*, Math. USSR-Sb. **18** (1972), 321–332.

[45] K. Winsor, *Hecke triangle groups and special hyperbolic elements*, preprint (2026), https://arxiv.org/abs/2605.30064.

[46] X. Zhang, *On representation of integers from thin subgroups of* $\mathrm{SL}_2(\mathbb{Z})$ *with parabolics*, https://arxiv.org/abs/1610.00770.

Part II
Research Archive and Proof Audits

II.1 Precise target and notation

Let $K = \mathbf{Q}(2\cos(\pi/7))$, let $G_7 \subset \mathrm{SL}_2(\mathcal{O}_K)$ be the heptagonal group, and let g be a recurrent positive return word. At the three real embeddings write

$$a = \log \rho(\sigma_0(g)), \qquad b = \log \rho(\sigma_1(g)), \qquad c = \log \rho(\sigma_2(g)).$$

The central Asai weight, transverse contraction, and Roth surplus are

$$\chi = -a + b + c, \qquad m = \min(b, c), \qquad \delta = m - \chi = a - \max(b, c).$$

In the residual case under discussion $\chi > 0$, and spectral domination gives $\delta > 0$. The manuscript proposes a rational point Y_g and a moving mixed Asai eigenline L_g satisfying

$$h(Y_g) = \chi + O(1), \qquad m_\infty(Y_g, L_g) := -\log d_\infty(Y_g, L_g) \geq 2m - O(1). \tag{655}$$

Thus the nominal approximation exponent is

$$\kappa = \frac{2m}{\chi} = 2 + \frac{2\delta}{\chi}. \tag{656}$$

The global theorem to be proved is

$$\mathbf{P}^1(K) = G_7\infty \sqcup \bigcup_{\substack{g \in G_7 \text{ special} \\ g \text{ hyperbolic}}} G_7 \,\mathrm{Fix}(g).$$

Two logical cautions are kept active throughout this file. First, the construction of the recurrent pairs (Y_g, L_g) with the uniform constants in (655) must be checked, not treated as a definition. Second, a fixed-target theorem cannot be applied separately to a sequence in which each target occurs only once.

II.2 Proved algebraic statements and exact computational certificates

Put $\lambda = 2\cos(\pi/7)$, $\alpha = \lambda^2 - 1$, and

$$A_1 = \begin{pmatrix} 1 & \lambda \\ 0 & 1 \end{pmatrix}, \qquad A_2 = \begin{pmatrix} \lambda & \alpha \\ 1 & \lambda \end{pmatrix}.$$

For a positive word, write its trace uniquely as $A + B\lambda + C\alpha$ with $A, B, C \in \mathbf{Z}_{\geq 0}$. The fusion identities

$$\lambda^2 = 1 + \alpha, \qquad \lambda\alpha = \lambda + \alpha, \qquad \alpha^2 = 1 + \lambda + \alpha$$

give a finite-state propagation rule for the defect $\partial = A - 2B - 3C$. The defect automaton proves that every nonparabolic positive word satisfies $A \leq 2B + 3C$. Comparing the conjugate traces and using monotonicity of $\mathrm{arcosh}(|t|/2)$ gives the following result.

Theorem II.2.1 (Uniform outer spectral gap)**.** *For every hyperbolic positive word with $\chi > 0$,*

$$\delta = a - \max(b, c) \geq \delta_7 := \log \frac{\sigma_0(\lambda)}{|\sigma_2(\lambda)|} = 0.368138295478646\ldots.$$

The bound is approached by $A_2 A_1^N$ and its symmetric variants.

This is the *outer-gap* result. Its proof is algebraic number theory (comparison of the three real conjugates and the unit/fusion identities) implemented through finite-state fusion combinatorics (the coefficient semiring and defect automaton). It proves the X-part $\delta \geq \delta_7$, not the Y-part consisting of construction and arithmetic control of the moving target.

Experiment II.2.2 (Exhaustive and random word tests)**.** The C++ program `curt10_spectral_tests.cpp` uses exact integer fusion coefficients. Exhaustion through length 11 checked $6^{11} = 362{,}797{,}056$ words. The smallest value at length 11 was $0.37048888\ldots$, on the expected one-defect parabolic-ray family. Five million additional exact-coefficient random words of lengths up to 20 gave minimum $0.368900685\ldots$. No violation of $\delta \geq \delta_7$ occurred. These computations motivated the defect automaton but are not used as a finite proof of the theorem.

Proposition II.2.3 (The outer gap is an integrated Hodge deficit)**.** *Let g be a positive hyperbolic return and let γ_g be its closed central geodesic, parametrized by hyperbolic arclength. For a fixed compact core $\mathcal{K}$ of the Teichmüller curve, put*

$$T_{\mathcal{K}}(g) = \text{length}\{t : \gamma_g(t) \in \mathcal{K}\}, \qquad \kappa_i = \inf_{\mathcal{K}} \big(1 - \|DF_i\|_{\text{hyp}}\big) > 0.$$

Then

$$\delta(g) \geq \frac{1}{2} \min(\kappa_1, \kappa_2)\, T_{\mathcal{K}}(g). \tag{657}$$

Proof. The central translation length of g is $2a$. Equivariance of F_i makes $F_i \circ \gamma_g$ a path from z_i to $\sigma_i(g) z_i$, so its length is at least the translation length $2b_i$ of $\sigma_i(g)$, where $b_i = b, c$. Schwarz–Pick gives

$$2b_i \leq \int_{\gamma_g} \|DF_i\|_{\text{hyp}}\, dt.$$

Subtracting from $2a = \text{length}(\gamma_g)$, and using the definition of κ_i on $\mathcal{K}$, gives

$$a - b_i \geq \frac{\kappa_i}{2} T_{\mathcal{K}}(g).$$

Take the minimum over $i = 1, 2$. □

Lemma II.2.4 (Induced symbols carry uniform thick time)**.** *There are a compact core $\mathcal{K}$ of the heptagonal orbifold and a constant $c_{\text{geo}} > 0$, depending only on that orbifold, such that every cyclic positive hyperbolic word satisfies*

$$T_{\mathcal{K}}(g) \geq c_{\text{geo}}\, \nu(g). \tag{658}$$

Proof. Use the standard finite cross-section for the geodesic flow on a Hecke triangle surface, obtained from the sides of a truncated heptagonal fundamental polygon; an explicit version, including its return-time function, is given in [27, Sections 4.3–4.4]. Choose the cusp horoballs with pairwise disjoint interiors. In the positive cutting sequence (a finite recoding of that cross-section), a maximal A_1- or A_6-run is exactly the part which winds through consecutive translates abutting

one horoball. Collapsing such a run is therefore the same as inducing the cross-section past one cusp excursion.

After this induction there are only finitely many compact transition types. The zero-thick-time arrows in their boundary transition graph have no directed cycle: such a cycle would remain on the boundary of one horoball, so its symbols would all be A_1, or all A_6, and would have been one induced symbol. Take M_7 larger than the number of states. On the compact natural-extension rectangles, the sum of the thick parts of any M_7 successive induced return times then has a positive minimum ℓ_7. Compactness therefore gives

$$T_{\mathcal{K}}(g) \geq \frac{\ell_7}{M_7}\nu(g) - O_7(1).$$

Apply this inequality to the cyclic word w^k. Its closed geodesic and induced complexity are the k-fold repetitions of those for w, so $T_{\mathcal{K}}(g^k) = kT_{\mathcal{K}}(g)$ and $\nu(w^k) = k\nu(w)$. Divide by k and let $k \to \infty$; this removes the additive constant and proves (658). □

Corollary II.2.5 (A proved coarse relative fusion gap)**.** *There is an effective constant $c_{\mathrm{rel}} > 0$, depending only on the two conjugate period maps and the chosen Hecke cross-section, such that*

$$\delta(g) \geq c_{\mathrm{rel}}\,\nu(g)$$

for every positive hyperbolic return. One may take

$$c_{\mathrm{rel}} = \frac{1}{2}c_{\mathrm{geo}}\min(\kappa_1, \kappa_2).$$

Proof. Combine Proposition II.2.3 with Lemma II.2.4. □

Lemma II.2.6 (Alternating-shear monotonicity)**.** *For $x > 0$, put*

$$U_x(n) = \begin{pmatrix} 1 & nx \\ 0 & 1 \end{pmatrix}, \qquad L_x(n) = \begin{pmatrix} 1 & 0 \\ nx & 1 \end{pmatrix}.$$

If $x_0 > x_1 > 0$ and $n_1, \dots, n_{2p} \geq 1$, then

$$\frac{\operatorname{tr}\prod_{j=1}^{p} U_{x_0}(n_{2j-1})L_{x_0}(n_{2j})}{\operatorname{tr}\prod_{j=1}^{p} U_{x_1}(n_{2j-1})L_{x_1}(n_{2j})} \geq \frac{\operatorname{tr}\big((U_{x_0}(1)L_{x_0}(1))^p\big)}{\operatorname{tr}\big((U_{x_1}(1)L_{x_1}(1))^p\big)}. \tag{659}$$

Proof. We record the elementary continuant monotonicity used here. If

$$M = L_x(a_1)U_x(a_2)\cdots U_x(a_{2q-2})L_x(a_{2q-1}),$$

write $y = x^2$, $N(y) = \operatorname{tr} M$, and $Q(y) = xM_{21}$. Expanding the two-state path sum gives

$$Q'(y)N(y) - Q(y)N'(y) \in \mathbf{R}_{>0}[y]. \tag{660}$$

For completeness, this is also a direct continuant induction. Write $M = \left(\begin{smallmatrix} A & xB \\ xC & D \end{smallmatrix}\right)$. Appending $U_x(b)L_x(c)$ replaces

$$(A, B, C, D) \quad \text{by} \quad \big(A(1+bcy) + cyB,\ bA + B, C(1+bcy) + cD,\ byC + D\big).$$

Starting with $(1, 0, a_1, 1)$, substitution and collection of like powers shows inductively that the Wronskian in (660) has positive coefficients. Equivalently, one may pair the closed two-state paths

contributing a negative term in the Wronskian with the path obtained by switching the two tails at their first disagreement; the switched path has one additional transition and the same coefficient monomial. The unpaired terms are the positive ones. Thus Q/N is strictly increasing in y.

Fix all run parameters except, say, n_1, and cyclically cut the trace immediately after the corresponding upper shear. With the remaining odd alternating product denoted by M,

$$\operatorname{tr}(U_x(n_1)M) = N(x^2) + n_1 Q(x^2).$$

Since Q/N increases with x, the ratio of this expression at x_0 and x_1 increases with n_1. Transposition gives the same conclusion for a lower shear. Decreasing the parameters one at a time to one proves (659). □

Theorem II.2.7 (Sharp relative gap on the two-cusp language)**.** *Let*

$$w = A_1^{n_1} A_6^{n_2} \cdots A_1^{n_{2p-1}} A_6^{n_{2p}}, \qquad n_j \geq 1,$$

up to cyclic interchange of A_1, A_6. Then

$$\delta(w) > p\delta_7 = \frac{\delta_7}{2}\nu(w). \tag{661}$$

The same inequality holds for every positive power of such a word.

Proof. Put $x_0 = \sigma_0(\lambda)$ and $x_2 = |\sigma_2(\lambda)|$. The trace of an alternating shear product is a polynomial in x^2 with positive coefficients. Since $0 < \sigma_1(\lambda) < x_2 < x_0$, its largest outer trace is the one at σ_2. Lemma II.2.6 reduces the trace ratio to the case in which all $n_j = 1$.

Let

$$A(x) = \operatorname{arcosh}(1 + x^2/2).$$

The matrix $U_x(1)L_x(1)$ has trace $2 + x^2$, and hence

$$\operatorname{tr}\big((U_x(1)L_x(1))^p\big) = 2\cosh(pA(x)).$$

For $A_0 > A_2 > 0$, the function

$$p \longmapsto \frac{1}{p}\big(\log\cosh(pA_0) - \log\cosh(pA_2)\big)$$

is increasing: after multiplication by p^2, its derivative is $h(pA_0) - h(pA_2)$, where $h(u) = u\tanh u - \log\cosh u$ and $h'(u) = u\operatorname{sech}^2 u > 0$. It is therefore enough to use $p = 1$. Finally,

$$\frac{2 + x_0^2}{2 + x_2^2} > \frac{x_0}{x_2},$$

because the difference after clearing denominators is $(x_0 - x_2)(x_0x_2 - 2) > 0$. Consequently the logarithmic trace ratio is strictly larger than $p\log(x_0/x_2) = p\delta_7$.

All three traces are larger than two in absolute value. Since $t \mapsto \operatorname{arcosh}(t/2) - \log t$ is increasing on $(2, \infty)$, the spectral-exponent difference is at least the logarithmic trace ratio. This proves (661). The assertion about powers follows either from the same proof or from homogeneity of translation lengths and induced complexity. □

Theorem II.2.8 (Sharp relative gap on a one-cusp many-block language)**.** *Let*

$$w = (A_2A_1^{n_1})(A_2A_1^{n_2})\cdots(A_2A_1^{n_s}), \qquad n_j \geq 0.$$

If w has positive character, then

$$\delta(w) \geq s\delta_7 \geq \frac{\delta_7}{2}\nu(w).$$

The transpose statement holds with A_2, A_1 replaced by A_5, A_6.

Proof. Put $x_i = |\sigma_i(\lambda)|$, $R_7 = x_0/x_2 = e^{\delta_7} = \lambda^2 - \lambda$ at the distinguished embedding, and $D = \mathrm{diag}(1, R_7)$. At a positive real parameter x, set

$$B_n(x) = A_2(x)A_1(x)^n = \begin{pmatrix} x & (n+1)x^2 - 1 \\ 1 & (n+1)x \end{pmatrix}.$$

Let $E_i(n)$ be the entrywise absolute value of $\sigma_i(A_2A_1^n)$. Directly from the displayed matrix,

$$B_n(x_0) \ \geq_{\mathrm{ent}} \ R_7 D^{-1} E_i(n) D, \qquad i = 1, 2, \quad n \geq 0. \tag{662}$$

For $i = 2$, the diagonal entries and the lower-left entry are equal on the two sides, while the difference of the upper-right entries is $R_7^2 - 1 > 0$. For $i = 1$, the only additional inequality is

$$(n+1)x_0^2 - 1 \geq R_7^2\big|(n+1)x_1^2 - 1\big|.$$

When $(n+1)x_1^2 \leq 1$, its left side minus its right side is increasing in n, and the $n = 0$ case is $x_0^2 - 1 > R_7^2(1 - x_1^2)$. When $(n+1)x_1^2 \geq 1$, it follows from $x_0^2 > R_7^2 x_1^2$. These two root inequalities follow immediately from the ordered roots of $X^3 - X^2 - 2X + 1$ (or by squaring the elementary cosine expressions).

Entrywise order is preserved under multiplication of nonnegative matrices. The diagonal gauges in (662) telescope, giving

$$\prod_{j=1}^{s} B_{n_j}(x_0) \ \geq_{\mathrm{ent}} \ R_7^s D^{-1} \left(\prod_{j=1}^{s} E_i(n_j) \right) D.$$

The triangle inequality for the two-state trace path sum gives

$$\left| \mathrm{tr} \prod_j \sigma_i(A_2A_1^{n_j}) \right| \leq \mathrm{tr} \prod_j E_i(n_j).$$

Taking traces in the preceding entrywise comparison therefore proves

$$\frac{\sigma_0(\mathrm{tr}\, w)}{\max_i |\sigma_i(\mathrm{tr}\, w)|} \geq R_7^s. \tag{663}$$

Positive character forces both outer matrices to be hyperbolic. As in Theorem II.2.7, monotonicity of $\mathrm{arcosh}(t/2) - \log t$ upgrades (663) to $\delta(w) \geq s\delta_7$. Finally the induced cyclic word contains exactly s nonparabolic symbols and at most s parabolic runs, so $\nu(w) \leq 2s$. Transposition proves the final assertion. □

Theorem II.2.9 (Half-alphabet relative gap)**.** *Let w be any cyclic word in A_1, A_2, A_3, and let s be the number of its A_2, A_3 letters. If w has positive character, then*

$$\delta(w) \geq s\delta_7 \geq \frac{\delta_7}{2}\nu(w).$$

The reflected statement holds for every word in A_4, A_5, A_6, with s counting A_4, A_5.

Proof. Retain $R_7 = x_0/x_2$ and $D_+ = \operatorname{diag}(1, R_7)$ from the preceding proof. If $E_i(A_j)$ denotes the entrywise absolute value of $\sigma_i(A_j)$, direct substitution in the six displayed generators at the start of this section gives

$$\begin{aligned} A_1(x_0) &\geq_{\text{ent}} D_+^{-1} E_i(A_1) D_+, \\ A_2(x_0) &\geq_{\text{ent}} R_7 D_+^{-1} E_i(A_2) D_+, \qquad (i = 1, 2). \\ A_3(x_0) &\geq_{\text{ent}} R_7 D_+^{-1} E_i(A_3) D_+ \end{aligned}$$

The only equalities at $i = 2$ are the parabolic upper-right entry and the lower-left entry of A_2; all other required differences reduce to the positive root inequalities used in (662).

Multiplying these inequalities telescopes the common gauge and produces the factor R_7^s. The outer trace is bounded by the trace of the product of the entrywise absolute matrices, so

$$\frac{\sigma_0(\operatorname{tr} w)}{\max_i |\sigma_i(\operatorname{tr} w)|} \geq R_7^s.$$

The trace-to-spectral argument gives $\delta(w) \geq s\delta_7$. Since a nonparabolic cyclic word in this alphabet has at most s maximal A_1-runs, $\nu(w) \leq 2s$. For the reflected alphabet use $D_- = \operatorname{diag}(R_7, 1)$; the four entrywise checks are the transposes of the preceding ones. □

Theorem II.2.10 (Two-gauge cyclic language)**.** *Consider the two-state automaton*

state +	$A_1 : + \to +$	$A_2 : + \to +$	$A_3 : + \to -$
state −	$A_6 : - \to -$	$A_5 : - \to -$	$A_4 : - \to +$

and let w be any cyclic word accepted by it. If $s(w)$ is the number of nonparabolic letters and w has positive character, then

$$\delta(w) \geq s(w)\delta_7 \geq \frac{\delta_7}{2}\nu(w).$$

Proof. Use the two gauges

$$D_+ = \operatorname{diag}(1, R_7), \qquad D_- = \operatorname{diag}(R_7, 1).$$

The loop inequalities are those in Theorem II.2.9. Direct entrywise substitution gives the two bridge inequalities

$$\begin{aligned} A_3(x_0) &\geq_{\text{ent}} R_7 D_+^{-1} E_i(A_3) D_-, \\ A_4(x_0) &\geq_{\text{ent}} R_7 D_-^{-1} E_i(A_4) D_+ \end{aligned} \qquad (i = 1, 2).$$

Multiply along the accepted cyclic state path. Consecutive gauges cancel and the initial and terminal gauges agree, so taking traces and using the outer path-sum triangle inequality gives a factor $R_7^{s(w)}$. The trace-to-spectral argument proves the first inequality. Every parabolic run in an accepted nonparabolic cyclic word is separated from the next one by a nonparabolic letter; hence there are at most $s(w)$ such runs and $\nu(w) \leq 2s(w)$. □

Experiment II.2.11 (Alternating-shear certificate)**.** The exact symbolic program `parabolic_relative_gap_tests` checked the coefficientwise positivity in (660) on 600 independently generated continuants with up to six upper–lower pairs. Exhaustion of all run parameters in $\{1, 2, 3, 4\}$ through four pairs found the minimum trace slack at the all-one word in every case. The respective minima above $p\delta_7$ were

$$0.02117101599\ldots, \quad 0.13920191471\ldots, \quad 0.21913836478\ldots, \quad 0.29316273174\ldots.$$

These computations audit the two algebraic steps of Theorem II.2.7; its proof itself is the continuant monotonicity and Chebyshev argument above.

Conjecture II.2.12 (Relative fusion gap)**.** *For a cyclic positive word w, form its* induced complexity *$\nu(w)$ by replacing every cyclically maximal run of A_1's or A_6's by one letter and retaining every occurrence of A_2, A_3, A_4, A_5. If w is hyperbolic and $\chi(w) > 0$, then*

$$\delta(w) \geq \frac{\delta_7}{2}\,\nu(w). \tag{664}$$

The constant is sharp if this statement is true: the words $A_2A_1^N$ have induced complexity two and outer gap tending to δ_7.

This is strictly stronger than the uniform outer-gap theorem only when the number of induced blocks grows. It is also the spectral statement which matches the remaining arithmetic problem. Bounded $\nu(w)$ means a fixed number of parabolic parameters and hence, after passing to a subsequence, one fixed connector type; this is already the scope of Theorem II.14.11. Unbounded $\nu(w)$ would instead produce an outer surplus linear in the number of connector crossings, exactly the budget missing from the balanced many-block case.

Experiment II.2.13 (Additivity of the fusion gap)**.** The augmented exact program `curt10_spectral_tests.cpp` checked all $6^{10} = 60{,}466{,}176$ positive words of length ten. For every word with $\chi > 0$, it found

$$\delta(w)/\nu(w) \geq 0.1854705\ldots > \delta_7/2 = 0.1840691\ldots.$$

The record was $A_1^9A_2$, and the ratio decreases to $\delta_7/2$ along $A_1^NA_2$, as predicted. The separate program `fusion_additive_gap_tests.py` sampled 12000 induced words with up to thirty nonparabolic blocks and parabolic exponents up to 10^4. In the sampled class containing one nonparabolic letter per induced block it found the stronger inequality

$$\delta(w) \geq \delta_7\, s,$$

where s is the number of those blocks; the smallest excess was $3.9 \cdot 10^{-9}$ for a one-block parabolic ray. Purely alternating parabolic channels have a different asymptotic charge: for $(A_1A_6)^r$, direct computation gives $\delta/(2r) = 0.2207243102\ldots > \delta_7/2$. An additional 30000-sample mixed test allowed arbitrary interleavings of nonparabolic letters and the two parabolic channels, induced complexity up to 50, and parabolic exponents up to 10^5. Of the 28807 samples with positive character, none violated (664); the smallest slack was $4.24 \cdot 10^{-11}$, again on the sharp two-block degeneration $A_2A_1^{88880}$. No test proves (664), but together they identify both sharp boundary languages and show that the proposed additive constant is not an artifact of random generic words.

The exact-coefficient implementation now uses checked `unsigned __int128` intermediates and marks a matrix invalid before any `uint64_t` wraparound. Repeating the exhaustive length-ten calculation reported zero arithmetic overflows. In contrast, a deliberately overlong 20000-sample run with lengths up to 100 flagged 8941 samples as outside the fixed-width range. A previously printed apparent violation from such a run was therefore discarded; arbitrary-precision reevaluation gave $\chi = -41.0015\ldots$, $\delta = 43.1067\ldots$, and relative slack $32.9829\ldots$, so it was purely an overflow artifact. All long-run claims above come from the arbitrary-precision Python program.

Route II.2.14 (Geometric mechanism behind the relative gap)**.** For the two conjugate period maps, Schwarz–Pick gives

$$a - b \geq \kappa_1 T_{\text{thick}} - O(1), \qquad a - c \geq \kappa_2 T_{\text{thick}} - O(1)$$

on a closed return, with positive constants on a fixed thick core. Every passage between induced parabolic channels crosses that core. Lemma II.2.4 and Corollary II.2.5 make this coarse mechanism rigorous. Conjecture II.2.12 asks additionally for the sharp arithmetic constant. Its stronger trace form, supported by the same computations, is

$$\log \frac{\sigma_0(\operatorname{tr} g)}{\max_i |\sigma_i(\operatorname{tr} g)|} \geq \frac{\delta_7}{2}\nu(g).$$

This trace inequality implies the spectral inequality because $t \mapsto \operatorname{arcosh}(t/2) - \log t$ is increasing for $t > 2$. A sharp proof can be sought by refining the entry-defect automaton: the principal trace is a two-state positive path sum, while each nonintegral edge loses at least δ_7 at the worst conjugate. The only paths with zero loss are the diagonal paths through A_1, A_6; cyclically pairing those paths with the off-diagonal paths is the finite-state injection still to be constructed.

Route II.2.15 (A finite charged-matching formulation)**.** Let $s(w)$ be the number of occurrences of A_2, A_3, A_4, A_5 in the induced cyclic word. Break the remaining parabolic symbols into maximal alternating A_1, A_6 strings and choose a maximum disjoint matching inside each string. If $m(w)$ is the number of matched parabolic pairs, then the elementary parity inequality

$$s(w) + m(w) \geq \frac{1}{2}\nu(w) \tag{665}$$

holds: every odd parabolic string loses one half under matching, and the number of such strings is at most the number of separating nonparabolic symbols.

The sharp relative gap would follow from the stronger trace partition-function inequality

$$\log \frac{\sigma_0(\operatorname{tr} w)}{\max_i |\sigma_i(\operatorname{tr} w)|} \geq \delta_7\big(s(w) + m(w)\big). \tag{666}$$

Theorem II.2.7 proves (666) when $s(w) = 0$. Exact exhaustion through length ten also verifies the stronger inequality

$$\log \frac{\sigma_0(\operatorname{tr} w)}{\max_i |\sigma_i(\operatorname{tr} w)|} \geq s(w)\delta_7$$

on every tested positive-character word; equality occurs on $A_2 A_1^N$. This initially suggested a finite two-state path-switching assertion: charge every nonparabolic tile once and every chosen A_1A_6 pair once. The only subunit edge ratios are the $(2,1)$-entry of A_2 and the $(1,2)$-entry of A_5; cyclicity forces them to occur in balanced transition paths. Experiment II.2.20 below shows that this sufficient strengthening is false: it double-counts charge where a charged tile meets a matched parabolic turn. The correct target charges one half per induced symbol, not both full charges simultaneously.

Experiment II.2.16 (The active-macro chamber dichotomy)**.** Pair an alternating parabolic string greedily and replace each pair by one of the macros

$$B_+ = A_1A_6, \qquad B_- = A_6A_1.$$

Together with A_2, A_3, A_4, A_5, these form a six-symbol active alphabet. The discarded stronger proposal assigned every symbol one unit of charge. Consecutive macro symbols in one parabolic string must have the same sign; the sign may change only across a charged letter.

The program `active_macro_gap_tests.cpp` numerically exhausts this admissible macro language. The unrestricted trace inequality is false: the shortest counterexample is

$$B_+A_5B_-A_2,$$

whose normalized trace-charge margin is $-0.104904\ldots$. Its character, however, is $-0.661635\ldots$. Exhaustion through active length eleven gives

length	1	2	3	4	5	6	7	8	9	10	11
failures	0	0	0	4	0	0	0	4	18	0	0

and every failure has negative character. The largest character among all failures is $-0.661635\ldots$; repeated copies of the shortest counterexample make it more negative.

In fact the negative-character certificate can be made algebraic. Every failure in the exhaustion satisfies

$$\max\{1,|T_1|\}\max\{1,|T_2|\} < T_0 - 1 < \rho(T_0), \qquad T_i = \sigma_i(\operatorname{tr} w).$$

Since $\rho(T_i) \le \max\{1,|T_i|\}$, this directly gives $\rho(T_1)\rho(T_2) < \rho(T_0)$, hence $\chi < 0$. The largest observed ratio $|T_1T_2|/(T_0-1)$ on a failure is only $0.630045\ldots$. Thus the second branch of the desired finite certificate can be expressed using algebraic trace products, without inverse hyperbolic functions.

At unit run parameters the data therefore suggested the chamber alternative

$$\log \frac{\sigma_0(\operatorname{tr} w)}{\max_i |\sigma_i(\operatorname{tr} w)|} \ge \delta_7\{s(w)+m(w)\} \quad \text{or} \quad \chi(w) \le 0. \tag{667}$$

The family in Experiment II.2.20 disproves even this chamber alternative once a run parameter is allowed to grow. Thus the unit-parameter calculation is retained only to identify the primitive overcharged cycle; it is not evidence for an all-parameter theorem.

Route II.2.17 (The $p = 11$ test rules out a uniform scalar subaction)**.** The scalar Hodge defect is an invariant-measure average, not a uniform period inequality. This matters before attempting to upgrade the negative heptagonal average by compactness or a calibrated-subaction argument.

For

$$W_8 = A_1^8A_6A_5A_6A_1A_2$$

one has $\operatorname{tr} W_8(x) = 72x^4 + 2$. If $u_i = 36\sigma_i(\lambda)^4 + 1$, elimination gives

$$\prod_i (U - u_i) = U^3 - 471U^2 + 34635U - 80821$$

and hence

$$\sum_i u_i^2 - 2u_0u_1u_2 - 1 = -9072.$$

Also

$$u_0^2 - u_0u_1u_2 = 36(1014\lambda^2 + 686\lambda - 2751) > 0.$$

It follows exactly that

$$\operatorname{arcosh}(u_1) + \operatorname{arcosh}(u_2) > \operatorname{arcosh}(u_0),$$

so the scalar period is

$$\chi_{\mathrm{sc}}(W_8) = 0.0638765195\ldots > 0.$$

Therefore an inequality $\varphi + v - v \circ T \leq 0$ on the full compactly induced geodesic coding is impossible: summing it around this orbit gives the opposite period inequality. Parabolic acceleration does not remove the obstruction because W_8 is already a finite hyperbolic period.

This failed route identifies, rather than merely relocates, the missing mathematics. The fixed-point discriminant has norm

$$2^{15}3^6 \cdot 43 \cdot 337,$$

so the two endpoints of W_8 are quadratic over K_7, not K_7-rational. Thus the trace-field restriction removes the exact positive periodic orbit, but it does not define a compact subsystem: K_7-points can approximate its endpoint and shadow it for an arbitrarily long finite interval. A proof must charge the height needed for repeated shadowing. This is precisely why the stopped problem is Arakelov-arithmetic and cannot be solved by a continuous Hodge transfer on the real natural extension alone.

The $p = 11$ comparison now has a clean interpretation. At $p = 7$, positive scalar periods exist but are exceptional relative to the Bowen–Margulis measure; the natural scalar average is $-1/5$. At $p = 11$, the natural scalar average is already $7/9$. Hence the prime-family phase transition concerns which invariant regime is typical, while the heptagonal theorem asks for arithmetic exclusion of the exceptional positive regimes along a trace-field rational ray.

Route II.2.18 (The unifying conjecture suggested by $7, 9, 11$)**.** The least conjectural statement common to the three experiments is an Oseledets alternative for trace-field endpoints. For odd q, every $x \in \mathbf{P}^1(K_q)$ should either terminate, become eventually periodic, or have, along compact returns, the uniformizing denominator exponent -1 and the outer exponents

$$\chi_m = \frac{q - 2m}{q - 2}.$$

The last alternative makes the scalar integral endpoint tensor have exponent

$$-1 + \sum_{m \neq 1} \chi_m = -\Delta_q.$$

This is strictly weaker than full Bowen–Margulis genericity: it asks for only finitely many logarithmic growth rates, not the limiting frequency of every digit word.

If $\Delta_q > 0$, integrality excludes the third alternative and gives termination or repetition. If $\Delta_q < 0$, integrality is compatible with it; a separate existence assertion predicts an aperiodic algebraic ray, with $q = 11, x = 2$ the first candidate. The exact $q = 9$ closure data and the four measured $q = 11$ slopes support this statement. The positive-period word W_8 at $q = 7$ shows that the conjecture is genuinely arithmetic: it fails on the full real symbolic system and can hold only because the endpoint is restricted to K_q.

Lemma II.2.19 (The primitive bad macro cycle is outside the Hodge chamber)**.** *Put*

$$H = (A_1A_6)A_5(A_6A_1)A_2.$$

Then $\chi(H) < 0$, and $\chi(H^k) = k\chi(H) < 0$ for every $k \geq 1$. The same holds for cyclic rotations, transposes, and the reflection $A_j \leftrightarrow A_{7-j}$.

Proof. Direct multiplication, using only the six displayed generator matrices, gives the trace polynomial

$$\operatorname{tr} H(x) = 16x^4 + 2.$$

For the three conjugates $x_i = \sigma_i(\lambda)$, put

$$u_i = \frac{1}{2}\operatorname{tr} H(x_i) = 8x_i^4 + 1, \qquad a_i = \operatorname{arcosh}(u_i).$$

All $u_i > 1$. Elimination with $x^3 - x^2 - 2x + 1 = 0$ gives the exact minimal polynomial

$$u^3 - 107u^2 + 1875u - 2281.$$

Consequently

$$u_0^2 + u_1^2 + u_2^2 - 2u_0u_1u_2 - 1 = 107^2 - 2(1875) - 2(2281) - 1 = 3136 > 0.$$

Moreover $u_0 - u_1u_2 > 0$: since $u_0u_1u_2 = 2281$, reduction at the central root gives

$$u_0^2 - 2281 = 8(230\lambda^2 + 154\lambda - 399) > 0$$

on $9/5 < \lambda < 181/100$. The preceding two inequalities imply

$$u_0 > u_1u_2 + \sqrt{(u_1^2 - 1)(u_2^2 - 1)} = \cosh(a_1 + a_2).$$

Thus $a_0 > a_1 + a_2$, or $\chi(H) = -a_0 + a_1 + a_2 < 0$. Translation exponents multiply under positive powers, proving the assertion for H^k. The listed symmetries preserve the three translation exponents. □

Experiment II.2.20 (The full charged–matching strengthening is false)**.** For $N \geq 1$, put

$$W_N = A_1^N A_6 A_5 A_6 A_1 A_2.$$

Direct multiplication gives the unusually simple exact trace polynomial

$$\operatorname{tr} W_N(x) = 8(N+1)x^4 + 2. \tag{668}$$

The induced word has $s = 2$ charged letters, four parabolic symbols, $m = 2$ matched pairs, and $\nu = 6$. At $N = 8$, exact interval evaluation gives

$$\chi(W_8) = 0.0638765\ldots > 0, \qquad \frac{\delta(W_8)}{\delta_7} = 3.9762015\ldots < 4 = s + m.$$

The trace-ratio quotient is likewise

$$\frac{1}{\delta_7}\log\frac{\sigma_0(\operatorname{tr} W_8)}{\max_i |\sigma_i(\operatorname{tr} W_8)|} = 3.9761186\ldots < 4.$$

Thus both (666) and the chamber alternative (667) are false. As $N \to \infty$, the quotient tends to 4 from below while $\chi(W_N) \to +\infty$, so this is a structural overcount rather than a small-word anomaly.

It does not contradict Conjecture II.2.12: that conjecture asks only for

$$\frac{\nu(W_N)}{2}\delta_7 = 3\delta_7.$$

In general, if u is the number of unmatched parabolic symbols, then

$$\nu = s + 2m + u, \qquad \frac{\nu}{2} = m + \frac{s+u}{2}.$$

The corrected finite target is therefore to charge each matched macro once and each charged or unmatched symbol one half:

$$\log \frac{\sigma_0(\operatorname{tr} w)}{\max_i |\sigma_i(\operatorname{tr} w)|} \geq \delta_7 \left(m(w) + \frac{s(w) + u(w)}{2} \right). \tag{669}$$

On the unit-parameter active alphabet, `active_macro_gap_tests.cpp` finds no failure of this corrected weight through length ten; the least margin is already positive at length one. This corrected weighting, not (666), is the finite-state formulation of the relative fusion conjecture.

Experiment II.2.21 (Exact residual-block certificate and its finite exception automaton)**.** Put $s_0 = \sqrt{R}$. Its minimal polynomial and the two embedding parameters are

$$s_0^6 - 4s_0^4 + 3s_0^2 + 1 = 0, \qquad L = -s_0^4 + 2s_0^2 + 1, \qquad t = -s_0^4 + 3s_0^2 - 1.$$

Thus every half-integral gauge factor is an integral power of s_0, and all block comparisons can be made exactly in $\mathbf{Z}[s_0]/(s_0^6 - 4s_0^4 + 3s_0^2 + 1)$.

The program `relative_fusion_block_certificate.py` expands each alternating parabolic run exponent as $n_i = 1 + y_i$. For a block containing p parabolic symbols followed by a charged word q of length ℓ, beginning in orientation e and ending in mode v, it checks every coefficient and every entry of

$$B(L)_{ij} - R^{(p+\ell)/2 + V_v(j) - V_e(i)} B(t)_{ij}, \qquad V_0 = (0, \tfrac{1}{2}), \quad V_1 = (\tfrac{1}{2}, 0).$$

Signs are certified by rational interval arithmetic at the isolated root $s_0 \in [961/800, 481/400]$; floating point is not used.

For $1 \leq p \leq 10$, $1 \leq \ell \leq 4$, all four charged letters, and both initial and terminal modes, this is 13600 exact transitions. Exactly 22 fail. Every failure uses only the weak letters A_2, A_5; there are no failures for $\ell = 3, 4$, and the complete list is

$e = 0$	$(1,2,0), (1,2,1), (1,5,0), (1,22,1), (1,25,0),$ $(2,2,1), (2,5,0), (3,2,1), (4,5,0), (5,2,1), (6,5,0)$
$e = 1$	$(1,2,1), (1,5,0), (1,5,1), (1,52,1), (1,55,0),$ $(2,2,1), (2,5,0), (3,5,0), (4,2,1), (5,5,0), (6,2,1).$

Here (p, q, v) records parabolic length, charged word, and terminal mode. The same program checks all 242 compatible ordered pairs of exceptions. Exactly 30 pairs still fail entrywise. In every such pair the first block is one of

$$(0,1,A_2,0), \quad (0,1,A_2,1), \quad (1,1,A_5,0), \quad (1,1,A_5,1).$$

The pure ten-parabolic block in either orientation has a nonnegative coefficient certificate of weight five, with 2304 coefficients. Consequently arbitrary parabolic strings can be peeled ten symbols at a time and the only remaining second-outer issue is the displayed finite exception language.

The endpoint failure is not a failure of the cyclic trace inequality. The same algebraic program now compares cyclic traces coefficient by coefficient in $\mathbf{Z}[s_0]$. It gives

cycle length	compatible cycles	coefficients	negative
1	12	180	0
2	122	20808	0
3	1332	2702160	0.

The zeros are repetitions of the sharp one-cusp rays. Thus all of these cyclic statements are exact; no floating-point comparison remains in the length-three audit. The calculation indicates that the remaining all-cycle lemma should use complementary trace paths, or add the full one-cusp gauges $\mathrm{diag}(1, R)$ and $\mathrm{diag}(R, 1)$, rather than enlarge the half-gauge search depth.

Route II.2.22 (Trace-parameter monotonicity: a global replacement for the incompatible gauges)**.** Put $t = |\sigma_2(\lambda)| = \lambda^2 - 2$, and define $A_j(x)$ by replacing λ by x and α by $x^2 - 1$ in the six positive generators. For a word $w = j_1 \cdots j_n$, put

$$F_w(x) = \mathrm{tr}\big(A_{j_1}(x) \cdots A_{j_n}(x)\big), \qquad s(w) = \#\{k : j_k \in \{2, 3, 4, 5\}\}.$$

Entrywise,

$$A_j(\lambda) = \sigma_0(A_j), \qquad A_j(t) = |\sigma_2(A_j)|.$$

Consequently the following purely real-polynomial inequality would give the complete σ_2-part of (666):

$$xF_w'(x) - s(w)F_w(x) \geq 0 \quad (t \leq x \leq \lambda). \tag{670}$$

Indeed, integration of (670) gives

$$\frac{F_w(\lambda)}{F_w(t)} \geq \left(\frac{\lambda}{t}\right)^{s(w)} = R_7^{s(w)},$$

and the outer trace is bounded by the positive path sum $F_w(t)$.

This formulation exposes exactly why the two diagonal gauges nearly worked. If $H_j = xA_j' - s(A_j)A_j$, where $s(A_j) = 0$ for $j = 1, 6$ and 1 otherwise, then H_1, H_3, H_4, H_6 are entrywise nonnegative. The only negative entries in the entire derivative are

$$(H_2)_{21} = -1, \qquad (H_5)_{12} = -1.$$

They are the same two subunit edges found by the charged-gauge audit. Thus (670) is precisely a finite two-state path-switching inequality. Pure all-bad cycles alternate A_2, A_5; their complete path sum is a Chebyshev polynomial, so the alternating-shear argument proves their nonnegativity once $x^4 \geq 2$. The remaining combinatorial step is to splice these Chebyshev cycles into an arbitrary trace path without losing the positive marked edges.

There is a companion outer-place issue. At σ_1, the entry α changes sign, so its absolute matrix is not $A_j(x)$ on the same positive one-parameter arc. It is enough to prove either

$$\sigma_0(\mathrm{tr}\, w) \geq R_7^{s(w)}\, \mathrm{tr} \prod_k |\sigma_1(A_{j_k})|$$

or the weaker outer-order dichotomy

$$b(w) > c(w) \quad \Longrightarrow \quad \chi(w) \leq 0.$$

Exact exhaustion finds (II.2.22) throughout the tested positive semigroup. Theorem II.2.33 now proves the stronger companion (II.2.22) for every positive word. Hence the endpoint form needed from the second outer place is now supplied by Theorem II.2.42. The stronger interval differential statement (670) remains a useful description of the polynomial cone, but it is no longer a logical obligation for either outer embedding.

Experiment II.2.23 (Exact Bernstein certificate for trace monotonicity)**.** The new program `trace_parameter_monotonicity_tests.py` constructs $F_w \in \mathbf{Z}[x]$ exactly and certifies

$$xF_w' - s(w)F_w \geq 0 \qquad (6/5 \leq x \leq 2)$$

by rational Bernstein coefficients with exact recursive subdivision. This interval strictly contains $[|\sigma_2(\lambda)|, \lambda]$. Exhaustion through length seven gave

$\lvert w\rvert$	1	2	3	4	5	6	7
charged words tested	4	32	208	1280	7744	46592	279808
uncertified	0	0	0	0	0	0	0
identically zero	2	4	6	8	10	12	14.

The zero cases are exactly the cyclic words $A_2A_1^{n-1}$ and $A_5A_6^{n-1}$, including their cyclic rotations. These are precisely the sharp rays on which the ratio is constant. The certificate is much stronger than evaluation at the two cubic embeddings: it verifies monotonicity on a rational interval and therefore identifies a plausible inductive cone in $\mathbf{Z}[x]$.

In fact recursive subdivision was never used: every nonzero tested gap already had nonnegative Bernstein coefficients on the whole interval. Under the projective change

$$z = \frac{5x-6}{10-5x}, \qquad x = \frac{6+10z}{5+5z},$$

this is ordinary coefficientwise positivity on $z \geq 0$. Clearing the common denominator sends the three possible positive entries to

$$25(1+z)^2, \qquad 5(1+z)(6+10z), \qquad 11+70z+75z^2,$$

all in $\mathbf{Z}_{\geq 0}[z]$. Thus a proof can be organized as a finite semiring induction after this change of parameter. An additional exact-Bernstein random test of 2000 words of lengths up to 50 also found no uncertified polynomial; its 23 zero cases were again sharp one-turn rays.

Experiment II.2.24 (Infinitesimal hyperbolic-length certificate)**.** For a positive word w, let $F_w(x)$ be its trace polynomial and let $\nu(w)$ be its cyclic induced complexity. The differential inequality

$$\frac{xF_w'(x)}{\sqrt{F_w(x)^2-4}} \geq \frac{\nu(w)}{2} \tag{671}$$

is exactly

$$\frac{d}{d\log x}\operatorname{arcosh}\frac{F_w(x)}{2} \geq \frac{\nu(w)}{2}.$$

It would integrate from $t = |\sigma_2(\lambda)|$ to λ to give the sharp second-outer fusion comparison. After squaring, its polynomial certificate is

$$4x^2F_w'(x)^2 - \nu(w)^2(F_w(x)^2-4) \geq 0. \tag{672}$$

The program `hyperbolic_length_derivative_tests.py` constructs this polynomial exactly and checks rational Bernstein coefficients on $[6/5, 2]$. Exhaustion gives

$\lvert w\rvert$	1	2	3	4	5	6	7
tested	4	34	214	1294	7774	46654	279934
failed	0	0	0	0	0	0	0.

Every certificate is positive on the full interval without subdivision. For lengths at least three, the least Bernstein coefficient is 16, attained on $A_1^{n-1}A_2$; at lengths one and two the least coefficients are $532/25$ and $1152/625$, respectively.

The stronger unsquared inequality $2xF_w' \geq \nu(w)F_w$ has only the primitive pure two-cusp failure A_1A_6 in exhaustion through length five. That primitive is repaired by the $F_w^2 - 4$ term in (672), and the full pure two-cusp language is already controlled by the alternating-continuant theorem. This suggests a proof split: a finite path-switching potential for every word containing a charged letter, followed by the existing continuant argument on the pure parabolic language. The experiment is an exact discovery certificate, not yet an all-length proof.

Experiment II.2.25 (Coefficientwise induced-derivative semiring)**.** Put $y = x^2 - 1$. Every trace polynomial has a unique decomposition

$$F_w(x) = P_{w,0}(y) + xP_{w,1}(y).$$

The exact program `trace_induced_derivative_semiring_tests.py` checks the strictly stronger assertion

$$2xF_w'(x) - \nu(w)F_w(x) = G_{w,0}(y) + xG_{w,1}(y), \qquad G_{w,0}, G_{w,1} \in \mathbf{Z}_{\geq 0}[y], \tag{673}$$

for every word containing at least one charged letter. Exact exhaustion gave

$\lvert w\rvert$	1	2	3	4	5	6	7
tested	4	32	208	1280	7744	46592	279808
failures	0	0	0	0	0	0	0
identically zero	0	4	6	8	10	12	14.

The zero words are the sharp one-turn families $A_2A_1^{n-1}$, $A_5A_6^{n-1}$, and their cyclic rotations. Every nonzero coefficient appearing from length three onward is at least one.

This form exposes a multiplicative mechanism absent from the squared Bernstein certificate. Define

$$\mathcal{L}_q(f) = 2x\frac{df}{dx} - qf.$$

Then

$$\mathcal{L}_{q+r}(fg) = \mathcal{L}_q(f)g + f\mathcal{L}_r(g). \tag{674}$$

The semiring $\mathbf{Z}_{\geq 0}[y] \oplus x\mathbf{Z}_{\geq 0}[y]$ is closed under multiplication, so any macro decomposition whose factors satisfy their assigned $\mathcal{L}_q$-inequalities proves the sum of their charges.

A one-path or one-entry proof is false: in A_2A_5, one complementary closed path uses the two unit entries and has zero derivative cost, whereas its complement carries the compensating y^2 term. This is exactly the pairing already used in Theorem II.2.37. The remaining all-length lemma can therefore be stated finitely: prove that the complementary-pair cone defined by (673) is preserved when an alternating A_1, A_6 run, with arbitrary positive exponents, is inserted between two charged chains. Equation (674) then adds the induced charges.

Theorem II.2.26 (Finite-macro derivative theorem for every charged open chain)**.** *Let Q be a word of length $s \geq 2$ in A_2, A_3, A_4, A_5. For every matrix entry,*

$$\mathcal{L}_s(Q_{ij}) = 2x\frac{dQ_{ij}}{dx} - sQ_{ij} \in \mathbf{Z}_{\geq 0}[y] \oplus x\mathbf{Z}_{\geq 0}[y], \qquad y = x^2 - 1. \tag{675}$$

For a one-letter charged cyclic word the same assertion holds for its trace. Consequently every charged cyclic word C of length s satisfies

$$2xF_C'(x) - sF_C(x) \geq 0 \qquad (x \geq 1).$$

Proof. The exact certificate `charged_derivative_macro_certificate.py` multiplies all $4^2 = 16$ two-letter and $4^3 = 64$ three-letter macros. It decomposes each of their four entries in $\mathbf{Z}[y] \oplus x\mathbf{Z}[y]$, and verifies coefficientwise positivity of $\mathcal{L}_2$, respectively $\mathcal{L}_3$. The certificate checks 1,362 positive integer coefficients; the least is one, and no coefficient is negative.

Every integer $s \geq 2$ is a sum of 2's and 3's. Split Q into consecutive macros of those lengths. Equation (674), entrywise matrix multiplication, and closure of $\mathbf{Z}_{\geq 0}[y] \oplus x\mathbf{Z}_{\geq 0}[y]$ under addition and multiplication prove (675). For $s = 1$, the four traces are $2x$ or $2y$, and the assertion is immediate. Taking the sum of the two diagonal entries proves the cyclic statement. □

Remark II.2.27 (Exact remaining fusion interface). Theorem II.2.26 proves the charged core with no interval restriction and no asymptotic estimate. The analogous two-letter certificate after adjoining A_1, A_6 has exactly four failures:

$$A_1A_2, \quad A_2A_1, \quad A_5A_6, \quad A_6A_5.$$

They are the sharp one-sided languages already handled by the continuant theorem. Nevertheless arbitrary mixed words do not always admit a cyclic tiling by good two- and three-letter macros; global complementary-path pairing is still required at this interface. Thus the finite-macro theorem is a complete simplification of the charged subproblem, not by itself a proof of (671) for all six letters.

Lemma II.2.28 (Invariant Perron-slope interval for a matrix and its transpose)**.** *Let $x > 1$, put $a = x^2 - 1$, and set*

$$M = \begin{pmatrix} x & a \\ 1 & x \end{pmatrix}, \qquad M^\dagger = M^t.$$

If P is a nonempty product of $M, M^\dagger$, then

$$\min(a, a^{-1}) \leq \frac{P_{12}}{P_{21}} \leq \max(a, a^{-1}). \tag{676}$$

Proof. On the positive projective coordinate z, the two generators act by

$$m(z) = \frac{xz + a}{z + x}, \qquad m^\dagger(z) = \frac{xz + 1}{az + x}.$$

Their positive fixed points are respectively $\sqrt{a}$ and $1/\sqrt{a}$, and

$$m(z) - z = \frac{a - z^2}{z + x}, \qquad m^\dagger(z) - z = \frac{1 - az^2}{az + x}.$$

Both maps are increasing. Hence both preserve the closed interval

$$I_a = [\min(\sqrt{a}, a^{-1/2}), \max(\sqrt{a}, a^{-1/2})]$$

and move every point in it toward their respective fixed point. The unique positive Perron fixed slope z_P of every product P therefore belongs to I_a. The transpose P^t is again a product of $M, M^\dagger$, so its Perron fixed slope z_{P^t} also belongs to I_a.

Write $P = (p_{ij})$ and let ρ be its Perron eigenvalue. Applied to the right Perron vectors $(z_P, 1)^t$ and $(z_{P^t}, 1)^t$, the first rows of P and P^t give

$$p_{12} = (\rho - p_{11})z_P, \qquad p_{21} = (\rho - p_{11})z_{P^t}.$$

Thus $p_{12}/p_{21} = z_P/z_{P^t}$. Taking the quotient of two numbers in I_a proves (676). □

Theorem II.2.29 (Trace monotonicity on the full A_2, A_5 semigroup)**.** *Let w be any word of length $s \geq 1$ in A_2, A_5, and define*

$$F_w(x) = \operatorname{tr} \prod_{j \in w} A_j(x).$$

For every $x > 1$ with $x^4 \geq 2$,

$$xF_w'(x) - sF_w(x) \geq 0. \tag{677}$$

Consequently

$$\frac{F_w(x_0)}{F_w(x_1)} \geq \left(\frac{x_0}{x_1}\right)^s \qquad (x_0 \geq x_1 \geq 2^{1/4}).$$

Proof. Put $M = A_2(x)$, so $A_5(x) = M^t$, and note

$$xM'(x) - M(x) = \begin{pmatrix} 0 & x^2+1 \\ -1 & 0 \end{pmatrix}, \qquad x(M^t)'(x) - M^t(x) = \begin{pmatrix} 0 & -1 \\ x^2+1 & 0 \end{pmatrix}.$$

Differentiate the trace of the product and cut it cyclically at the marked factor. If $Q = (q_{ij})$ is the product of all the other factors, the marked contribution is respectively

$$(x^2+1)q_{21} - q_{12}, \qquad (x^2+1)q_{12} - q_{21}.$$

For $s = 1$, both are zero. For $s > 1$, Lemma II.2.28 applies to Q. Since

$$\max(a, a^{-1}) \leq x^2 + 1$$

is automatic for $a = x^2 - 1 \geq 1$, while for $a \leq 1$ it is equivalent to $x^4 \geq 2$, both displayed marked contributions are nonnegative. Their sum is exactly $xF_w' - sF_w$, proving (677). Integrating $(\log F_w)' \geq s/x$ gives the last assertion. □

Remark II.2.30. This proof resolves the all-bad alternating cycles without pairing one closed path with its state complement. The complement pairing fails for long cycles because its weight can be exponentially smaller. The Perron-slope interval instead sums the entire two-state partition function before comparing the two off-diagonal directions. It is therefore the appropriate macro-certificate to splice between marked parabolic transitions.

Theorem II.2.31 (The two-charge parabolic splice)**.** *Let P, Q be arbitrary words in $A_1(x), A_6(x)$, including the empty word, and put*

$$F_{P,Q}(x) = \operatorname{tr}\big(A_2(x)P(x)A_5(x)Q(x)\big).$$

If $x > 1$ and $x^4 \geq 2$, then

$$xF_{P,Q}'(x) - 2F_{P,Q}(x) \geq 0. \tag{678}$$

The same assertion holds with A_2, A_5 interchanged.

Proof. Write $M = A_2(x)$, so $A_5(x) = M^t$. Every entry of a word in

$$A_1 = \begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix}, \qquad A_6 = \begin{pmatrix} 1 & 0 \\ x & 1 \end{pmatrix}$$

is a polynomial with nonnegative integral coefficients. Its diagonal entries contain only even powers of x, with constant term 1, and its off-diagonal entries contain only odd powers, hence have degree at least one. Expand P, Q coefficientwise in the four matrix units. For a term $cx^d E_{ij}$ from P and a term $c'x^e E_{kl}$ from Q, put $D = d + e$. Its contribution to the left side of (678) is $cc'x^D$ times

$$\mathcal{K}_{ij,kl}(D) = x\frac{d}{dx}\operatorname{tr}(ME_{ij}M^tE_{kl}) + (D-2)\operatorname{tr}(ME_{ij}M^tE_{kl}).$$

The sixteen kernels are

$(ij)\backslash(kl)$	00	01	10	11
00	Dx^2	$x(D-1)$	$x(D-1)$	$D-2$
01	$x\{D(x^2-1)+x^2+1\}$	$D(x^2-1)+2$	Dx^2	$x(D-1)$
10	$x\{D(x^2-1)+x^2+1\}$	Dx^2	$D(x^2-1)+2$	$x(D-1)$
11	$(x^2-1)\{D(x^2-1)+2x^2+2\}$	$x\{D(x^2-1)+x^2+1\}$	$x\{D(x^2-1)+x^2+1\}$	Dx^2.

Degree parity makes every displayed contribution nonnegative except possibly $\mathcal{K}_{00,11}(0) = -2$. Indeed an off-diagonal matrix unit forces $D \geq 1$, whereas diagonal units force an even $D \geq 0$. The exceptional term occurs exactly once, from the constant E_{00} in P and the constant E_{11} in Q. The constant E_{11} in P and E_{00} in Q contributes

$$\mathcal{K}_{11,00}(0) = 2(x^4 - 1).$$

Their sum is $2(x^4-2) \geq 0$. All remaining terms are nonnegative, which proves (678). Transposition and cyclicity give the assertion with A_2, A_5 interchanged. □

Route II.2.32 (Why a single infinitesimal gauge cannot finish the core)**.** A natural way to make the path switching automatic is to seek one matrix $H(x)$ such that

$$xA_j' - s(A_j)A_j + HA_j - A_jH \geq_{\text{ent}} 0 \qquad (j = 1, 2, 5, 6).$$

The commutators disappear after differentiating the trace of a product, so (II.2.32) would prove the core inequality in one line. This possibility is, however, exactly impossible, even pointwise. Modulo a scalar write $H = \left(\begin{smallmatrix} d & b \\ c & -d \end{smallmatrix}\right)$. For A_1, the two diagonal entries in (II.2.32) are $-cx, cx$, hence $c = 0$. For A_6 they are $bx, -bx$, hence $b = 0$. The lower-left entry for A_2 is then $-2d - 1$, while its A_1 upper-right companion forces $2d + 1 \geq 0$; thus $d = -1/2$. Conversely the upper-right entry for A_5, together with the A_6 lower-left entry, forces $d = +1/2$. This contradiction is the dual certificate for the two bad edges. A proof must therefore switch gauges nonlocally; Theorem II.2.31 performs the first complete such switch, and shows that the residual slack is exactly $2(x^4 - 2)$.

Theorem II.2.33 (First-outer comparison for all six generators)**.** *Let W be a cyclic word in $A_1, \ldots, A_6$, and let $s \geq 1$ be its number of A_2, A_3, A_4, A_5 letters. Put*

$$L = 2\cos(\pi/7), \qquad y = 2\cos(3\pi/7), \qquad R = \frac{L}{|2\cos(5\pi/7)|} = L^2 - L = 1 + y.$$

If $E_1(A_j) = |\sigma_1(A_j)|$ entrywise, then

$$\operatorname{tr} W(L) \geq R^s \operatorname{tr} \prod_{j \in W} E_1(A_j). \tag{679}$$

Consequently the actual first-outer trace satisfies the same bound in absolute value. In particular this proves the companion obligation (II.2.22) for every positive word.

Proof. Set $A = L^2 - 1$ and $\beta = 1 - y^2$. The four central charged matrices have entries among $1, L, A$, and their absolute outer companions have entries among $1, y, \beta$. In particular, the pair A_2, A_5 is

$$M_0 = \begin{pmatrix} L & A \\ 1 & L \end{pmatrix}, \quad M_0^t, \qquad M_1 = \begin{pmatrix} y & \beta \\ 1 & y \end{pmatrix}, \quad M_1^t,$$

while A_3, A_4 and their outer companions use only L, A and y, β, respectively. First consider corresponding open charged products Q_0, Q_1 of length r. At length one every entry of Q_0 is at least the corresponding entry of Q_1. If m_r is the least of the four entries of Q_0, and M_r the greatest entry of Q_1, then the least column sum among all four central generators is $L + 1$, whereas the greatest column sum among all four outer generators is 2β. Indeed the other relevant column sums are

$$1 + y = R, \qquad y + \beta < R, \qquad L + A > L + 1, \qquad 2A > L + 1.$$

It follows inductively that

$$m_r \geq (L + 1)^{r-1}, \qquad M_r \leq (2\beta)^{r-1}.$$

The cubic relation $L^3 - L^2 - 2L + 1 = 0$ gives

$$L + 1 - 2\beta R = (L - 1)(2L - 3) > 0.$$

Hence every corresponding open-chain entry satisfies

$$(Q_0)_{ij} \geq R^{r-1} (Q_1)_{ij}. \tag{680}$$

We also need the extra factor for a charged cyclic word with no parabolic cut. A central product of length r has every row sum at least $(L + 1)^r$, so its Perron root, and therefore its trace, is at least $(L + 1)^r$. The outer product has matrix one-norm at most $(2\beta)^r$, hence trace at most $2(2\beta)^r$. For $r \geq 4$ the desired ratio follows from

$$(L + 1)^4 - 2(2\beta R)^4 = -9747L^2 + 21901L - 7812 > 0.$$

The 4+16+64 charged cyclic words of lengths one, two, and three are the finite exact base. Reducing their trace differences modulo $L^3 - L^2 - 2L + 1$ and taking Bernstein coefficients on the rational isolating interval $[9/5, 181/100]$, the smallest coefficient is $54/25$, attained at the length-one word A_3 (and its transpose). The exact calculation is reproduced by `first_outer_charged_certificate.py`. The fourth-power quadratic displayed above is positive on the same interval, so the row-sum induction covers every remaining length.

Now write the cyclic word as $C_1 P_1 \cdots C_s P_s$, with charged C_i and $P_i \in \langle A_1, A_6 \rangle$. Every entry of $P_i(x)$ is a polynomial in x with nonnegative integral coefficients and constant matrix I. Expand $P_i = I + (P_i - I)$ and then expand the latter in monomial matrix units. The all-identity term is the charged cyclic case just proved. A term with $k \geq 1$ nonidentity gaps has total parabolic degree

$D \geq k$; the matrix units cut its charged part into k open chains of total length s. By (680), its central-to-outer ratio is at least

$$\left(\frac{L}{y}\right)^{D} R^{s-k} \geq R^{s},$$

because $L/y > R$. Termwise summation proves (679). Finally the triangle inequality bounds the absolute trace of the signed outer product by the trace of its entrywise-absolute product. □

Lemma II.2.34 (A quotient-recursion comparison)**.** *Let $c \geq 1$, let $A \geq cB > 2c$, and let two positive sequences u_k, v_k satisfy*

$$u_{k+1} = Au_k - u_{k-1}, \qquad v_{k+1} = Bv_k - v_{k-1}.$$

If

$$\frac{u_1}{u_0} \geq c\frac{v_1}{v_0},$$

then, for every $k \geq 1$ for which the displayed quotients are positive,

$$\frac{u_k/u_{k-1}}{v_k/v_{k-1}} \geq c, \qquad \frac{u_k/u_0}{v_k/v_0} \geq c^k.$$

Proof. Put $r_k = u_k/u_{k-1}$ and $s_k = v_k/v_{k-1}$. The recurrences give

$$r_{k+1} = A - r_k^{-1}, \qquad s_{k+1} = B - s_k^{-1}.$$

If $r_k \geq cs_k$, then

$$r_{k+1} \geq cB - \frac{1}{cs_k} \geq cB - \frac{c}{s_k} = cs_{k+1}.$$

The first assertion follows by induction, and multiplying the quotient inequalities gives the second. □

Theorem II.2.35 (All-length second-outer comparison on an open charged chain)**.** *Put*

$$L = 2\cos(\pi/7), \quad t = L^2 - 2 = |\sigma_2(L)|, \quad R = L/t = L^2 - L,$$

and

$$M(x) = A_2(x) = \begin{pmatrix} x & x^2 - 1 \\ 1 & x \end{pmatrix}, \qquad N(x) = A_5(x) = M(x)^t.$$

For $\epsilon = +, -$, put

$$D_+ = \mathrm{diag}(1, R), \qquad D_- = \mathrm{diag}(R, 1).$$

If Q is any word of length $s \geq 1$ in M, N, and ϵ_-, ϵ_+ are respectively the signs of its first and last letters, then

$$Q(L) \geq_{\mathrm{ent}} R^s D_{\epsilon_-}^{-1} Q(t) D_{\epsilon_+}. \tag{681}$$

In particular,

$$Q(L) \geq_{\mathrm{ent}} R^{s-1} Q(t). \tag{682}$$

Proof. We first prove (681) for an alternating word. Transposition reduces to a word starting in M. Put

$$P(x) = M(x)N(x), \qquad \tau(x) = \operatorname{tr} P(x) = x^4 + 2.$$

An even alternating word is $P(x)^k$; an odd one is $P(x)^k M(x)$. Every entry of either family is positive and, as a function of k, satisfies the two-step recurrence with coefficient $\tau(x)$.

All input needed for comparing the recurrences is the following finite list. Reduction modulo $L^3 - L^2 - 2L + 1$ gives

$$\begin{aligned}
\tau(L) - R^2\tau(t) &= -2L^2 + 6L - 1 > 0,\\
P_{00}(L) - R^3 P_{00}(t) &= -3L^2 + 8L - 2 > 0,\\
P_{11}(L) - R P_{11}(t) &= -L^2 + L + 2 > 0,\\
P_{01}(L) - R^2 P_{01}(t) &= L > 0,\\
(L^2 - 1) - R^2(t^2 - 1) &= 2L^2 - 3L > 0.
\end{aligned} \tag{683}$$

Here $P_{01} = P_{10}$. For the odd family put $q_k(x) = (P(x)^k M(x))_{ij}$. The four first-quotient inequalities

$$\frac{q_1(L)}{q_0(L)} \geq R^2 \frac{q_1(t)}{q_0(t)}$$

reduce, after cross multiplication, to

$$2L^2 + 3L - 2, \qquad 0, \qquad 3L - 1, \qquad 2L^2 + 3L - 2$$

in the order $00, 01, 10, 11$. Thus all are nonnegative. Positivity of every displayed polynomial follows already from $9/5 < L < 181/100$. The script `second_outer_alternating_certificate.py` reproduces all nine exact reductions.

Apply Lemma II.2.34 with $c = R^2$, $A = \tau(L)$, and $B = \tau(t)$. For the diagonal entries of P^k, the second and third base inequalities give respectively

$$(P(L)^k)_{00} \geq R^{2k+1}(P(t)^k)_{00}, \qquad (P(L)^k)_{11} \geq R^{2k-1}(P(t)^k)_{11}.$$

For an off-diagonal entry use

$$(P^k)_{01} = U_{k-1}(\tau/2)P_{01};$$

the quotient recursion and the fourth base inequality give the factor R^{2k}. These four exponents are exactly those in

$$P(L)^k \geq R^{2k} D_+^{-1} P(t)^k D_-.$$

For $P^k M$, the four ratios at $k = 0$ are bounded below by

$$R, \qquad R^2, \qquad 1, \qquad R,$$

respectively; the only nontrivial one is the fifth base inequality. The four first-quotient calculations and the same recursion lemma add a factor R^2 at every step. Consequently

$$P(L)^k M(L) \geq R^{2k+1} D_+^{-1} P(t)^k M(t) D_+.$$

This proves the gauged assertion for every alternating block.

Now cut an arbitrary M, N-word between every two equal adjacent letters. Each resulting block alternates, and the last sign of one block equals the first sign of the next. On multiplying the block inequalities, all intermediate diagonal gauges cancel. This proves (681). Finally every entry of $D_{\epsilon_-}^{-1} Q(t) D_{\epsilon_+}$ is at least R^{-1} times the corresponding entry of $Q(t)$, which proves (682). □

Theorem II.2.36 (Second-outer comparison for the four-letter core)**.** *Let W be a cyclic word in A_1, A_2, A_5, A_6, and let $s \geq 1$ be its number of A_2, A_5 letters. Then*

$$\operatorname{tr} W(L) \geq R^s \operatorname{tr} \prod_{j \in W} |\sigma_2(A_j)| \geq R^s |\sigma_2(\operatorname{tr} W)|. \tag{684}$$

Proof. Write the word cyclically as $C_1 P_1 \cdots C_s P_s$, where $C_i \in \{A_2, A_5\}$ and $P_i \in \langle A_1, A_6 \rangle$. Every entry of $P_i(x)$ is a polynomial in x with nonnegative integral coefficients and constant matrix I. Expand each $P_i = I + (P_i - I)$, and expand every nonconstant part in monomial matrix units.

The all-identity term is a cyclic A_2, A_5-word and has central to outer ratio at least R^s by Theorem II.2.29. In a remaining term suppose $k \geq 1$ gaps contribute nonidentity monomials of total degree D. Then $D \geq k$, and the matrix units cut the charged cycle into k nonempty open chains of total length s. Equation (682) contributes R^{s-k}, while the monomials contribute $(L/t)^D = R^D$. The total ratio is therefore at least

$$R^{D+s-k} \geq R^s.$$

Termwise summation proves the first inequality. The second is the triangle inequality applied at the outer embedding. □

Theorem II.2.37 (Complementary-path comparison for every charged cyclic word)**.** *Let C be a cyclic word of length $s \geq 1$ in A_2, A_3, A_4, A_5. Then*

$$\operatorname{tr} C(L) \geq R^s \operatorname{tr} C(t). \tag{685}$$

The same argument gives, for every open charged word Q,

$$Q(L)_{ij} + Q(L)_{\bar{i}\bar{j}} \geq R^s \big(Q(t)_{ij} + Q(t)_{\bar{i}\bar{j}}\big), \qquad \bar{i} = 1 - i, \quad \bar{j} = 1 - j. \tag{686}$$

Proof. Write

$$a = L^2 - 1, \qquad b = t^2 - 1.$$

The cubic equation gives

$$L = Rt, \qquad a = Lt, \qquad a - R^2 = -(L-1)(L-2) > 0,$$

while $L, t, a > 1$ and $0 < b < 1$.

Expand the trace as a sum over cyclic paths in the two matrix states, and pair a path $u = (u_1, \ldots, u_s)$ with its state complement $\bar{u} = (1 - u_1, \ldots, 1 - u_s)$. At a diagonal step, the path and its complement use the two diagonal entries. For A_2, A_5 those entries are both L centrally and both t outside, so the step contributes exactly R. For A_3, A_4 they are both a centrally and both b outside, so it contributes at least R. We may therefore remove all d diagonal positions at a cost R^d.

At a flip, the unordered pair of path weights is

	central	outer
A_2, A_5	$\{a, 1\}$	$\{b, 1\}$
A_3, A_4	$\{a, L\}$	$\{b, t\}$.

Suppose there are w weak flips of the first kind and m strong flips of the second kind. Factor L^m from the sum of the two central path monomials and t^m from the corresponding outer sum. In the remaining central monomials the flip ratios are a at weak positions and $a/L = t$ at strong positions. If the first monomial is A, the complementary one is B, and

$$AB = a^w t^m.$$

Consequently

$$A + B \geq 2a^{w/2}t^{m/2} \geq 2R^w.$$

In the two outer monomials, the corresponding ratios are b and b/t, both in $(0, 1)$. Each outer monomial is therefore at most one, and their sum is at most two. Since $L^m = R^m t^m$, the central pair is at least R^{w+m} times the outer pair. Restoring the diagonal positions gives $R^{d+w+m} = R^s$.

Summing over complementary cyclic path pairs proves (685). For an open word, complement also complements both endpoints; the identical pairing proves (686). □

Route II.2.38 (The gauge table behind the A_3, A_4 interface)**.** Theorem II.2.36 is an endpoint theorem, so it is already sufficient for the second outer embedding; the stronger differential statement below is no longer needed for that core. It also shows precisely how to extend the proof to all six letters. With D_+, D_- as above, let $e_j(u, v)$ be the greatest real exponent for which

$$A_j(L) \geq_{\text{ent}} R^{e_j(u,v)} D_u^{-1} A_j(t) D_v.$$

The following integer lower bounds are exact:

	$v = +$	$v = -$
$A_2,\ u = +$	1	0
$A_2,\ u = -$	0	-1

	$v = +$	$v = -$
$A_3,\ u = +$	2	1
$A_3,\ u = -$	1	0

and, by transposition,

	$v = +$	$v = -$
$A_4,\ u = +$	0	1
$A_4,\ u = -$	1	2

	$v = +$	$v = -$
$A_5,\ u = +$	-1	0
$A_5,\ u = -$	0	1.

Thus A_3 has exactly one more unit of gauge slack than A_2, and A_4 exactly one more than A_5. Local multiplication alone loses the charge on a perfectly alternating A_2, A_5 skeleton; Theorem II.2.35 supplies precisely that missing alternating macro. Conversely, the only new bad local sandwiches after adjoining the strong letters are

$$A_2A_4A_2, \qquad A_5A_3A_5$$

and their overlaps. Direct algebraic endpoint evaluation gives more than three units of gauged slack on each sandwich. High-precision exhaustion of alternating sign blocks through length twelve finds the least normalized comparison when every strong letter is replaced by its weak partner $A_3 \mapsto A_2,\ A_4 \mapsto A_5$. The next two lemmas prove this strong-replacement assertion at every length. No parabolic exponent and no first-outer estimate enters their proof.

Lemma II.2.39 (Cross-embedding projective cone)**.** *Put $K_0 = L + 1$. Let Q be a nonempty word in A_2, A_3, A_4, A_5. For either row of Q, let z_L and z_t be the ratios second-entry/first-entry at L and t, respectively. Then*

$$K_0^{-1} \leq \frac{z_L}{z_t} \leq K_0. \tag{687}$$

The same assertion holds for the ratios bottom-entry/top-entry in either column.

Proof. For a positive matrix $A = (a_{ij})$, its action on a column slope is

$$f_A(z) = \frac{a_{10} + a_{11}z}{a_{00} + a_{01}z}.$$

Assume $z_L = \mu z_t$ with $\mu \in [K_0^{-1}, K_0]$. For each of the four charged generators, cross multiplication shows

$$K_0^{-1} \le \frac{f_{A_j(L)}(\mu z)}{f_{A_j(t)}(z)} \le K_0 \qquad (z \ge 0).$$

Both cross differences are affine in μ, so it is enough to use $\mu = K_0^{\pm 1}$. After clearing the positive denominator, each is a quadratic in z. Reduction of its coefficients modulo $L^3 - L^2 - 2L + 1$ gives nonnegative quadratic polynomials in L; their exact rational Bernstein coefficients on $[9/5, 181/100]$ are positive. This finite calculation is printed by `strong_replacement_certificate.py`. Boundary slopes 0 and ∞ follow by continuity, so induction on the word proves the column assertion. The charged alphabet is closed under transposition, which proves the row assertion. □

Lemma II.2.40 (Strong replacement)**.** *Let P, Q be charged words, possibly empty, and fix a matrix entry. Replacing one occurrence of A_2 by A_3 cannot decrease its central-to-second-outer ratio:*

$$\frac{(P(L)A_3(L)Q(L))_{ij}}{(P(t)A_3(t)Q(t))_{ij}} \ge \frac{(P(L)A_2(L)Q(L))_{ij}}{(P(t)A_2(t)Q(t))_{ij}}. \tag{688}$$

Likewise replacing A_5 by A_4 cannot decrease the ratio.

Proof. For the first assertion, scale the relevant prefix row and suffix column so their first coordinates are one. Write their outer slopes as $r, z > 0$, and their central slopes as $\mu r, \nu z$. Lemma II.2.39 gives

$$\mu, \nu \in [K_0^{-1}, K_0].$$

The two bilinear evaluations, after the common row and column scales are removed, are

$$F_2(x; r, z) = x + (x^2 - 1)z + r(1 + xz),$$
$$F_3(x; r, z) = (x^2 - 1)(1 + z) + r\{x + (x^2 - 1)z\}.$$

Thus (688) is equivalent to

$$F_3(L; \mu r, \nu z)F_2(t; r, z) - F_2(L; \mu r, \nu z)F_3(t; r, z) \ge 0.$$

The left side is multi-affine in μ, ν. It is therefore enough to check their four corner values $K_0^{\pm 1}$. At each corner, clearing the positive denominator and reducing the coefficients modulo the cubic gives a polynomial in r, z all of whose coefficients are positive on $9/5 \le L \le 181/100$. The exact certificate checks 72 nonzero coefficient polynomials; the least of all their rational Bernstein coefficients is $5339/10000$. It is reproduced by `strong_replacement_certificate.py`. Empty prefixes or suffixes follow from the same formula by taking boundary slopes. Transposing and reversing the entire word turns $A_4 \mapsto A_5$ into $A_3 \mapsto A_2$, proving the second assertion. □

Theorem II.2.41 (All charged open chains at the second outer embedding)**.** *If Q is any word of length $s \ge 1$ in A_2, A_3, A_4, A_5, then*

$$Q(L) \ge_{\text{ent}} R^{s-1}Q(t). \tag{689}$$

Proof. Successively replace every A_3 by A_2 and every A_4 by A_5. Lemma II.2.40 says that no central-to-outer entry ratio increases during this reduction. The resulting word belongs to the A_2, A_5 semigroup, so Theorem II.2.35 gives the lower bound R^{s-1} for every entry. The same bound therefore held before the replacements. □

Theorem II.2.42 (Second-outer comparison for all six generators)**.** *Let W be a cyclic word in $A_1, \dots, A_6$, and let $s \geq 1$ count its A_2, A_3, A_4, A_5 letters. Then*

$$\operatorname{tr} W(L) \geq R^s \operatorname{tr} \prod_{j \in W} |\sigma_2(A_j)| \geq R^s |\sigma_2(\operatorname{tr} W)|. \tag{690}$$

Proof. Write $W = C_1 P_1 \cdots C_s P_s$, where the C_i are charged and $P_i \in \langle A_1, A_6 \rangle$. Expand $P_i = I + (P_i - I)$ and then expand each nonconstant part in monomial matrix units. The all-identity term is Theorem II.2.37. A term with $k \geq 1$ nonidentity gaps has total parabolic degree $D \geq k$; its matrix units cut the charged cycle into k nonempty open chains of total length s. Theorem II.2.41 contributes R^{s-k}, and the parabolic monomials contribute $(L/t)^D = R^D$. Hence the termwise ratio is at least

$$R^{D+s-k} \geq R^s.$$

Summation proves the positive path comparison, and the outer triangle inequality proves the final assertion. □

Corollary II.2.43 (Sharp nonparabolic charge for every positive word)**.** *Let W be a cyclic positive word with positive character, and let s count its A_2, A_3, A_4, A_5 letters. Then*

$$\delta(W) \geq s\delta_7. \tag{691}$$

Proof. Theorems II.2.33 and II.2.42 give

$$\frac{\sigma_0(\operatorname{tr} W)}{\max_{i=1,2} |\sigma_i(\operatorname{tr} W)|} \geq R^s.$$

For $T > 2$, the function $\operatorname{arcosh}(T/2) - \log T$ is increasing. Passing from trace ratios to hyperbolic translation lengths therefore gives

$$a - \max(b, c) \geq s \log R = s\delta_7.$$

The left side is $\delta(W)$. □

Route II.2.44 (Stronger differential open charged-chain inequality)**.** Let $Q(x)$ be a word of length $s \geq 1$ in $A_2(x), A_5(x)$. The endpoint theorem above is sufficient for the heptagonal embedding comparison. A stronger one-parameter assertion, still useful for identifying the full polynomial cone, is

$$xQ'(x) - (s-1)Q(x) \geq_{\text{ent}} 0 \qquad (x^4 \geq 2). \tag{692}$$

This differential assertion is retained as a lemma-to-prove, rather than used as a proved result. Its loss of one unit is sharp: for $Q = A_2$, the lower-left entry in (692) is zero, and the transpose gives the other sharp orientation. It is substantially weaker than the full endpoint-potential macro (694), but the following reduction shows that it is already sufficient for cyclic traces.

Proposition II.2.45 (The open lemma would prove the four-letter core)**.** *Assume (692). If W is any cyclic word in A_1, A_2, A_5, A_6, and $s(W)$ counts its A_2, A_5 letters, then*

$$x \frac{d}{dx} \operatorname{tr} W(x) - s(W) \operatorname{tr} W(x) \geq 0 \qquad (x^4 \geq 2).$$

Proof. Write the cyclic word as

$$C_1P_1C_2P_2\cdots C_sP_s, \qquad C_i\in\{A_2,A_5\}, \quad P_i\in\langle A_1,A_6\rangle.$$

Every entry of P_i is a polynomial with nonnegative coefficients, its constant matrix is the identity, and every monomial in $R_i=P_i-I$ has positive x-degree. Expand the trace after writing each $P_i=I+R_i$.

The term in which every gap contributes I is a closed A_2,A_5-word and is nonnegative by Theorem II.2.29. In any other term, suppose $k\ge 1$ gaps contribute monomials $c_jx^{d_j}E_{a_jb_j}$ from R_i, and put $D=\sum_j d_j\ge k$. These k matrix units cut the charged cyclic word into k nonempty open charged chains $Q_1,\dots,Q_k$. Repeatedly using $E_{ab}QE_{cd}=Q_{bc}E_{ad}$ shows that the trace term factors as

$$cx^D\prod_{j=1}^{k}(Q_j)_{u_jv_j}, \qquad c\in\mathbf{Z}_{>0}.$$

If s_j is the length of Q_j, equation (692) says that its selected entry has logarithmic derivative at least s_j-1. The logarithmic derivative of the whole term, after subtracting $s=\sum_j s_j$, is therefore at least

$$D+\sum_{j=1}^{k}(s_j-1)-s=D-k\ge 0.$$

Every term in the expansion is nonnegative, proving the assertion. □

Experiment II.2.46 (Exact open-chain Bernstein audit)**.** The program `trace_open_charged_tests.py` constructs all four entries of $xQ'-(s-1)Q$ exactly and certifies them on the rational interval $[6/5,2]$ by Bernstein coefficients. Exhaustion through length twelve tested $2^{12}=4096$ words at the last length (and 32760 matrix entries cumulatively). There were no uncertified entries and no recursive subdivisions: every certificate was already positive on the whole interval. The only zero entries occurred in lengths one and two, on the sharp one-orientation chains. Thus the remaining four-letter differential trace problem has been reduced from arbitrary parabolic insertions to the two-generator open inequality (692); no parabolic exponent remains.

The optional full-alphabet mode has now performed the same exact test for all A_2,A_3,A_4,A_5-words through length eight. At the last length it tested $4^8=65536$ words and 262144 entry polynomials. Again there were no zero entries, no recursive subdivisions, and no uncertified polynomials. Theorem II.2.37 proves the paired-entry shadow of this stronger assertion for every length; the individual-entry loss-one inequality is the remaining conversion.

Route II.2.47 (One-dimensional Ising form of the open lemma)**.** The open inequality has an exact ferromagnetic formulation which may be more tractable than entrywise polynomial induction. Put $a=x^2-1$, $\theta=a/(a+1)=1-x^{-2}$, and normalize

$$B_+=x^{-1}A_2=\begin{pmatrix}1 & a/x\\ 1/x & 1\end{pmatrix}, \qquad B_-=x^{-1}A_5=\begin{pmatrix}1 & 1/x\\ a/x & 1\end{pmatrix}.$$

Index rows and columns by spins $u,v\in\{+1,-1\}$. For $\epsilon=+1,-1$, respectively, the entries are exactly

$$B_\epsilon(u,v)=\exp\{C+Juv+\epsilon h(u-v)\}, \quad C=\frac{1}{4}\log\theta, \quad J=-\frac{1}{4}\log\theta>0, \quad h=\frac{1}{4}\log a.$$

Thus an entry of a product is the partition function of an inhomogeneous ferromagnetic Ising chain with fixed endpoint spins. Since

$$\dot{C} = \frac{1}{2a}, \qquad \dot{J} = -\frac{1}{2a}, \qquad \dot{h} = \frac{x^2}{2a}, \quad \dot{} = x\frac{d}{dx},$$

equation (692) is equivalent to the following single correlation inequality for every sign sequence ϵ_i and both fixed endpoints:

$$\mathbb{E}\left[1 + \frac{1}{2a}\sum_i (1 - u_i u_{i+1}) + \frac{x^2}{2a}\sum_i \epsilon_i (u_i - u_{i+1})\right] \geq 0. \tag{693}$$

The expectation uses the displayed Gibbs weights. The last sum telescopes to endpoint fields plus fields supported only where ϵ_i changes. Hence same-sign runs are already the fixed-gauge case, while every genuine difficulty is concentrated at a sign-change interface. This is exactly the interface structure seen in the two-charge kernel and in the four-facet row cone. A proof of (693) by the one-dimensional transfer recursion would prove the entire A_1, A_2, A_5, A_6 trace language via Proposition II.2.45.

Conjecture II.2.48 (Open-path differential macro on the four-letter core)**.** *Let $W(x)$ be a word in A_1, A_2, A_5, A_6, let $s(W)$ count its A_2, A_5 letters, and assume $s(W) > 0$. Give A_2 sign $+1$ and A_5 sign -1; let ϵ_- and ϵ_+ be the signs of the first and last such letters in the linear word. Put*

$$K_\epsilon = \frac{1}{2}\operatorname{diag}(-\epsilon, \epsilon).$$

Then, for $x^4 \geq 2$,

$$xW'(x) - s(W)W(x) + K_{\epsilon_-}W(x) - W(x)K_{\epsilon_+} \geq_{\text{ent}} 0. \tag{694}$$

This is an open-path strengthening of Theorem II.2.29. If a cyclic core word has two cyclically consecutive equal A_2, A_5-signs, cut between them. Then $\epsilon_- = \epsilon_+$, the two boundary potentials in (694) have zero trace, and the desired trace monotonicity follows immediately. Hence only the cyclic skeleton which alternates A_2, A_5 perfectly remains. With no parabolic insertions that skeleton is already Theorem II.2.29. With insertions, its blocks are

$$A_2P_1A_5P_2A_2P_3\cdots, \qquad P_i \in \langle A_1, A_6\rangle;$$

this is the exact Chebyshev–continuant splice still to prove.

Experiment II.2.49 (Exact differential-macro audit)**.** The program `trace_differential_macro_tests.py` makes the projective substitution

$$z = (5x - 6)/(10 - 5x)$$

and clears the common denominator. The three nonzero entry weights become

$$W_1 = 25(1 + z)^2, \quad W_x = 10(3 + 5z)(1 + z), \quad W_\alpha = (1 + 5z)(11 + 15z).$$

For the cleared charged logarithmic derivative, their local defects are

	W_1	W_x	W_α
uncharged	0	$2W_x$	$16(3 + 5z)^2$
charged	$-2W_1$	0	$2(125z^2 + 170z + 61)$.

Thus the only negative local terms are again the $A_2(2,1)$ and $A_5(1,2)$ edges. The program propagates the derivative and both boundary potentials exactly in $\mathbf{Z}[z]$. Exhaustion gives

$\lvert W\rvert$	1	2	3	4	5	6	7
tested	2	12	56	240	992	4032	16256
failures	0	0	0	0	0	0	0.

This does not substitute for the alternating-segment induction, but it shows that the stronger entrywise statement, not merely its trace, is the stable finite-cone invariant.

Route II.2.50 (Comparison with the nearest-wall pencil after the trace test)**.** The trace derivative and the nearest-wall calculation in `curt11.tex` describe the same archimedean charge at different levels. Equation (670) constructs the charge directly in the positive fusion transfer matrix, while the nearest-wall pencil identifies its Hodge determinant line. They are not interchangeable. The diagonal tensor evaluation proposed in `curt11.tex` is nonlinear, and a scalar nearest-wall metric twist cancels against the GIT quotient minimum by Proposition II.14.39. Thus the trace inequality can prove the spectral budget, but it does not by itself pay the moving integral target height.

The viable synthesis is local and evaluation-dependent. Decompose a return into induced blocks, use the trace monotonicity to assign outer charge to each block, and use the compression identity (807) to attach the corresponding finite divisor to the actual remainder state. Recycled primes are then paid by Corollary II.14.53; finite support closes by Theorem II.9.14. What remains is the same time-ordered statement in both languages: time-truncated reduced boundary divisors cannot absorb the positive trace/Hodge charge on infinitely many balanced blocks. This is not a new scalar metric on Maculan's one-return Grassmannian; it is a dissipative covariant along the Rosen path.

Experiment II.2.51 (Exact Rosen-map test)**.** The program `curt10_rosen_tests.py` iterated the exact Rosen map on 36,111 bounded-height field elements. It found 31,257 terminating and 4,854 periodic cases, no unresolved case, and maximal detection time 174. This supports the proposed classification on the tested box, but does not construct the induced global partition or prove the moving-target estimate (655).

II.3 Route I: ordinary moving-target Roth

Vojta's moving-target Roth theorem rules out, for each *fixed* $\varepsilon > 0$,

$$m_S(Q_1 + \cdots + Q_r, P) \geq (2+\varepsilon)h(P)$$

when the sum of the target heights is $o(h(P))$. In the present notation this would use (656) on a subsequence if

$$h(L_g) = o(\chi(g)) \quad \text{and} \quad \inf_g \frac{\delta(g)}{\chi(g)} > 0. \tag{695}$$

The ordinary eigenline bound available from matrix coefficients is only

$$h(L_g) = O(a+b+c+1),$$

and can be much larger than χ. Consequently the standard theorem does not apply until (695) is independently established in an intrinsic normalization.

Route II.3.1 (Verdict)**.** This route works if the Hodge twist proves (695). Target height $o(\chi)$ alone is not enough: along $A_2A_1^N$ the excess $2\delta/\chi$ tends to zero, so no one fixed $\varepsilon > 0$ applies to an infinite tail. The word *normalized* alone alters neither hypothesis of Roth's theorem.

II.4 Route II: Vojta's quantitative theorem with positive target height

The technical theorem in Vojta's proof is stronger than the usual little-o corollary. Let Φ be the standard normal distribution function. For parameters $q \geq 0$, $x > 0$, $D > 0$, and an integer $r \geq 0$, put

$$y = \frac{1}{D}\left(x(2+D) + \frac{2q\, e^{-x^2/2}}{\sqrt{2\pi}\,[1 - r\Phi(-x)]}\right). \tag{696}$$

If

$$r\Phi(-x) < 1, \qquad r\Phi(-x) + \Phi(-y) > 1, \tag{697}$$

then there is no infinite sequence with total target height at most $qh(P) + o(h(P))$ and proximity at least $(2+D)h(P)$.

For a stable numerical optimization set

$$u = 1 - r\Phi(-x), \quad x = -\Phi^{-1}\left(\frac{1-u}{r}\right), \quad B = -\Phi^{-1}(u).$$

Condition (697) is $y < B$. Therefore, for fixed (q, r, u), it suffices that

$$D > D_{\mathrm{req}}(q, r, u) := \frac{2x + 2q\phi(x)/u}{B - x}, \qquad \phi(x) = \frac{e^{-x^2/2}}{\sqrt{2\pi}}, \tag{698}$$

provided $B > x$. This is the formula used in the computation below.

Experiment II.4.1 (Gaussian optimization)**.** For each displayed q, the calculation minimized (698) over $3 \leq r \leq 10{,}000$ and $10^{-15} \leq u < 1$, followed by golden-section refinement. The optimizing integer was $r = 3$ throughout the positive q range shown.

q	least observed D_{req}	required exponent $2 + D$
10^{-8}	0.1935677493	2.1935677493
10^{-6}	0.2430092816	2.2430092816
10^{-4}	0.3497964796	2.3497964796
10^{-3}	0.4734243628	2.4734243628
10^{-2}	0.7844816804	2.7844816804
$3 \cdot 10^{-2}$	1.1420348026	3.1420348026
10^{-1}	2.0162339904	4.0162339904
$3 \cdot 10^{-1}$	4.0487071514	6.0487071514
1	10.5690492381	12.5690492381

For $q = 0$, taking $u \to 0$ makes $D_{\mathrm{req}} \to 0$, recovering the usual $2 + \varepsilon$ theorem. The parameters are fixed before the infinite sequence is considered; allowing $D = D_N \to 0$ does not give one finiteness conclusion. Retaining a fixed positive q, moreover, produces a fixed positive exponent loss.

For the exact family $w_N = A_2 A_1^N$,

$$\chi_N = \log N - 2\log \sigma_0(\lambda) + o(1), \qquad \delta_N \to \delta_7,$$

so the available excess exponent is

$$\kappa_N - 2 = \frac{2\delta_7 + o(1)}{\log N} \longrightarrow 0.$$

It follows that no fixed positive target-height ratio q can close this family via Vojta's quantitative theorem. Even $q = 10^{-4}$ eventually asks for $0.3498\ldots$ of excess exponent while the supplied excess tends to zero.

Route II.4.2 (Verdict)**.** The quantitative theorem is useful on a uniformly supercritical subsequence after proving a sufficiently small target-height ratio. It cannot close the near-critical family even when the target charge is formally zero, because its theorem is a fixed-parameter assertion. It does not replace either the Hodge relative-height estimate or a boundary-scale argument.

II.5 Route III: optimal and conjectural symmetric Roth inequalities

An initial coefficient-one heuristic,

$$m_S(L,Y) \overset{?}{\leq} 2h(Y) + h(L) + o(h(Y) + h(L)),$$

is too optimistic when the target has degree $q > 1$. A degree-q algebraic unit can concentrate essentially $qh(L)$ at one selected embedding. Thus the additive Roth–Liouville inequality one can reasonably seek is

$$m_S(L,Y) \leq (2+\varepsilon)h(Y) + qh_{\mathrm{rel}}(L) + o(h(Y) + h_{\mathrm{rel}}(L)). \tag{699}$$

Maculan's effective theorem has a substantially larger target-height coefficient, of order q/ε, so (699) is already an optimistic limiting form. Substitution of (655) into (699) gives the finite-(g) budget

$$qh_{\mathrm{rel}}(L) + \varepsilon\chi + o(\chi) < 2\delta. \tag{700}$$

For an infinite moving sequence the same fixed $\varepsilon > 0$ must be used throughout. It is therefore not legitimate to delete the $\varepsilon\chi$ term by choosing a new ε for each target.

There is a still stronger conjectural comparison. Schmidt's symmetric Roth conjecture predicts, for algebraic numbers in two prescribed fields, finiteness of

$$|\alpha - \beta| < \max\{M(\alpha), M(\beta)\}^{-2-\varepsilon},$$

where M is Mahler measure. In logarithmic-height notation its best-case analogue here is

$$m_S(L,Y) \leq (2+\varepsilon)\max\{h(Y), qh_{\mathrm{rel}}(L)\} + O(1). \tag{701}$$

This would close on any subsequence on which, for some fixed $\eta > 0$,

$$\max\{h(Y), qh_{\mathrm{rel}}(L)\} \leq (1-\eta)m. \tag{702}$$

The conjecture in its standard form keeps the two number fields fixed, whereas the eigenline splitting field can vary with g. We therefore use (701) only as a best-case numerical test, not as an available implication.

Route II.5.1 (Interim verdict)**.** The two honest tests are (700), with one fixed ε, for an additive moving-target theorem and (702) for the much stronger symmetric conjecture. The next section shows that ordinary eigenline height fails even the second test on the exact near-critical family. Since $h(Y) = \chi = m - \delta + O(1)$, even a zero target height fails (702) when $\delta/m \to 0$.

II.6 The critical-scale obstruction and the Legendre near miss

The distinction between exponent information and a constant-factor improvement is decisive here. Along the extremal family,

$$-\log d(Y_N, L_N) = 2h(Y_N) + 2\delta_7 + O(1),$$

so, after suppressing the bounded metric factor, the scale is

$$d(Y_N, L_N) \asymp C_7 H(Y_N)^{-2}, \qquad C_7 = e^{-2\delta_7} = 0.47889371719708358\ldots. \tag{703}$$

Every fixed-ε Roth statement compares this with $H^{-2-\varepsilon}$, and therefore says nothing. This remains true for the conjectural symmetric Roth inequality and for the usual formulation of Vojta's conjecture: their exceptional sets depend on a fixed $\varepsilon > 0$. The obstruction persists even if the moving target is assigned height zero.

This is not an accident of proof technology. Ordinary continued-fraction convergents naturally attain C/q^2 with $C = O(1)$. Hence a generic boundary-Roth assertion forbidding (703) would be false. What can be true is a *recognition theorem*: below a sharp Legendre constant the approximant must carry continued-fraction symbolic structure.

There is a useful classical calibration. In the usual normalization, Legendre's sufficient inequality is

$$\left|x - \frac{p}{q}\right| < \frac{1}{2q^2},$$

not merely a constant smaller than one. On the other hand, Hurwitz's theorem gives infinitely many reduced fractions with

$$\left|x - \frac{p}{q}\right| < \frac{1}{\sqrt{5}q^2}$$

for every irrational real x. Consequently even a constant strictly below one at exponent two is not, by itself, a finiteness phenomenon. Crossing a Legendre threshold only identifies the approximant as a convergent; it does not make an algebraic target rational or periodic. Any closing argument must use the resulting symbolic identification as new information.

For Rosen fractions of odd index k, Nakada [23] proves that the Legendre constant is

$$L_k = \frac{R_k}{R_k + 1}, \qquad R_k^2 + (2 - \lambda_k)R_k - 1 = 0, \qquad R_k > 0.$$

At $k = 7$,

$$R_7 = 0.90586048649843892\ldots, \qquad L_7 = 0.47530262205222591\ldots,$$
$$C_7 - L_7 = 0.00359109514485767\ldots, \qquad \frac{C_7}{L_7} = 1.00755538677517\ldots.$$

Thus the bare outer-gap constant misses the sharp Rosen Legendre threshold by only 0.7555%. Equivalently, Legendre recognition would follow from an effective surplus

$$\delta > -\tfrac{1}{2}\log L_7 = 0.371902\ldots,$$

whereas the limiting outer surplus is $0.368138\ldots$. Finite members of the family can lie on the favorable side, but its tail cannot be handled by the bare spectral inequality.

Allowing mediant Rosen convergents changes this conclusion substantially. Kraaikamp–Nakada–Schmidt [24] prove that for odd index the Legendre constant for the union of principal and mediant Rosen convergents is

$$L_7^{\mathrm{med}} = \lambda - R_7 = 0.89607724930639933\ldots.$$

Consequently

$$C_7 = 0.4788937\ldots < L_7^{\mathrm{med}}.$$

Thus the uniform spectral constant is strong enough for *mediant recognition*, though not for principal-convergent recognition. For the finite extremal words the principal threshold holds exactly through $A_2A_1^7$; it first fails at $A_2A_1^8$:

N	6	7	8	9
$e^{-2\delta_N}$	0.4737910	0.4748767	0.4756484	0.4762168.

The mediant threshold retains a large uniform margin on the full tail.

There are two further qualifications. Nakada's theorem recognizes approximants $p/q = g\infty$ already lying in the Hecke cusp orbit; it does not apply to an arbitrary K-rational or $\mathbf{Q}$-rational projection of an Asai tensor. Also, the unspecified $O(1)$ in the Hodge distance changes C_7, so a coordinate- and metric-exact computation is mandatory. The promising new target is consequently not a stronger Roth theorem but a *Rosen recognition lemma*: construct a cusp approximant from the return, put its exact error below L_7^{med}/q^2, classify it as a principal or mediant convergent, and use the resulting prefix overlap to force a finite or periodic itinerary.

Route II.6.1 (Verdict)**.** The constant-scale observation redirects the residual proof from Diophantine finiteness to symbolic recognition. Mediant recognition has ample numerical room; what is missing is the exact construction of the relevant G_7-rational approximant from the Hodge return. The final closing step must then prove that the recognized principal/mediant convergent creates a repeated Rosen natural-extension state, not merely observe that good approximants exist.

II.7 Exact group-theoretic reformulation

The recognition problem has a useful formulation which removes all metric ambiguity. Put $\Gamma = \mathrm{PSL}_2(\mathcal{O}_K)$, $H = G_7$, and let $P < \Gamma$ be the stabilizer of ∞. Since K has class number one, Γ acts transitively on $\mathbf{P}^1(K)$, and

$$\mathbf{P}^1(K) = \Gamma/P, \qquad H\backslash\mathbf{P}^1(K) = H\backslash\Gamma/P.$$

Proposition II.7.1 (Stabilizer form of the classification)**.** *The desired cusp/fixed-point classification is equivalent to*

$$\mathrm{Stab}_{G_7}(x) \neq \{1\} \quad \textit{for every } x \in \mathbf{P}^1(K). \tag{704}$$

Equivalently, for every $g \in \Gamma$,

$$H \cap gPg^{-1} \neq \{1\}.$$

Proof. If $x = g\infty$, its stabilizer in H is $H \cap gPg^{-1}$. A nonidentity element of a Fuchsian triangle group fixing a real boundary point is parabolic or hyperbolic. In the parabolic case its fixed point lies in the unique cusp orbit $H\infty$. In the hyperbolic case, an element fixing $x \in K$ has an eigenvector $(x, 1)^t$ and eigenvalue $cx + d \in K$; hence it is special. Conversely, both a cusp and a special hyperbolic fixed point have a nontrivial stabilizer. □

Thus the residual statement is a thin-group double-coset theorem: every vertex of the G_7-action on $\mathbf{P}^1(K)$ lies on a nontrivial stabilizer cycle. Hodge theory produces useful weights on this graph, algebraic number theory supplies the integral lattice and units in P, and Rosen principal/mediant recognition supplies its local finite-state edges. This is the precise sense in which the problem is a hybrid rather than simply "number theory or Hodge theory."

II.7.1 What can be obtained from the whole thin group

The six Rosen letters are a coding of G_7; they are not the structural reason for the proposed classification. Two global facts about the embedding of the group are more relevant.

Proposition II.7.2 (Full Galois Zariski closure)**.** *The image of*

$$\rho_{\mathrm{Gal}} : G_7 \longrightarrow \mathrm{SL}_2(\mathbf{R})^3, \qquad g \longmapsto (\sigma_0(g), \sigma_1(g), \sigma_2(g)),$$

is Zariski dense.

Proof. Each projection contains the nontrivial upper unipotent $\sigma_i(T)$ and its lower-unipotent conjugate by S, and is therefore Zariski dense in SL_2. The identity component of the full closure is a semisimple subdirect product. If it were proper, the algebraic Goursat lemma would identify two simple factors. An automorphism of PSL_2 preserves squared trace, whereas the three squared traces of the one element ST are

$$\lambda_0^2, \qquad \lambda_1^2, \qquad \lambda_2^2,$$

and are pairwise distinct. No factors are identified. □

This has a negative and a positive consequence. Negatively, no polynomial identity on the three matrix factors can encode the classification: an identity holding on G_7 would hold on the full ambient product. Generator calculations can establish inequalities on a chosen return semigroup, but cannot by themselves force the rational Borel intersection (704). Positively, the problem is naturally an orbital arithmetic problem for a Zariski-dense thin group, so its invariant form is visible on the rational flag variety rather than in a particular alphabet.

There is a stronger topological statement in the two outer factors.

Proposition II.7.3 (Joint outer density, including congruence fibres)**.** *The image of*

$$G_7 \longrightarrow \mathrm{PSL}_2(\mathbf{R}) \times \mathrm{PSL}_2(\mathbf{R}), \qquad g \longmapsto (\sigma_1(g), \sigma_2(g)),$$

is topologically dense. Every finite-index subgroup has the same property. Hence, whenever reduction modulo an ideal $\mathfrak{a}$ is surjective, the lifts of each prescribed finite residue class are dense in the two outer real factors.

Proof. Let $\alpha = \lambda^2 - 1$. The element

$$A_3 = \begin{pmatrix} \alpha & \alpha \\ \lambda & \alpha \end{pmatrix}$$

has trace 2α. It is hyperbolic in the distinguished factor, but

$$|2\sigma_1(\alpha)| < 2, \qquad |2\sigma_2(\alpha)| < 2,$$

so both outer images are elliptic. Each has infinite order. Indeed, a finite-order outer image would give $\sigma_i(A_3^n) = \pm I$; injectivity of σ_i would then give $A_3^n = \pm I$, contradicting central hyperbolicity.

Thus each outer projection is nondiscrete. It is nonelementary because it contains the upper parabolic T and its lower-parabolic conjugate by S. The classification of closed subgroups of $\mathrm{PSL}_2(\mathbf{R})$ makes each projection dense.

Let H be the closure of the joint image. The powers of A_3 form an infinite subset of a compact product of two elliptic circles, so H is nondiscrete. Its projections are full. Normality of H° shows that each projection of the identity component is trivial or full. If one is trivial, the other full factor and the surjectivity of H give the full product. If both are full, topological Goursat gives either the full product or the graph of an automorphism. Such a graph is self-normalizing, while an automorphism preserves squared traces and $\lambda_1^2 \neq \lambda_2^2$ are the squared traces of ST in the two factors. The graph case is impossible.

For a finite-index subgroup, the connected outer product is a finite union of translates of the closure of its image. That closure is finite-index and hence open, so it is the whole connected group. Apply this to a congruence kernel and multiply one chosen lift of a residue class by that kernel. □

This is a genuinely global use of the thin group. It says that neither the finite congruence conditions nor the two outer Hodge directions cause a qualitative failure of approximation. The unproved part is quantitative: one needs such a simultaneous lift with distinguished hyperbolic displacement

$$R = (1 + o(1)) \log \mathrm{Nm}\, \mathfrak{a}$$

and outer Cartan sum at most $(2/5 + o(1))R$. Topological density permits an arbitrarily large central cost and therefore does not exactify the incidence. In this form the target is an optimal quantitative weak- approximation theorem for a thin diagonal embedding, rather than a statement about any preferred generator.

The remaining quantitative demand has an invariant numerical scale. For $x = [P] \in \mathbf{P}^1(K)$, $e = (1, 0)^t$, a good prime $\mathfrak{p}$, and fixed θ, C, set

$$\mathcal{R}_{\theta,C}(x; \mathfrak{p}) = \inf\{R(g) : g \in G_7,\ \det(P, ge) \in \mathfrak{p},\ a_1(g) + a_2(g) \leq \theta R(g) + C\}. \tag{705}$$

Joint outer density in every congruence fibre proves that this infimum is finite. Indeed, after fixing the residue class, take distinct lifts whose outer images approach the identity. Their central images escape because the Fuchsian representation is discrete. Thus the outer Cartan sum is bounded while the right side of the Hodge window grows.

Let

$$\kappa_\theta(x) = \inf_{C \geq 0} \liminf_{\mathrm{N}\,\mathfrak{p} \to \infty} \frac{\mathcal{R}_{\theta,C}(x; \mathfrak{p})}{\log \mathrm{N}\,\mathfrak{p}}, \tag{706}$$

over good primes. This is an invariant of the complete thin orbit; no coding or preferred generators occur in its definition.

Proposition II.7.4 (Exact covering wall)**.** *If x is not a cusp, then*

$$\kappa_\theta(x) \geq (1/2 + \theta)^{-1}.$$

If x is a cusp, then $\kappa_\theta(x) = 0$. Thus at the Hodge value $\theta = 2/5$ the two possibilities are separated by the exact lower bound $10/9$.

Proof. For a noncusp, $d_g = \det(P, ge)$ is nonzero for every g. If it lies in $\mathfrak{p}$, the finite-place inequality and the three real-place operator bounds give

$$\log \mathrm{N}\,\mathfrak{p} \leq \log |\,\mathrm{N}\, d_g| \leq a_0(g) + a_1(g) + a_2(g) + O_P(1) \leq (1/2 + \theta)R(g) + C + O_P(1).$$

Taking infima and lower limits proves the assertion. If $x = h\infty$, then $d_h = 0$; choose C once so that h lies in the Hodge window. The same h is admissible at every prime, giving exponent zero. □

The entropy of the Fuchsian lattice is one and the projective quotient has $\mathrm{N}\,\mathfrak{p}+1$ points, so the information-theoretic covering exponent is one. The theorem would follow from the genuinely thin-group inverse assertion

$$\kappa_{2/5}(x)>1 \quad\Longrightarrow\quad G_7\cap B_x(K)\neq\{1\}. \tag{707}$$

Indeed, a rational flag with trivial stabilizer would have exponent at most one by the inverse assertion and at least $10/9$ by the product formula. A special fixed point is correctly retained: its exponent is at least $10/9$, and its nontrivial Borel intersection is exactly the exception allowed on the right of (707). This formulation isolates the central difficulty better than the phrase "optimal lifting." One needs a pointwise inverse theorem for the coherent rational section, not a universal diameter bound and not an estimate for the six generators.

Global Arakelov sieve: sum every prime before closing

The preceding exponent suggests a route which does not select one *time-truncated* prime. Let

$$\Omega_R(\delta)=\{g: R\le R(g)<R+1,\ a_1(g)+a_2(g)\le(2/5+\delta)R\},\qquad M_R=\#\Omega_R(\delta),$$

and for a primitive P put

$$D_g(P)=\det(P,gP),\qquad N_{\mathfrak{p}}(P;R)=\#\{g\in\Omega_R(\delta): g[P]=[P]\pmod{\mathfrak{p}}\}.$$

If there is no projectively nonidentity stabilizer, every $D_g(P)$ in this large-radius shell is nonzero and

$$\log|\,\mathrm{N}\,D_g(P)|\le(9/10+\delta)R+O_P(1). \tag{708}$$

Proposition II.7.5 (A signed level above $9/10$ closes globally)**.** *Let $9/10+\delta<u<1$, assume $M_R=e^{R-o(R)}$, and suppose*

$$\left|\sum_{\substack{\mathfrak{p}\ \text{good}\\ \mathrm{N}\,\mathfrak{p}\le e^{uR}}}(\log\mathrm{N}\,\mathfrak{p})\left(N_{\mathfrak{p}}(P;R)-\frac{M_R}{\mathrm{N}\,\mathfrak{p}+1}\right)\right|=o(RM_R). \tag{709}$$

Then $G_7\cap B_P(K)$ contains a projectively nonidentity element.

Proof. Assume the conclusion false. Since $\mathfrak{p}\mid D_g(P)$ is precisely the displayed projective congruence, summing over all primes first gives

$$\begin{aligned}\sum_{\substack{\mathfrak{p}\ \text{good}\\ \mathrm{N}\,\mathfrak{p}\le e^{uR}}}(\log\mathrm{N}\,\mathfrak{p})N_{\mathfrak{p}}(P;R)&=\sum_{g\in\Omega_R(\delta)}\sum_{\substack{\mathfrak{p}\mid D_g(P)\\ \mathfrak{p}\ \text{good},\ \mathrm{N}\,\mathfrak{p}\le e^{uR}}}\log\mathrm{N}\,\mathfrak{p}\\&\le\sum_{g\in\Omega_R(\delta)}\log|\,\mathrm{N}\,D_g(P)|\\&\le(9/10+\delta)RM_R+O_P(M_R).\end{aligned}$$

The prime-ideal Mertens formula gives

$$\sum_{\substack{\mathfrak{p}\ \text{good}\\ \mathrm{N}\,\mathfrak{p}\le e^{uR}}}\frac{\log\mathrm{N}\,\mathfrak{p}}{\mathrm{N}\,\mathfrak{p}+1}=(u+o(1))R.$$

Equation (709) therefore makes the first member of the preceding display equal to $(u+o(1))RM_R$, a contradiction. □

The absolute values in the earlier formulation were unnecessary: the proof uses only the signed sum in (709). This is the global-height version of the thin-group strategy. It does not watch prime factors being born and it does not ask a chosen short word to hit a chosen prime. It compares, in one identity, the complete finite divisor of every incidence with the distribution of the complete Hodge-balanced shell in all projective congruence quotients.

The signed assumption follows from a fixed-section variance estimate. Set

$$E_{\mathfrak{p}}(P;R) = N_{\mathfrak{p}}(P;R) - \frac{M_R}{\mathrm{N}\,\mathfrak{p}+1}.$$

At a good prime, let

$$e_{P,\mathfrak{p}} = \delta_{[P \bmod \mathfrak{p}]} - \frac{\mathbf{1}}{\mathrm{N}\,\mathfrak{p}+1}, \qquad \mu_R = M_R^{-1} \sum_{g\in\Omega_R(\delta)} \delta_g.$$

The projective permutation representation decomposes as $\mathbf{1} \oplus \mathrm{St}_{\mathfrak{p}}$, and therefore

$$E_{\mathfrak{p}}(P;R) = M_R \left\langle \mathrm{St}_{\mathfrak{p}}(\mu_R) e_{P,\mathfrak{p}}, e_{P,\mathfrak{p}} \right\rangle. \tag{710}$$

Thus the variance below is exactly a large sieve for one coherent Steinberg newvector across the complete prime family.

It also has a common-divisor expansion which makes the group law visible. Put

$$I_{\mathfrak{p}}(g) = \mathbf{1}_{\{\mathfrak{p}\mid D_g(P)\}}.$$

Then

$$I_{\mathfrak{p}}(g) I_{\mathfrak{p}}(h) = I_{\mathfrak{p}}(h) I_{\mathfrak{p}}(h^{-1}g). \tag{711}$$

Indeed, two elements fixing $[P]$ modulo $\mathfrak{p}$ have quotient fixing $[P]$, and the converse follows after multiplication by h. Consequently every common prime divisor of $D_g(P)$ and $D_h(P)$ divides $D_{h^{-1}g}(P)$, and

$$\sum_{\mathrm{N}\,\mathfrak{p}\le Q} (\log \mathrm{N}\,\mathfrak{p}) I_{\mathfrak{p}}(g) I_{\mathfrak{p}}(h) = \sum_{\substack{\mathfrak{p}\mid (D_g(P), D_h(P)) \\ \mathrm{N}\,\mathfrak{p}\le Q}} \log \mathrm{N}\,\mathfrak{p}. \tag{712}$$

Expanding the variance is therefore a centered average common-divisor energy for all pairs in the shell. This is an intrinsic thin-group convolution problem: it involves $h^{-1}g$, the complete shell, and all projective congruence representations, with no preferred generators.

This identifies the cleanest standard conjectural comparison. A level-one Barban–Davenport–Halberstam/Elliott–Halberstam principle for the Hodge-weighted affine-sieve sequence $D_g(P)$ gives the next conjecture. Quotient-wise Ramanujan does not: it bounds each Steinberg coefficient separately and provides no orthogonality in $\mathfrak{p}$. Bourgain–Kontorovich's beyond-expansion method shows that bilinear dispersion can cross the expansion barrier for particular thin orbits, but the fixed coherent Borel section, the two Hodge cocycles, and the required level $u > 9/10$ are additional here. Hence the condition below is standard-shaped rather than a disguised form of Roth or scalar GRH.

Conjecture II.7.6 (Single-prime fixed-section BDH)**.** *For every fixed $u < 1$, if $[P]$ has trivial projective stabilizer, then*

$$\sum_{\substack{\mathfrak{p}\ \text{good} \\ \mathrm{N}\,\mathfrak{p}\le e^{uR}}} (\log \mathrm{N}\,\mathfrak{p}) |E_{\mathfrak{p}}(P;R)|^2 \le M_R e^{o(R)}. \tag{713}$$

Cauchy–Schwarz and the prime-ideal theorem bound the signed error in (709) by

$$e^{uR/2+o(R)}M_R^{1/2}e^{o(R)} = e^{(1+u)R/2+o(R)} = o(RM_R),$$

because $u < 1$ and $M_R = e^{R-o(R)}$. Thus (713) conditionally completes the classification.

This variance formulation also measures the obstruction without a pointwise conjecture. Choose $9/10 + \delta < v < u$. If the rational flag has no stabilizer, the norm bound (708) forces

$$N_{\mathfrak{p}}(P;R) = 0 \qquad (e^{vR} \le \mathrm{N}\,\mathfrak{p} \le 2e^{vR}).$$

On that interval $E_{\mathfrak{p}} = -M_R/(\mathrm{N}\,\mathfrak{p}+1)$, and hence

$$\sum_{\mathrm{N}\,\mathfrak{p}\le e^{uR}} (\log \mathrm{N}\,\mathfrak{p})|E_{\mathfrak{p}}|^2 \ge M_R^2 e^{-vR+o(R)} = e^{(2-v)R-o(R)}. \tag{714}$$

The expected BDH scale is $e^{R+o(R)}$, so the exact excess is $e^{(1-v)R}$. This is a coherent fixed-section hole across a whole dyadic family of primes; it is not a defect of one generator or one bad modulus.

An absolute Elliott–Halberstam estimate is stronger. For example, the pointed relative square-root estimate

$$\left|N_{\mathfrak{p}}(P;R) - \frac{M_R}{\mathrm{N}\,\mathfrak{p}+1}\right| \ll \left(\frac{M_R}{\mathrm{N}\,\mathfrak{p}}\right)^{1/2} (\mathrm{N}\,\mathfrak{p})^{\epsilon}$$

does imply (709) for every $u < 1$, because its weighted sum is

$$O\Big((M_R e^{uR})^{1/2+o(1)}\Big) = o(M_R).$$

Thus the global averaging does not remove the analytic gap, but it weakens its required form from pointwise relative Lindelöf to fixed-section BDH. The latter asks only for square-root cancellation after summing over the entire prime family. This is also a direct synthesis of the Arakelov and thin-group viewpoints.

Proposition II.7.7 (High-level half-shells are injective but sparse)**.** *Let $\mathcal{A}_R, \mathcal{B}_R$ be compatible block families with*

$$|\mathcal{A}_R|, |\mathcal{B}_R| = e^{R/2-o(R)}, \qquad R(a), R(b) \le R/2 + O(1),$$

and impose the blockwise Hodge bounds

$$a_1(a) + a_2(a),\ a_1(b) + a_2(b) \le (1/5+\epsilon)R.$$

If $[P]$ has trivial projective stabilizer, then for every prime with

$$\mathrm{N}\,\mathfrak{p} > e^{(9/10+3\epsilon)R}$$

the two maps

$$a \mapsto [a^{-1}P \bmod \mathfrak{p}], \qquad b \mapsto [bP \bmod \mathfrak{p}]$$

are injective.

Proof. If two left endpoints coincide modulo $\mathfrak{p}$, then $\mathfrak{p}$ divides the nonzero algebraic integer

$$\Delta_{a,a'} = \det(a^{-1}P, a'^{-1}P).$$

The central embeddings of the two blocks contribute at most $R/2$ to the logarithmic norm of this determinant. The two outer embeddings together contribute at most

$$a_1(a) + a_2(a) + a_1(a') + a_2(a') \le (2/5 + 2\epsilon)R.$$

Thus

$$\log |\operatorname{N} \Delta_{a,a'}| \le (9/10 + 2\epsilon)R + O_P(1),$$

which is incompatible with divisibility by the displayed prime for large R. The right endpoints are identical, using $\det(bP, b'P)$. □

This proves that the remaining failure is not caused by multiplicities inside either endpoint measure. Let $\alpha_{\mathfrak{p}}$ and $\beta_{\mathfrak{p}}$ be their normalized projective measures and let $\mathbf{u}$ be uniform measure. Injectivity gives exactly

$$\|\alpha_{\mathfrak{p}} - \mathbf{u}\|_2^2 = \frac{1}{|\mathcal{A}_R|} - \frac{1}{\operatorname{N}\mathfrak{p} + 1}, \qquad \|\beta_{\mathfrak{p}} - \mathbf{u}\|_2^2 = \frac{1}{|\mathcal{B}_R|} - \frac{1}{\operatorname{N}\mathfrak{p} + 1}.$$

Their collision probability is

$$\langle \alpha_{\mathfrak{p}}, \beta_{\mathfrak{p}} \rangle = \frac{1}{\operatorname{N}\mathfrak{p} + 1} + \langle \alpha_{\mathfrak{p}} - \mathbf{u}, \beta_{\mathfrak{p}} - \mathbf{u} \rangle.$$

For $q = \operatorname{N}\mathfrak{p} = e^{uR+o(R)}$, Cauchy–Schwarz controls the last term only by $e^{-R/2+o(R)}$, while the desired main term is $e^{-uR+o(R)}$. At $u = 9/10$, the endpoint sets have size $q^{5/9+o(1)}$: larger than $q^{1/2}$, but still small enough for two arbitrary subsets to be disjoint.

The relative square-root prediction for the unnormalized collision count is

$$\frac{|\mathcal{A}_R||\mathcal{B}_R|}{q} + O\left(\sqrt{\frac{|\mathcal{A}_R||\mathcal{B}_R|}{q}}\, q^{o(1)}\right).$$

After normalization, its error is $(q|\mathcal{A}_R||\mathcal{B}_R|)^{-1/2+o(1)}$, which is the elementary L^2 product multiplied by the missing $q^{-1/2}$. This gives a precise appreciation of the impasse: all same-side collisions have already been eliminated by the global product formula, yet a lower bound for the cross-intersection of the two sparse thin-block images is still required. Neither generator inequalities nor marginal expansion addresses that cross-correlation.

Route II.7.8 (The outer-length sum does not characterize special holonomy)**.** The basic special element A_1A_6 has a K-rational eigenline, and the unit product formula gives the attractive equality

$$\ell_1(A_1A_6) + \ell_2(A_1A_6) = \ell_0(A_1A_6).$$

It is natural to ask whether $\ell_1(g) + \ell_2(g) \le \ell_0(g)$ holds for every hyperbolic element, with equality detecting special holonomy. The next member of the same full-group family disproves it. Writing

$$U(t) = \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}, \qquad L(t) = \begin{pmatrix} 1 & 0 \\ t & 1 \end{pmatrix},$$

one has

$$g = A_1^2A_6 = U(2\lambda)L(\lambda), \qquad \operatorname{tr}(\sigma_i g) = 2 + 2\lambda_i^2.$$

Consequently

$$\ell_i(g) = 2\operatorname{arcosh}(1 + \lambda_i^2),$$

and certified evaluation in the isolating intervals for the three roots of $X^3 - X^2 - 2X + 1$ gives

$$\ell_0(g) = 4.2503946429\ldots, \quad \ell_1(g) = 1.2388670024\ldots, \quad \ell_2(g) = 3.1809537139\ldots.$$

Thus $\ell_1(g) + \ell_2(g) > \ell_0(g)$. Longer words $A_1^n A_6$ approach the parabolic regime in which the two outer lengths together can be nearly twice the central length. The exact equality at $n = 1$ is therefore a statement about the evaluated unit eigenline, not an operator-length extremum. A whole-group proof still has to subtract the cusp coordinate v_g or retain the logarithmic conormal angle; translation lengths alone discard precisely that datum.

A tempting Steinberg centrality shortcut, and why it fails

The joint-density theorem suggests averaging a whole Hodge-balanced shell rather than following a preferred word. At a good split prime, let $Q = |\mathcal{O}_K/\mathfrak{p}|$, let $X = \mathbb{P}^1(\mathbb{F}_Q)$, and for $y \in X$ put

$$e_y = \delta_y - \frac{1}{Q+1}\mathbf{1}.$$

For a finite nonnegative measure ν on G_7, write $M = \nu(G_7)$ and let $\bar{\nu}$ be its reduction. The exact permutation-representation decomposition gives

$$\sum_g \nu(g)\mathbf{1}_{\{\bar{g}\zeta = x\}} = \frac{M}{Q+1} + \langle \mathrm{St}_{\mathfrak{p}}(\bar{\nu})e_\zeta, e_x\rangle\,. \tag{715}$$

If $\bar{\nu}$ were a central element of the finite group algebra, Schur's lemma would give $\mathrm{St}_{\mathfrak{p}}(\bar{\nu}) = c_\nu I$. Since

$$\langle e_\zeta, e_x\rangle = -\frac{1}{Q+1} \qquad (x \neq \zeta),$$

every off-diagonal target would then have the same count $(M - c_\nu)/(Q+1)$. This would supply for free the factor $Q^{-1/2}$ missing from the operator-norm estimate.

There is, however, an exact obstruction to producing such a measure by a finite geometric shell. Apart from the projective identity, every conjugacy class in the non-elementary free product $G_7 \simeq C_2 * C_7$ is infinite. Indeed, the centralizer of a nonidentity element is finite cyclic for torsion and virtually cyclic for an infinite-order element, and therefore has infinite index. Hence a nonnegative conjugation-invariant ℓ^1-measure on G_7 is supported at the identity. Translation length is conjugacy invariant, but the set of all conjugates having a fixed translation length is infinite; imposing base-point displacement to make the shell finite breaks conjugacy invariance.

The direct way to centralize after reduction is to average over finite conjugators. Every term hgh^{-1} then contains the displacement of h twice. Even granting conjugators at the information-theoretic projective scale $(1 + o(1))\log Q$, this construction has a $(2 + o(1))\log Q$ budget before the middle element is charged, whereas exact closing requires less than $(10/9 - o(1))\log Q$ in the Hodge-typical window. This does not rule out a different source of finite-quotient centrality; it proves that literal conjugacy averaging is another form of the factor-two construction.

The surviving spectral target is now precise. For a probability shell ν_R of mass $e^{R-o(R)}$, with $Q = e^{R/(1+\eta)+o(R)}$, one needs the individual coefficient bound

$$|\langle \mathrm{St}_{\mathfrak{p}}(\bar{\nu}_R)e_\zeta, e_x\rangle| \ll e^{-R/2}Q^{-1/2+o(1)}. \tag{716}$$

The operator-norm, or Ramanujan, scale supplies only $e^{-R/2}$. The extra $Q^{-1/2}$ is delocalization of one propagated Steinberg vector at the coherent rational coordinate x. Equation (716) is therefore

the exact finite-representation meaning of the earlier "Lindelöf-sized local limit." It is a statement about an individual matrix coefficient, not a stronger norm bound for the generators.

Define the Galois flag

$$\Delta_K(x) = (\sigma_0 x, \sigma_1 x, \sigma_2 x) \in \mathbf{P}^1(\mathbf{R})^3.$$

Proposition II.7.1 says exactly that every such rational flag has a projectively nontrivial G_7-stabilizer. If $H_x \in \mathrm{SL}_2(\mathcal{O}_K)$ has first column representing x, the lower coordinate of $H_x^{-1} g^{-1} e$ is

$$r_g = \det(x, g^{-1}e).$$

Thus the remainder is the lower Bruhat coordinate of the thin cusp orbit relative to the rational Borel B_x. This definition involves the whole orbit $G_7 e$, not the choice of Rosen generators.

For the Hodge geometry one must change from the raw Galois bases to positively oriented symplectic bases. Put

$$\varepsilon_i = \mathrm{sign}(\lambda_i), \qquad R_i = \begin{pmatrix} \varepsilon_i & 0 \\ 0 & 1 \end{pmatrix}, \qquad \rho_i(g) = R_i \sigma_i(g) R_i^{-1}, \qquad x_i^\dagger = \varepsilon_i \sigma_i(x).$$

Then $\rho_i(T)$ is translation by $|\lambda_i| > 0$. This correction is forced at $i = 2$: a holomorphic self-map of the upper half-plane cannot intertwine cusp translations of opposite orientation. The signed flag $(x_0^\dagger, x_1^\dagger, x_2^\dagger)$ is obtained from $\Delta_K(x)$ by a fixed projective conjugation, so it has exactly the same stabilizer. Traces, lengths, and Zariski closure are unchanged.

The second global fact is Hodge-theoretic domination.

Proposition II.7.9 (Global two-factor length domination)**.** *Let $g \in G_7$ be hyperbolic in the distinguished representation, and let $\ell_i(g)$ be the translation length of $\sigma_i(g)$ on $\mathbb{H}$, with length zero when the outer element is elliptic. Then*

$$\ell_i(g) < \ell_0(g), \qquad i = 1, 2,$$

and consequently

$$2\ell_0(g) - \ell_1(g) - \ell_2(g) > 0.$$

Proof. Let z lie on the axis of $\sigma_0(g)$. Equivariance of the modular map F_i and strict Schwarz–Pick contraction give

$$\ell_i(g) \le d_{\mathbb{H}}\big(F_i(z), \rho_i(g) F_i(z)\big) = d_{\mathbb{H}}\big(F_i(z), F_i(gz)\big) < d_{\mathbb{H}}(z, gz) = \ell_0(g).$$

Adding the two inequalities proves the second assertion. □

Fix a standard cusp neighborhood in the triangle quotient and write $L_{\mathrm{thick}}(g)$ for the length of the part of the closed geodesic of g outside it. The same argument proves the quantitative inequality

$$2\ell_0(g) - \ell_1(g) - \ell_2(g) \ge \kappa_7 L_{\mathrm{thick}}(g) \tag{717}$$

for a constant $\kappa_7 > 0$. Indeed, on the compact thick part the two hyperbolic derivative norms have upper bounds $k_i < 1$, while on the cusp Schwarz–Pick gives one; integrating along the axis gives

$$\ell_i(g) \le \ell_0(g) - (1 - k_i) L_{\mathrm{thick}}(g).$$

Thus the coarse relative spectral gap is a direct theorem about the modular embedding of the whole group.

This proposition proves, without a word calculation, the eigenvalue gap for

$$\rho_\star = \operatorname{Sym}^2 \rho_0 \oplus (\rho_1 \otimes \rho_2).$$

Parabolic elements are the neutral boundary because all three translation lengths vanish. The detailed fusion estimates compare L_{thick} with particular symbolic complexities and sharpen the constant; they do not create the gap.

Route II.7.10 (The full cusp-transvection Voronoi domain)**.** Let $e = (1,0)^t$, $f = (0,1)^t$. For every $g \in G_7$, write

$$c = ge, \qquad d = gf, \qquad U_{c,n} = gT^n g^{-1}.$$

If $P = uc + rd$, then

$$U_{c,n} P = (u + n\lambda r)c + rd. \tag{718}$$

Thus the full group, rather than a selected alphabet, supplies one integer family of shears for every cusp $c \in G_7 e$.

Define

$$\mathscr{R}_{G_7} = \{x \in \mathbf{P}^1(K) : h(\gamma x) \geq h(x) \text{ for all } \gamma \in G_7\}.$$

Every orbit meets this set: by Northcott finiteness, the orbit has a point of least Weil height. The classification is therefore equivalent to

$$x \in \mathscr{R}_{G_7} \quad \Longrightarrow \quad \operatorname{Stab}_{G_7}(x) \neq \{1\}. \tag{719}$$

At a reduced point the inequalities

$$h([(u + n\lambda r)c + rd]) \geq h([uc + rd]) \qquad (c \in G_7 e,\ n \in \mathbf{Z}) \tag{720}$$

hold simultaneously. If one r vanishes, the point is a cusp. The remaining target is to show that a noncuspidal simultaneous minimum has only finitely many reduced complete Borel frames; two equal frames give the required stabilizer.

This is a genuine thin-group reduction problem. The usual Rosen move checks a local face associated with one chosen cusp frame. Equation (720) imposes every conjugate face at once. Its possible noncompactness is also sharply identified: ordinary height quotients out the rank-two unit-frame lattice, whereas the Hodge line $\mathcal{N}$ has nonzero degree along that lattice. A successful proof should establish Hodge properness of this global Voronoi domain, with a zero-degree ray forcing a repeated unit frame. This is the same obstruction as the stopped-incidence problem, but it no longer privileges the six generators.

Lemma II.7.11 (One-ideal zero test)**.** *Let $\mathfrak{a} \subset \mathcal{O}_K$ be a nonzero integral ideal and let $r \in \mathfrak{a}$. If*

$$|\operatorname{N}_{K/\mathbf{Q}}(r)| < \operatorname{N}(\mathfrak{a}),$$

then $r = 0$.

Proof. If $r \neq 0$, then $(r) \subseteq \mathfrak{a}$, so reversal of ideal inclusion under norm gives

$$|\operatorname{N}_{K/\mathbf{Q}}(r)| = \operatorname{N}((r)) \geq \operatorname{N}(\mathfrak{a}).$$

□

Route II.7.12 (Adelic congruence mixing plus Hodge shrinking)**.** This gives a second genuinely thin-group route. For $x = [P] \in \mathbf{P}^1(K)$, write

$$r_g = \det(P, g^{-1}e).$$

The congruence $r_g \equiv 0 \pmod{\mathfrak{a}}$ says that the cusp $g^{-1}\infty$ lies in the reduction of the rational Borel B_x modulo the single ideal $\mathfrak{a}$. Strong approximation for the Zariski-dense thin group supplies such congruence incidences away from a fixed finite set of bad ideals. Expansion of the congruence quotients is the corresponding quantitative group input. Neither assertion by itself controls the archimedean size of r_g.

The precise sufficient strengthening is the following adelic shrinking-target statement. Unless $B_x \cap G_7$ is already projectively nontrivial, there are $\epsilon > 0$, C_x, arbitrarily large ideals $\mathfrak{a}$, and $g \in G_7$ such that

$$r_g \in \mathfrak{a}, \tag{721}$$

$$\log|\,\mathrm{N}(r_g)| \le (1-\epsilon)\log \mathrm{N}(\mathfrak{a}) + C_x. \tag{722}$$

For $\mathrm{N}(\mathfrak{a})$ large, (722) is strictly smaller than $\log \mathrm{N}(\mathfrak{a})$; Lemma II.7.11 then gives $r_g = 0$, so x is a cusp. Hence this statement implies the full cusp/special classification.

This formulation is not prime-by-prime divisor bookkeeping. The ideal $\mathfrak{a}$ packages all chosen finite places at once, and (722) is a single Arakelov inequality comparing its finite degree with the three real metrics. It also says exactly what ordinary super-approximation lacks: a congruence solution must be found inside the archimedean Hodge-shrinking region, not merely inside a short word ball. A route to (721)–(722) would combine uniform mixing on congruence covers of the triangle surface with the thick-part contraction (717). The exceptional failure of mixing in a rational Borel is not noise; it is precisely the desired special stabilizer alternative.

Route II.7.13 (Thin-group route: rational-flag conicality)**.** The desired conclusion separates into two statements.

1. Prove that at first norm-hitting prefixes the evaluated Borel incidence vector determined by $\Delta_K(x)$ lies in the nonpositive bundle of the relatively dominated representation $\rho_\star$. Equivalently, prove the stopped incidence lower bound while retaining the full flag, not merely an operator norm.
2. Apply lattice discreteness. A strictly negative integral vector forces a parabolic endpoint; a bounded nonpositive integral vector repeats at two reduced returns and the intervening group element lies in B_x. Proposition II.7.1 then gives the cusp or special alternative.

The first statement is the rational-flag conicality theorem still to be proved. The second is the integral nonpositive-weight argument proved later in Proposition II.19.3. This division uses the actual thin group: relative domination supplies its limit geometry, the rational Borel supplies the arithmetic flag, and integrality converts conical recurrence into an exact stabilizer. The Rosen alphabet is needed only to select first returns and certify the final repetition.

Route II.7.14 (The incidence is already a group cocycle)**.** There is an exact identity which sharpens the transport problem. Write

$$g = \begin{pmatrix} a & b \\ c & d \end{pmatrix}, \qquad x = p/q, \qquad j_g(x) = cx + d.$$

For $e = (1,0)^t$, $\ell_P = (-q, p)$, and $r_g = \ell_P(g^{-1}e)$, one has

$$\frac{r_g}{r_0} = \frac{cp + dq}{q} = j_g(x). \tag{723}$$

Indeed, $g^{-1}e = (d, -c)^t$, so both numerator and denominator acquire the same minus sign. Hence the remainder incidence satisfies the standard group law

$$j_{gh}(x) = j_g(hx)j_h(x). \tag{724}$$

The finite scalar needed by the argument is therefore not an unknown normal function: it is the denominator automorphy factor for the action of the entire thin group on its rational flag bundle.

Its Hodge normalization is also canonical. Put

$$\Theta_g^{\mathrm{rat}}(x) = \log|\sigma_1 j_g(x)| + \log|\sigma_2 j_g(x)| - 2\log|\sigma_0 j_g(x)|.$$

Write

$$\rho_i(g) = \begin{pmatrix} a_{i,g} & b_{i,g} \\ c_{i,g} & d_{i,g} \end{pmatrix}, \qquad j_{i,g}^{\dagger}(w) = c_{i,g}w + d_{i,g}.$$

Then

$$j_{i,g}^{\dagger}(x_i^{\dagger}) = \sigma_i(j_g(x)).$$

With $F_0(z) = z$, put

$$\Theta_g^{\mathrm{Hod}}(z) = \log|j_{1,g}^{\dagger}(F_1 z)| + \log|j_{2,g}^{\dagger}(F_2 z)| - 2\log|j_{0,g}^{\dagger}(z)|.$$

Both are additive cocycles; consequently

$$\mathcal{A}_g(z, x) = \Theta_g^{\mathrm{rat}}(x) - \Theta_g^{\mathrm{Hod}}(z) \tag{725}$$

is an explicit cocycle on the fiber product of the rational flag bundle and the modular embedding. It is the automorphy factor of the ratio between the tautological line of character $(-2, 1, 1)$ and $\mathcal{N}$.

If g fixes x, then $j_g(x) = \mu \in \mathcal{O}_K^{\times}$, and the product formula gives

$$\Theta_g^{\mathrm{rat}}(x) = -3\log|\sigma_0\mu|.$$

Thus its zero period is parabolic and its nonzero period is exactly the unit direction of the special alternative. This is substantial progress over the formal groupoid tensor: the group-wide algebraic multiplier and its equality cases are now explicit.

The corresponding Green potential is equally explicit. Put

$$\Phi(w, \xi) = \log\frac{(\operatorname{Im} w)^2}{|w - \xi|^2},$$

using the other affine chart at $\xi = \infty$, and define

$$G_x(z) = \Phi(F_1 z, x_1^{\dagger}) + \Phi(F_2 z, x_2^{\dagger}) - 2\Phi(z, x_0^{\dagger}).$$

The Möbius identities for imaginary part and difference give

$$G_{gx}(gz) - G_x(z) = 2\mathcal{A}_g(z, x). \tag{726}$$

If $\omega_{\mathrm{hyp}} = -dd^c \log \operatorname{Im} w$, then

$$dd^c G_x = 2\big(2\omega_{\mathrm{hyp}} - F_1^*\omega_{\mathrm{hyp}} - F_2^*\omega_{\mathrm{hyp}}\big) > 0. \tag{727}$$

Indeed, $\log|w-\xi|^2$ is harmonic on $\mathbb{H}$, and strictness is Schwarz–Pick for the two modular maps. Also

$$\Phi(w,\xi)=\log\operatorname{Im} w-\log\frac{|w-\xi|^2}{\operatorname{Im} w},$$

so this potential is exactly the Hodge term minus the angular Busemann term computed earlier.

The open integral model is elementary once the correct cocycle is visible. The tautological line on $\mathbf{P}^1_K$, trivialized by the second homogeneous coordinate q, has transition

$$\frac{q(gP)}{q(P)}=\frac{cp+dq}{q}=j_g(x)=\frac{r_g}{r_0}. \tag{728}$$

Since $g\in \mathrm{SL}_2(\mathcal{O}_{K_v})$ at every finite place, it preserves the primitive tautological lattice. Hence the finite transition divisor is exactly $(r_g)-(r_0)$. Equations (724) and (726) are the two gluing laws for one adelically metrized line on the open action groupoid.

The signed normalization also settles its standard-cusp asymptotic. Writing $y=\operatorname{Im} z$, the modular-embedding cusp expansion is

$$F_i(z)=\frac{|\lambda_i|}{\lambda_0}z+B_{i,0}+O(e^{-2\pi y/\lambda_0}).$$

Therefore, for every fixed real ξ in either projective chart,

$$\Phi(F_i(z),\xi)=O(y^{-1}),\qquad \Phi(z,\xi)=O(y^{-1}),\qquad G_x(z)\longrightarrow 0.$$

Thus the real metric, its strict interior curvature, the finite lattice, and the normalized standard-cusp extension are all explicit.

The Higgs section gives a useful holomorphic factorization. In period coordinates, the section $\Xi\in H^0(\mathcal{N}^{-2})$ is represented by $F_1'F_2'$, after normalizing its cusp leading coefficient. Hence

$$\Psi_x(z)=\frac{F_1'(z)(z-x_0^\dagger)}{F_1(z)-x_1^\dagger}\frac{F_2'(z)(z-x_0^\dagger)}{F_2(z)-x_2^\dagger} \tag{729}$$

satisfies the following exact whole-group law. If $\Psi_{i,x}=F_i'(z)(z-x_0^\dagger)/(F_i(z)-x_i^\dagger)$, then

$$\Psi_{i,gx}(gz)=\Psi_{i,x}(z)\frac{j_{0,g}^\dagger(z)}{j_{i,g}^\dagger(F_i(z))}\frac{j_{i,g}^\dagger(x_i^\dagger)}{j_{0,g}^\dagger(x_0^\dagger)}. \tag{730}$$

Indeed, differentiated equivariance supplies the square of the first denominator ratio, while the Möbius difference formula cancels one copy and supplies the endpoint ratio. Multiplying the two identities gives

$$\frac{\Psi_{gx}(gz)}{\Psi_x(z)}=\frac{j_{0,g}^\dagger(z)^2}{j_{1,g}^\dagger(F_1(z))j_{2,g}^\dagger(F_2(z))}\frac{\sigma_1(j_g(x))\sigma_2(j_g(x))}{j_g(x)^2}. \tag{731}$$

The logarithmic modulus is exactly $\Theta_g^{\mathrm{rat}}-\Theta_g^{\mathrm{Hod}}$. In particular this is one multiplicative cocycle for every $g\in G_7$, not a generator charge. When $gx=x$, $j_g(x)=\mu$ is a unit and the endpoint factor is $\sigma_1(\mu)\sigma_2(\mu)/\mu^2$, with logarithmic modulus $-3\log|\sigma_0\mu|$. Thus its nonzero periods are exactly the unit holonomy of the special alternative.

The same section satisfies

$$\log|\Psi_x|^2=G_x+\log\|\Xi\|^2.$$

Its zeros are the full lift of $3P_7$, independent of x. On the vertical geodesic $z = x_0^\dagger + iy$, Schwarz–Pick gives

$$|F_i'(z)| \le \frac{\operatorname{Im} F_i(z)}{\operatorname{Im} z}, \qquad |F_i(z) - x_i^\dagger| \ge \operatorname{Im} F_i(z),$$

and therefore $|\Psi_x(z)| \le 1$, strictly in the interior. Its cusp limit is one. Distributionally,

$$dd^c\big(G_x + \log \|\Xi\|^2\big) = \delta_{3P_7}.$$

This identifies the group-wide source of every compact charge: it is the Green divisor of the complete G_7-orbit of the order-seven point. The charged Rosen letters only record crossings of a particular tessellation near that orbit.

The estimate $|\Psi_x| < 1$ is only pointwise. Multiplying it at successive returns would be an error: the relative cocycle is a difference of values of $\log|\Psi_x|$, so the intermediate terms telescope. Repeated thick returns can therefore be offset by an unbounded change of the unit frame. The global height jump must control precisely this resetting term.

What remains is the global stopped boundary step: prove that the cusp charts glue in the relevant toroidal or Deligne extension with an effective height jump at first norm-record regions, and turn that height jump into the stopped all-place degree inequality. Curvature positivity and local boundedness alone do not control a moving rational boundary flag.

Route II.7.15 (Arithmetic Livšic formulation of the missing transport)**.** There is a group-wide way to formulate the transport section τ_g. On the geodesic flow of the triangle quotient, let

$$\varphi(v) = 2 - \|DF_1(v)\|_{\mathrm{hyp}} - \|DF_2(v)\|_{\mathrm{hyp}}.$$

It is positive in the interior and its integral through the thick part is bounded below by (717). For a fixed rational flag x, changing the cusp frame changes the Busemann function by the exact incidence coboundary

$$b_x(gv) - b_x(v) = \log\left|\frac{\ell_x(g^{-1}e)}{\ell_x(e)}\right| = \log\left|\frac{r_g}{r_0}\right|.$$

The stopped degree cocycle is the sum of the positive Hodge term and this incidence coboundary, with all three archimedean places and the finite ideal summed by the product formula.

For a compact hyperbolic return map, a positive Livšic or calibrated subaction theorem converts inequalities on all closed-orbit periods into a transfer function which makes the cocycle nonnegative on every arrow. In the present cusped setting one first induces past parabolic runs. The periodic points of the resulting rational-flag skew product are exactly the nontrivial intersections with B_x: unipotent periods are cusps and unit periods are special cycles.

The preceding denominator identity fixes the algebraic part of τ_g. One must realize the explicit cocycle $\mathcal{A}$ as an *arithmetic* Livšic transfer, not merely a real Hölder function:

1. its integral model must have the divisor of the already specified algebraic multiplier r_g/r_0;
2. its real metric must solve the calibrated Hodge cocycle equation;
3. its coefficient height must be paid by the thick-part surplus, not by the total height of the return word; and
4. vanishing of the transfer system must retain the complete rational flag, so its zero-period locus is exactly the Borel stabilizer locus.

The identity $\mathcal{N}^{-2} \simeq \mathcal{O}(3P_7)$ supplies the geometric line and the divisor on which such a transfer can live. Arithmetic holonomy is relevant because it is a method for replacing an analytic continuation

operator by integral jets with controlled coefficient height. In these terms, the Hodge-normalized Roth-plus statement is an arithmetic Livšic theorem at the critical vanishing slope. This synthesis uses closed-orbit data of the whole thin group; finite-state fusion is only one way to verify the real period inequality.

Proposition II.7.16 (Why ordinary Livšic theory is circular here)**.** *Let $\Gamma = G_7$, let $x \in \mathbf{P}^1(K)$, and put $\Gamma_x = \mathrm{Stab}_\Gamma(x)$. The denominator multiplier $j_g(x)$ is a cocycle on the transformation groupoid $\Gamma \ltimes \mathbf{P}^1(K)$, but it is not a scalar cocycle on the compact triangle quotient with x held fixed. The diagonal quotient carrying that fixed flag is*

$$\Gamma\backslash\big(T^1\mathbb{H} \times \Gamma x\big) \simeq \Gamma_x\backslash T^1\mathbb{H}. \tag{732}$$

Consequently its nonconstant closed geodesics are represented by hyperbolic elements of Γ_x. If Γ_x is projectively trivial, there are no such closed geodesics, and a periodic-orbit criterion for a Livšic transfer is vacuous.

Proof. The group law is

$$j_{gh}(x) = j_g(hx)j_h(x).$$

The first factor on the right is evaluated at hx, not at x. Thus the multiplier descends with a fixed endpoint only after restricting to elements which fix that endpoint. The orbit Γx is Γ/Γ_x; the standard induction identity gives (732). Closed geodesics on this cover are conjugacy classes of hyperbolic elements of Γ_x, which proves the last assertion. □

This proposition changes the status of the preceding route. The Green potential and the groupoid cocycle are useful exact formulas, but an ordinary positive Livšic theorem cannot create the desired stabilizer: the compact periodic system needed to invoke it already exists only after such a stabilizer has been supplied.

The noncircular replacement is the finite pointed congruence tower. For an ideal $\mathfrak{a}$, put

$$\Gamma_x(\mathfrak{a}) = \{g \in \Gamma : g[x] = [x] \pmod{\mathfrak{a}}\}.$$

This subgroup has finite index, its cover has genuine closed geodesics, and strong approximation identifies its index with a projective congruence-orbit size. Moreover

$$\Gamma_x = \bigcap_{\mathfrak{a}} \Gamma_x(\mathfrak{a}), \tag{733}$$

because $g[x] = [x]$ modulo every ideal is equivalent to $\det(P, gP) = 0$. The real problem is therefore uniform: find short, Hodge-balanced closed loops on these finite covers quickly enough that the product formula forces one loop into the intersection (733). This is precisely the pointed optimal-lifting problem below. It uses the congruence tower of the whole thin group and does not assume the desired periodic orbit at the outset.

Route II.7.17 (The transport belongs to a filtered flat block)**.** The symbol τ_g must not be read as a transport in $\mathcal{N}$ alone. Such a transport cannot have the required periods: deck equivariance makes the Hodge metric periodic around a closed orbit, while the Schwarz–Pick defect has a strictly positive hyperbolic period. The smallest flat carrier containing the relevant Hodge quotient is

$$\mathbb{E}_\star = \mathrm{Sym}^2 \mathbb{V}_0 \oplus (\mathbb{V}_1 \otimes \mathbb{V}_2).$$

The highest Hodge lines in the two summands are

$$A = \mathcal{L}_0^{\otimes 2}, \qquad B = \mathcal{L}_1 \otimes \mathcal{L}_2, \qquad \mathrm{Hom}(A, B) = \mathcal{N}.$$

An integral lift of the explicit cocycle $\mathcal{A}$ would take one flat vector, project it to both A and B, and form the ratio of those two projections. The flat vector, rather than the quotient line, retains the lattice and the complete rational flag; the quotient records exactly the negative Hodge weight. To speak of integrality one must keep the full Galois orbit of $\mathbb{E}_\star$, not one selected real block.

This corrects the earlier formal groupoid-line construction. It also gives a precise test for arithmetic holonomy: construct simultaneous jets for the two projections of a common integral flat vector, with common denominators, and prove that their height loss is bounded by the thick-part spectral surplus. Separate approximations to the two Hodge lines would not suffice, because their quotient could lose both the lattice and the endpoint flag.

Experiment II.7.18 (Exact stabilizer certificates)**.** The program `rosen_certificate_tests.py` tracks the exact field state and the exact accumulated G_7 matrix. On termination it verifies that the accumulated matrix sends the input to 0. On repetition, say $P_b x = P_a x$, it constructs $h = P_a^{-1} P_b$, checks $\det((x,1)^t, h(x,1)^t) = 0$ exactly, and verifies

$$\mu + \mu^{-1} = \mathrm{tr}(h), \qquad \mu \in K,$$

so the certificate is a special element, not merely a floating-point period. In 2,000 random tests with coefficient bound 1000 and denominator bound 300, there were 1,726 cusp certificates and 274 special certificates, no unresolved state, and maximum detection time 236. The longest sample was

$$x = (184 + 842\lambda - (\lambda^2 - 1))/213;$$

its eventual period had length 2 and exact eigenvalue $\mu = \lambda - 2$.

A second run of 5000 inputs with coefficient bound 5000, denominator bound 1000, and seed 77 gave 4350 cusp and 650 special certificates, again with no unresolved state. The period histogram was 641 of length two and nine of length six. Only two exact eigenvalues occurred: $\lambda - 2$ and $-417 + 507\lambda - 221\alpha$. The maximum detection time was 266, at

$$x = (-1212 - 1367\lambda + 2270\alpha)/281,$$

which terminated at a cusp.

II.8 Discarded route: the already-known 2-adic image

This section is retained only as an audit of a route which does *not* address the global theorem. The image of G_7 at 2, including the two residual orbits and their lifting through powers of 2, was already understood in the earlier work cited below. Recomputing it supplies a useful consistency check on the exact programs, but no new archimedean rigidity. In particular, the 2-adic closure forgets the distinction between a cusp and a real special cycle; the exact point 671 below is a counterexample to interpreting its two local orbits as the two global types. Consequently none of the results in this section is used in the main route.

The prime 2 is inert in K, because

$$X^3 - X^2 - 2X + 1 \equiv X^3 + X^2 + 1 \pmod 2$$

has no root in $\mathbf{F}_2$. Thus $\mathcal{O}_K/2\mathcal{O}_K \simeq \mathbf{F}_8$, and the image of λ generates $\mathbf{F}_8^\times$. Moreover

$$\alpha = \lambda^2 - 1 \equiv \lambda^2 + 1 = \lambda^3 \pmod 2.$$

Proposition II.8.1 (Two residual orbits)**.** *The reduction of G_7 modulo 2 is the dihedral group D_{14}. Its action on $\mathbf{P}^1(\mathbf{F}_8)$ has two orbits:*

$$\mathcal{O}_{\mathrm{cusp}} = \{\infty, 0, 1, \lambda, \lambda^2, \lambda^5, \lambda^6\}, \qquad \mathcal{O}_{\mathrm{exc}} = \{\lambda^3, \lambda^4\}.$$

The two exceptional points are the roots of the reduction of the fixed-point equation of M_0.

Proof. Modulo 2, both generators S and T are involutions, while ST has order seven. Their action gives the stated $7 + 2$ orbit decomposition by direct application to ∞. The fixed-point equation of $M_0 = \left(\begin{smallmatrix} -1 & \lambda \\ \lambda & -1-\lambda^2 \end{smallmatrix}\right)$ reduces, after division by λ, to

$$X^2 + \lambda X + 1 = 0.$$

Its roots in $\mathbf{F}_8$ are λ^3, λ^4, since their product is $\lambda^7 = 1$ and their sum is λ. □

This congruence calculation does not prove that every lift in either residual orbit has the corresponding global behavior. Calegari computes the exact adelic closure of the triangle group, including its 2-adic image, in [56]; see in particular Theorems 1.1 and 1.3. The earlier finite-quotient orbit questions also appear in Hanson–Merberg–Towse–Yudovina [62]. The computations below merely reproduce the $n = 14$ projective outcome certificates and explicit orbit counts through level 2^6.

Experiment II.8.2 (Rank-two unit lattice)**.** The program `rosen_unit_lattice_tests.py` tested all

$$x = \lambda^r \alpha^s, \qquad -40 \le r, s \le 40,$$

using exact Rosen states and exact cusp/fixer certificates. Of the 6561 units, 4685 were cusps and 1876 were special; no state was unresolved. In every case, with no exception,

$$x \text{ was special} \quad \Longleftrightarrow \quad r + 3s \equiv 3 \text{ or } 4 \pmod 7,$$

which is exactly the condition that its reduction lie in $\mathcal{O}_{\mathrm{exc}}$. There were 1864 period-two certificates and twelve period-six certificates. The only two exact eigenvalues observed were

$$\lambda - 2 = -(\lambda^2 + \lambda)^{-1} \quad \text{and} \quad -417 + 507\lambda - 221\alpha.$$

The longest tested unit orbit took 716 steps, at $(r, s) = (-32, -40)$, and terminated at the cusp.

Experiment II.8.3 (Outcome versus reduction modulo 2)**.** The program `mod2_outcome_tests.py` compared exact Rosen certificates with the residual orbit for 2000 random points of coefficient bound 3000 and denominator bound 600. The complete contingency table was

	cusp certificate	special certificate
$\mathcal{O}_{\mathrm{cusp}}$	1735	0
$\mathcal{O}_{\mathrm{exc}}$	0	265.

There were no unresolved states. Thus reduction modulo 2 perfectly predicted the two *types* in this sample, although it does not distinguish the different special G_7-orbits.

Remark II.8.4 (The residual orbits do not label the global type). The preceding contingency table is a sampling phenomenon, not a local–global rule. The exact point $671 \equiv 1 \pmod 2$ reduces to $\mathcal{O}_{\mathrm{cusp}}$, but its Rosen orbit has an exact period of length 23. The program `rosen_certificate_tests.py` returns after 61 steps a special fixer with eigenvalue

$$34291233 - 42755048\lambda + 19025852\alpha \in K.$$

Thus 671 is special and is not a cusp. Archimedean special cycles can therefore lie above the seven-point as well as the two-point residual orbit.

Theorem II.8.5 (Exactly two orbits modulo every power of 2). *For every $n \geq 1$, the action of G_7 on $\mathbf{P}^1(\mathcal{O}_K/2^n)$ has exactly two orbits, lifting $\mathcal{O}_{\mathrm{cusp}}$ and $\mathcal{O}_{\mathrm{exc}}$, of respective sizes*

$$7 \cdot 8^{n-1} \qquad \text{and} \qquad 2 \cdot 8^{n-1}.$$

Proof. Put $\overline{G} = G_7 \bmod 2 = D_{14}$, and $X = \lambda E_{12} \in \mathfrak{sl}_2(\mathbf{F}_8)$. Exact row reduction over $\mathbf{F}_2$ gives

$$\dim_{\mathbf{F}_2} \langle \mathrm{Ad}(g)X : g \in \overline{G} \rangle = 7.$$

For a residual projective point v, the infinitesimal action is

$$\mathrm{Ad}(g)X \longmapsto \det\big(v, \mathrm{Ad}(g)Xv\big) \in \mathbf{F}_8.$$

At $v = \infty$ and $v = \alpha$, exact row reduction gives rank three over $\mathbf{F}_2$, hence this map is onto $\mathbf{F}_8$. By $\overline{G}$-equivariance it is onto at every point of the corresponding two residual orbits.

Now

$$T^{2^n} = I + 2^n \lambda E_{12}.$$

Conjugating by lifts of the fourteen elements of $\overline{G}$, and multiplying the resulting matrices, realizes modulo 2^{n+1} every element of the seven-dimensional span above. These matrices fix every projective point modulo 2^n, while surjectivity of the infinitesimal action makes their stabilizer transitive on the eight lifts of that point modulo 2^{n+1}. Starting from the $7+2$ decomposition at $n = 1$, induction shows that each orbit acquires exactly eight lifts and that no new orbit appears.

All finite-field ranks in this proof are independently reproduced by `mod2_lie_certificate.py`. □

Corollary II.8.6 (The complete 2-adic orbit decomposition). *Let $\overline{G}_7$ be the closure of G_7 in $\mathrm{PSL}_2(\mathcal{O}_{K,2})$. Then*

$$\mathbf{P}^1(\mathcal{O}_{K,2}) = \overline{G}_7 \infty \ \sqcup \ \overline{G}_7 \alpha.$$

Proof. The theorem gives transitivity on each compatible finite-level orbit. For two points with the same residual type, the sets of elements of $\overline{G}_7$ carrying one to the other modulo 2^n are nonempty, closed, and nested. Compactness gives an element in their intersection. The two orbits remain disjoint because their reductions modulo 2 are disjoint. □

Corollary II.8.7 (The remaining global rigidity statement). *Every $x \in \mathbf{P}^1(K) \subset \mathbf{P}^1(K_2)$ has a nontrivial stabilizer in $\overline{G}_7$. The desired global classification is equivalent to*

$$\mathrm{Stab}_{\overline{G}_7}(x) \cap G_7 \neq \{1\} \qquad \text{for every } x \in \mathbf{P}^1(K). \tag{734}$$

Proof. If $x = \gamma\infty$ with $\gamma \in \overline{G}_7$, conjugate a nontrivial element of the pro-2 closure of the parabolic stabilizer of ∞. If $x = \gamma\alpha$, conjugate M_0. This proves the first assertion. Condition (734) says exactly that x has a nontrivial discrete G_7-stabilizer; Proposition II.7.1 then identifies it as a cusp or a special hyperbolic fixed point. □

Thus the congruence/local classification is complete but orthogonal to the question at issue. Forcing a compact 2-adic stabilizer to meet the discrete triangle group is simply another formulation of the missing global rigidity; the local calculation does not make that implication easier.

Experiment II.8.8 (Finite-ring verification)**.** The exact finite-ring program `mod2_power_orbits.py` represents $\mathbf{P}^1(\mathcal{O}_K/2^n)$ in its two local charts and performs breadth-first search under S, T. It independently returned

n	$\lvert\mathbf{P}^1(\mathcal{O}_K/2^n)\rvert$	$\lvert\mathcal{O}^{(n)}_{\mathrm{cusp}}\rvert$	$\lvert\mathcal{O}^{(n)}_{\mathrm{exc}}\rvert$
1	9	7	2
2	72	56	16
3	576	448	128
4	4608	3584	1024
5	36864	28672	8192
6	294912	229376	65536.

in agreement with the theorem.

Strong approximation by itself cannot replace an archimedean height inequality, since a dense congruence closure does not imply equality of thin-group orbits. The local state therefore contributes no usable descent datum.

II.8.1 A precise 2-adic height criterion for the global theorem

For $x \in \mathbf{P}^1(K)$, let $Z_x = \infty$ if its reduction lies in $\mathcal{O}_{\mathrm{cusp}}$, and let $Z_x = \alpha$ if its reduction lies in $\mathcal{O}_{\mathrm{exc}}$. The theorem above gives, for every n, some $y_n \in G_7 Z_x$ congruent to x modulo 2^n. Mere existence of these congruence lifts is not enough; their archimedean height is decisive.

Proposition II.8.9 (Short congruence lifts prove the global result)**.** *Suppose that for every $x \in \mathbf{P}^1(K)$ there are arbitrarily large n and points $y_n \in G_7 Z_x$, congruent to x modulo 2^n, such that*

$$h(y_n) \le (1-\eta) n \log 2 + O_x(1) \tag{735}$$

for some $\eta = \eta(x) > 0$. Then

$$\mathbf{P}^1(K) = G_7\infty \sqcup \bigcup_{g \text{ special hyperbolic}} G_7 \operatorname{Fix}(g).$$

In fact, the stronger two-orbit statement $\mathbf{P}^1(K) = G_7\infty \sqcup G_7\alpha$ follows.

Proof. The local chordal proximity at the unique prime above 2 satisfies

$$m_2(x, y_n) \ge n \log 2 - O_x(1).$$

The global height inequality for two distinct algebraic points in the same number field gives

$$m_2(x, y_n) \le h(x) + h(y_n) + O(1).$$

Together with (735), these inequalities are incompatible for large n unless $x = y_n$. If $Z_x = \infty$, this makes x a cusp. If $Z_x = \alpha$, then $x \in G_7\alpha$; since α is fixed by the special hyperbolic element M_0, every point in its orbit is a special hyperbolic fixed point. □

This criterion shows exactly what additional statement would be needed, but also why it is not a viable route. Already 671 lies over $\mathcal{O}_{\mathrm{cusp}}$ but is special rather than cuspidal, so (735) must fail for $x = 671$ and $Z_x = \infty$. The proposed short-lift statement would have to encode the missing archimedean special cycle in advance. It therefore gives no reduction of the problem and is dropped.

II.9 Route V: a cubic Panti block criterion

Panti's decreasing-block proof in quadratic trace fields [28] suggests a direct finite-state route. Its height calculation extends cleanly to three real embeddings, but the conclusion splits into the three Hodge regimes rather than giving uniform descent.

Let $A = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in \mathrm{SL}_2(\mathcal{O}_K)$, and at the embedding σ_i put

$$f_{A,i}(t) = \frac{\max\{|a_i t + b_i|, |c_i t + d_i|\}}{\max\{|t|, 1\}}.$$

For an interval $I \subset \mathbf{P}^1(\mathbf{R})$, set

$$m_0(A, I) = \inf_{t \in I} f_{A,0}(t), \qquad m_i(A) = \inf_{t \in \mathbf{P}^1(\mathbf{R})} f_{A,i}(t) \qquad (i = 1, 2).$$

The unrestricted minima have the elementary exact formula

$$m_i(A) = \frac{1}{\|\sigma_i(A)^{-1}\|_{\infty \to \infty}} = \frac{1}{\max\{|d_i| + |b_i|, |c_i| + |a_i|\}}. \tag{736}$$

Proposition II.9.1 (Cubic decreasing-block test)**.** *If*

$$m_0(A, I) m_1(A) m_2(A) > 1, \tag{737}$$

then $h(Ax) > h(x)$ for every $x \in K$ with $\sigma_0(x) \in I$. If the product is less than 1, then such x for which $h(Ax) < h(x)$ are dense in I. Hence the inverse branch A^{-1} is uniformly height-decreasing on its cylinder exactly on the favorable side of (737); equality is the neutral boundary.

Proof. Choose a homogeneous pair (p, q) for x. Since A is unimodular, the fractional ideal (p, q) and therefore the entire finite-place part of the projective height are unchanged. The archimedean height difference is

$$3(h(Ax) - h(x)) = \sum_{i=0}^{2} \log f_{A,i}(\sigma_i(x)).$$

The lower bound in (737) proves the first assertion. For the converse, choose three real arguments at which the product is less than one and use weak approximation in the totally real field K to approximate them simultaneously by the three conjugates of one $x \in K$. Continuity gives the strict reverse height inequality. Formula (736) is the conorm identity $\inf_{\|v\|_\infty = 1} \|Av\|_\infty = \|A^{-1}\|^{-1}$. □

For $A = A_2 A_1^N$ and the unaccelerated positive interval $I = [0, \infty]$, direct evaluation gives

N	1	3	8	10	100
$m_0 m_1 m_2$	.261986	.088800	.015540	.009948	.00009992.

Thus raw Weil-height descent fails strongly on the parabolic ray, exactly as predicted by the positive central Hodge weight. Restricting to a compact accelerated cylinder raises m_0, but the asymptotic sign is still governed by $-a + b + c = \chi$. This calculation explains both why Panti's quadratic proof works (there is only one outer conjugate and spectral domination wins) and why its untwisted cubic analogue does not.

Experiment II.9.2 (Algebraic norm descent fails pointwise)**.** The program `rosen_norm_descent_tests.py` tested the exact, projectively unit-invariant candidate

$$E(p,q) = |\,\mathrm{N}(p)| + |\,\mathrm{N}(q)|$$

on 500 primitive pairs with coefficient bound 5000. Every orbit closed, but E was not monotone: the largest one-step factor was $29171.21\ldots$, and the largest record-to-initial factor was $2336.83\ldots$. On the logarithmic scale the largest observed record exponent was much smaller,

$$\frac{\log(1+\max_n E(p_n,q_n))}{\log(1+E(p_0,q_0))} = 1.32716989\ldots.$$

This suggests a polynomial norm barrier, but E alone cannot prove Northcott finiteness because it does not see relative unit directions. Those directions are exactly what the mod–2 unit-lattice experiment isolates.

Experiment II.9.3 (High-precision height-overshoot test)**.** The program `rosen_height_barrier_tests.py` follows exact field vectors and evaluates the three archimedean terms at high precision. For

$$x = (184 + 842\lambda - (\lambda^2 - 1))/213,$$

the orbit reaches a special 2-cycle after 236 steps. A first run which froze the three conjugates at only 100 decimal digits falsely reported a height rise of 97.95. Recomputing the conjugates themselves at 1000 digits gives instead

$$\max_n (h(T^n x) - h(x)) = 0.1465150319857895\ldots, \qquad \min_n (h(T^n x) - h(x)) = -6.5732556571533900\ldots.$$

In 50 random inputs with coefficient bound 1000 and denominator bound 300, the greatest rise was $1.64393510164478\ldots$. Five further batches, comprising 375 inputs with coefficient bound 5000 and denominator bound 1000, had greatest batch rises

$$1.2044, \quad 2.6007, \quad 2.1310, \quad 2.79643, \quad 2.26973.$$

All 425 orbits terminated or repeated, and the largest verified rise was $2.79643\ldots$. The corrected computation therefore supports, rather than refutes, a bounded-overshoot theorem. It also demonstrates that fixed low-precision embeddings are unsafe once exact field coefficients have hundreds of digits.

Proposition II.9.4 (No universal additive height barrier)**.** *For every $C > 0$, there is an $x \in K$ for which*

$$h(Tx) > h(x) + C.$$

Proof. Let $B_n = \left(\begin{smallmatrix} n\lambda & -1 \\ 1 & 0 \end{smallmatrix}\right)$, the positive-input Rosen branch with digit $-n$. Put $\lambda_i = \sigma_i(\lambda)$. At the distinguished embedding choose t_0 arbitrarily close to $1/(n\lambda_0)$, in the interior of the corresponding digit cylinder. At the other two embeddings choose $|t_1|, |t_2|$ arbitrarily large. For the max norm, the three local expansion factors tend to

$$\frac{1}{n\lambda_0}, \qquad n|\lambda_1|, \qquad n|\lambda_2|,$$

respectively. Their product is

$$\frac{n|\lambda_1\lambda_2|}{\lambda_0} = \frac{n}{\lambda_0^2},$$

because $|\lambda_0\lambda_1\lambda_2| = 1$. Weak approximation in the totally real field K supplies an $x \in K$ whose three conjugates are as close to these choices as desired. Since $B_n \in \mathrm{SL}_2(\mathcal{O}_K)$, it preserves the finite-place max norms. Hence

$$h(Tx) - h(x) = \frac{1}{3}\sum_{i=0}^{2} \log \frac{\|\sigma_i(B_n)(\sigma_i(x),1)^t\|_\infty}{\|(\sigma_i(x),1)^t\|_\infty} = \frac{1}{3}\log\frac{n}{\lambda_0^2} + o(1).$$

Taking n large proves the assertion. □

Experiment II.9.5 (Weak-approximation witnesses)**.** The program `weak_approx_barrier.py` solves the three-embedding interpolation problem, rationalizes the coefficients to obtain an exact element of K, verifies the chosen Rosen digit exactly, and compares the one-step height difference with the asymptotic formula above:

n	digit	$h(Tx) - h(x)$	$\frac{1}{3}\log(n/\lambda_0^2)$
10	−10	0.3700403954	0.3749532938
100	−100	1.1424334806	1.1424816582
1000	−1000	1.9100095408	1.9100100225
10000	−10000	2.6775383820	2.6775383868
100000	−100000	3.4450667511	3.4450667511

The error decays by approximately a factor 100 when n is multiplied by ten, as predicted by placing the outer arguments at size n.

The random tests above did not see these adversarially correlated embeddings. They remain useful evidence about typical orbits, but cannot support an additive universal theorem.

Experiment II.9.6 (Adversarial blocks and the sharp factor two)**.** Two new tests target the relative-unit directions which random sampling misses. The program `rosen_block_excursion_tests.py` prescribes a finite Rosen digit block, chooses the two outer projective directions which maximize its conjugate operator norms, and uses the smallest tested common denominator which realizes the requested block exactly. For a single digit $-N$ it found

N	$h(x)$	$h(Tx)$	$h(Tx)/(1+h(x))$
10	6.214608	6.676037	0.925350
100	9.903491	11.056791	1.014060
1000	13.815511	15.726667	1.061500
10000	19.113828	21.791471	1.083407

and the alternating and separated two/four-block families stayed below 1.077 in the tested range.

More dangerous examples arise without a large rational denominator. Let r be the cusp with Rosen word $(-30,-1,30,1)$, and put

$$x_k = r + \lambda^{-k}\alpha^{-2k/3} \qquad (3 \mid k).$$

The perturbing unit is exponentially small at the distinguished embedding and exponentially large at the other two. Exact iteration by `rosen_unit_perturbation_tests.py` gave

k	$h(x_k)$	$\max_n h(T^n x_k)$	$\max_n h(T^n x_k)/(1+h(x_k))$	certificate
36	20.284688	33.388182	1.568648	special
96	42.856451	75.213368	1.714990	cusp
192	78.971242	146.647004	1.833747	cusp
300	119.600381	226.438558	1.877594	special
480	187.315614	362.013761	1.922378	cusp

A separate scan of 132 unit-lattice slopes at exponent 120 found maximum ratio 1.784712; every orbit again had an exact cusp or special certificate. The first table explains the unbounded one-step additive rise, while the unit ray appears to approach the factor 2. It therefore falsifies any small additive overshoot model and supplies a sharp candidate for the global multiplicative exponent.

Proposition II.9.7 (Uniform one-step quadratic bound)**.** *There is a constant C_7 such that every nonzero $x \in K \cap [-\lambda/2, \lambda/2)$ satisfies*

$$H(Tx) \le C_7 H(x)^2, \qquad \textit{equivalently} \qquad h(Tx) \le 2h(x) + \log C_7.$$

Proof. Choose a primitive integral representative $x = [p : q]$, and put

$$P_i = \max(|p_i|, |q_i|), \qquad H(x)^3 = P_0 P_1 P_2.$$

Write

$$y = Tx = -\frac{1}{x} - n\lambda$$

with $x, y \in [-\lambda_0/2, \lambda_0/2)$. The corresponding integral pair is

$$(p', q') = (-q - n\lambda p, p).$$

At the distinguished embedding both $|x_0|$ and $|y_0|$ are less than one. Hence

$$\frac{P_0'}{P_0} = |x_0|.$$

At either outer embedding, the branch matrix has max-norm $O(1 + |n|)$, so

$$\frac{P_i'}{P_i} \le C(1 + |n|), \qquad i = 1, 2.$$

The nearest-λ condition and boundedness of y_0 give

$$1 + |n| \le \frac{C'}{|x_0|}.$$

Since the branch matrix is integral unimodular, the finite-place factors of the height do not change. Multiplying the three archimedean estimates therefore gives

$$\frac{H(y)^3}{H(x)^3} \le \frac{C''}{|x_0|}.$$

Finally, the local contribution of the distinguished embedding to $h(1/x) = h(x)$ yields

$$\frac{1}{3} \log^+ \frac{1}{|x_0|} \le h(x), \qquad \text{hence} \qquad |x_0|^{-1} \le H(x)^3.$$

Thus $H(y)^3 \le C'' H(x)^6$, proving the assertion. □

Proposition II.9.8 (The arithmetic slack in the quadratic bound)**.** *Let $x = p/q$ be represented by coprime $p, q \in \mathcal{O}_K$, with $|x_0| < 1$, and define*

$$\mathscr{S}(x) = h(x) - \frac{1}{3} \log \frac{|q_0|}{|p_0|}.$$

Then $\mathscr{S}(x) \geq 0$,

$$\mathscr{S}(x) \geq \frac{1}{3}\log|\mathrm{N}_{K/\mathbf{Q}}(p)|,$$

and

$$h(Tx) \leq 2h(x) - \mathscr{S}(x) + O_7(1) \leq 2h(x) - \frac{1}{3}\log|\mathrm{N}(p)| + O_7(1).$$

Proof. View $q/p = 1/x$ in affine height coordinates. Since the ideals (p) and (q) are coprime, its finite-place height contribution is $\frac{1}{3}\log|\mathrm{N}(p)|$. At the distinguished real place its positive contribution is $\frac{1}{3}\log|q_0/p_0|$. The remaining two archimedean contributions are nonnegative. Since $h(q/p) = h(x)$, subtracting the distinguished term proves nonnegativity and (II.9.8).

The proof of Proposition II.9.7, before applying the final Liouville estimate, gives

$$h(Tx) - h(x) \leq \frac{1}{3}\log\frac{1}{|x_0|} + O_7(1) = h(x) - \mathscr{S}(x) + O_7(1).$$

This is (II.9.8); the norm form follows from (II.9.8). □

Thus a branch can approach the sharp factor two with a fixed bound on its slack only when its numerator has bounded algebraic norm and its two outer affine contributions are small. Up to multiplication by a unit there are only finitely many such numerators. If branches with a fixed bound on their total slack compound, three consecutive exact Rosen remainders have the form

$$r_j = a_j\varepsilon_j, \qquad r_{j+1} = a_{j+1}\varepsilon_{j+1}, \qquad r_{j+2} = a_{j+2}\varepsilon_{j+2},$$

where the a_i range over a finite norm-bounded set and the ε_i are units, while

$$r_j + n_{j+1}\lambda r_{j+1} + r_{j+2} = 0. \tag{unit recurrence}$$

This is the precise unit-equation remnant of the global problem. It is far smaller than an arbitrary moving-target Roth problem: the coefficient types are finite and the exponent lattice has rank two. The argument below in fact solves the larger finite-prime-support remnant completely.

Lemma II.9.9 (Rational-sum S-unit finiteness)**.** *Let $L/\mathbf{Q}$ be a cyclic cubic field, let $c \in L^\times$, and let $A, B \subset L^\times$ be finite. There are only finitely many triples*

$$(u, v, n) \in (A\mathcal{O}_L^\times) \times (B\mathcal{O}_L^\times) \times \mathbf{Q}^\times \quad \textit{such that} \quad u + v = nc.$$

Proof. Let τ generate $\mathrm{Gal}(L/\mathbf{Q})$, and put $z = u/c$, $w = v/c$. Since $z + w = n \in \mathbf{Q}$,

$$z + w - \tau z - \tau w = 0. \tag{738}$$

As u and v range over the indicated finite unit cosets, all four terms in (738) lie in one finitely generated multiplicative subgroup of $L^\times$. The projective form of the S-unit theorem therefore gives only finitely many nondegenerate projective solutions of (738); see [29].

We check both the word "projective" and all degeneracies. If two solutions determine the same projective four-tuple, say

$$(z', w', -\tau z', -\tau w') = t(z, w, -\tau z, -\tau w),$$

then the first and third coordinates give $\tau(t) = t$, so $t \in \mathbf{Q}^\times$. For fixed coefficient types, the quotient $t = z'/z$ is a fixed element times a unit. Taking norms gives t^3 equal to one of finitely

many nonzero rational numbers, up to sign. Hence only finitely many such t occur. Thus the nondegenerate actual pairs (z, w), not just their projective classes, are finite.

If (738) is degenerate, a vanishing proper subsum has two terms: a three-term subsum would force the remaining nonzero term to vanish. Pairing the four terms gives precisely three possibilities. The pairing $z + w = 0$ forces $n = 0$, which has been excluded. The pairing

$$z = \tau z, \qquad w = \tau w$$

puts z, w in $\mathbf{Q}$. The crossed pairing

$$z = \tau w, \qquad w = \tau z$$

gives $z = \tau^2 z$, and hence again $z, w \in \mathbf{Q}$. Finally, for a fixed coefficient type, $z = a\varepsilon/c \in \mathbf{Q}$ implies

$$z^3 \,\mathrm{N}(c/a) = \mathrm{N}(\varepsilon) = \pm 1,$$

so only finitely many rational z occur; the same holds for w. This proves finiteness also in every degenerate case. □

Lemma II.9.10 (Norm-form divisibility of every Rosen digit)**.** *For three consecutive nonzero integral remainders satisfying*

$$r_{j-1} + n_j \lambda r_j + r_{j+1} = 0,$$

one has the exact divisibility in $\mathbf{Z}$

$$n_j \mid \mathrm{N}(r_{j-1}) + \mathrm{N}(r_{j+1}). \tag{739}$$

Consequently, if $|\,\mathrm{N}(r_{j-1})|, |\,\mathrm{N}(r_{j+1})| \le M$*, then either* $|n_j| \le 2M$*, or the two signed norms cancel exactly:* $\mathrm{N}(r_{j-1}) = -\,\mathrm{N}(r_{j+1})$.

Proof. Because λ, r_{j-1}, r_j are algebraic integers, the norm form

$$P_j(X) = \mathrm{N}_{K/\mathbf{Q}}(r_{j-1} + X\lambda r_j)$$

belongs to $\mathbf{Z}[X]$. Its constant term is $P_j(0) = \mathrm{N}(r_{j-1})$, whereas the recurrence and the odd degree give

$$P_j(n_j) = \mathrm{N}(-r_{j+1}) = -\,\mathrm{N}(r_{j+1}).$$

For every integral polynomial, $P_j(n_j) - P_j(0)$ is divisible by n_j, which proves (739). The final alternative is immediate. □

This identity identifies the only dangerous bounded-norm branch more sharply than qualitative S-unit finiteness does. A large digit between bounded-norm remainders forces exact cancellation of the two outer signed norms. In the all-unit core, every transition with digit of absolute value greater than two is therefore necessarily of opposite-norm type. The four-term Galois unit equation in Lemma II.9.9 is needed precisely to make that exceptional branch finite.

Lemma II.9.11 (The Rosen interval removes projective S-unit scaling)**.** *Let* $L/\mathbf{Q}$ *be cyclic cubic, let* $c \in L^\times$*, fix one real embedding, and let* $\Gamma \subset L^\times$ *be a finitely generated Galois-stable group. Fix* $\theta > 0$*. Among the solutions*

$$u + v = nc, \qquad u, v \in \Gamma, \qquad n \in \mathbf{Z} \setminus \{0\}, \qquad |1/u| \le \theta, \quad |v| \le \theta,$$

there are only finitely many for which the associated four-term Galois relation is nondegenerate.

Proof. Put $z = u/c$, $w = v/c$, and let τ generate the Galois group. As before,

$$z + w - \tau z - \tau w = 0.$$

All four terms lie in a fixed finitely generated multiplicative group, so the projective S-unit theorem gives finitely many nondegenerate projective four-tuples. Fix one such class and a representative $(z_0, w_0, -\tau z_0, -\tau w_0)$. Every actual solution in the class is obtained by multiplication by a scalar t. Comparing the first and third coordinates with Galois conjugation gives $\tau(t) = t$, hence $t \in \mathbf{Q}^\times$.

The two interval conditions become

$$\frac{1}{|cz_0|\theta} \le |t| \le \frac{\theta}{|cw_0|}$$

at the fixed real embedding. Moreover $n = t(z_0 + w_0)$ is a nonzero integer. The displayed compact interval therefore permits only finitely many values of n, and each determines $t = n/(z_0 + w_0)$. Thus each projective class contains only finitely many actual Rosen-compatible solutions. □

Lemma II.9.12 (Two consecutive unit equations remove the scalar)**.** *For a Rosen transition put*

$$z_j = \frac{r_{j-1}}{\lambda r_j}, \qquad w_j = \frac{r_{j+1}}{\lambda r_j}; \qquad z_j + w_j = -n_j.$$

Fix the two projective classes

$$(z_j, w_j, -\tau z_j, -\tau w_j), \qquad (z_{j+1}, w_{j+1}, -\tau z_{j+1}, -\tau w_{j+1}).$$

If both four-term relations are nondegenerate, the two classes admit only finitely many actual consecutive Rosen transitions.

Proof. Choose representatives $(z_0, w_0, -\tau z_0, -\tau w_0)$ and $(z_1, w_1, -\tau z_1, -\tau w_1)$. The Galois compatibility argument in Lemma II.9.11 shows that an actual pair in these classes has

$$(z_j, w_j) = t(z_0, w_0), \qquad (z_{j+1}, w_{j+1}) = s(z_1, w_1), \qquad s, t \in \mathbf{Q}^\times.$$

Consecutive remainders give

$$z_{j+1} = \frac{1}{\lambda^2 w_j},$$

so $st = (\lambda^2 w_0 z_1)^{-1}$ is fixed by the two projective classes. Consequently

$$n_j n_{j+1} = st(z_0 + w_0)(z_1 + w_1)$$

is fixed. For a compatible actual block this is a nonzero integer, and it has only finitely many factorizations as the product of the two nonzero integers n_j, n_{j+1}. The first digit determines t, and the reciprocal identity then determines s. □

Lemma II.9.13 (Degenerate transitions are separated by two steps)**.** *If the four-term relation at transition j is degenerate, then the relations at both $j+1$ and $j+2$ are nondegenerate.*

Proof. The pairing classification in Lemma II.9.9 gives $z_j, w_j \in \mathbf{Q}^\times$. This makes the next relation nondegenerate, since $z_{j+1} = 1/(\lambda^2 w_j) \notin \mathbf{Q}$. If the relation at $j+2$ were degenerate, then $z_{j+2} = 1/(\lambda v_{j+1})$ would be rational. On the other hand

$$\lambda v_{j+1} = -n_{j+1}\lambda^2 - \frac{1}{w_j}.$$

The digit n_{j+1} is nonzero: a nonzero complete quotient lies in $[-\lambda_0/2, \lambda_0/2)$, so its negative reciprocal cannot select the zero digit. Since $\lambda^2 \notin \mathbf{Q}$, the last display is not rational, a contradiction. □

Theorem II.9.14 (Finite-prime-support Rosen orbits close)**.** *Let $x \in K$, and write its primitive exact Rosen remainders as*

$$x_j = r_j/r_{j-1}, \qquad r_{j-1} + n_j\lambda r_j + r_{j+1} = 0.$$

If there is a finite set S of prime ideals of $\mathcal{O}_K$ such that every prime divisor of every nonzero r_j belongs to S, then the orbit terminates or an exact state repeats. In the latter case x is in the G_7-orbit of a fixed point of a special hyperbolic element.

Proof. Enlarge S to be Galois-stable and put $\Gamma = \mathcal{O}_{K,S}^{\times}$. At a nonterminal transition set

$$u_j = \frac{r_{j-1}}{r_j}, \qquad v_j = \frac{r_{j+1}}{r_j}.$$

Then $u_j, v_j \in \Gamma$, $u_j + v_j = -n_j\lambda$, and the two normalized states give

$$|1/u_j| = |x_j| \le \lambda_0/2, \qquad |v_j| = |x_{j+1}| \le \lambda_0/2.$$

The projective S-unit theorem leaves only finitely many nondegenerate projective four-term classes. Lemma II.9.13 says that an infinite orbit contains infinitely many pairs of consecutive nondegenerate transitions. Lemma II.9.12 then removes the rational projective scaling and leaves only finitely many actual two-transition blocks. One exact state therefore repeats. (This also gives a proof independent of the archimedean interval bound in Lemma II.9.11.)

The nonempty word between equal states fixes that state. Every Rosen branch has derivative $1/x_j^2 > 4/\lambda_0^2 > 1$, so this word is neither scalar nor parabolic and is therefore hyperbolic. Conjugating by the finite prefix gives an element of G_7 fixing x; its eigenvalue on $(x,1)^t$ is $cx + d \in K$, so it is special. □

Proposition II.9.15 (The prime-ideal birth and return automaton)**.** *Assume that the initial remainder pair is primitive. For every prime ideal $\mathfrak{p}$ and every j,*

$$\min\{v_{\mathfrak{p}}(r_{j-1}), v_{\mathfrak{p}}(r_j)\} = 0.$$

Moreover:

1. *if $\mathfrak{p} \mid r_j$, then $\mathfrak{p} \nmid r_{j-1}r_{j+1}$;*
2. *if $\mathfrak{p} \mid r_{j-1}$, then $\mathfrak{p} \mid r_{j+1}$ exactly when $\mathfrak{p} \mid (n_j)$, and*

$$\min\{v_{\mathfrak{p}}(r_{j-1}), v_{\mathfrak{p}}(r_{j+1})\} \le v_{\mathfrak{p}}(n_j);$$

3. *if $\mathfrak{p} \nmid r_{j-1}r_j$, a birth at r_{j+1} occurs exactly when*

$$n_j \equiv -\frac{r_{j-1}}{\lambda r_j} \pmod{\mathfrak{p}}.$$

If $0 \le a < b$, put

$$G_{a,b} = B(n_{b-1}) \cdots B(n_a), \qquad B(n) = \begin{pmatrix} -n\lambda & -1 \\ 1 & 0 \end{pmatrix}.$$

Every return $\mathfrak{p} \mid r_a, r_b$ forces

$$\mathfrak{p} \mid (G_{a,b})_{12}.$$

Thus a prime can be recycled only through a divisor of the intervening return continuant; otherwise its later occurrence is a genuinely new birth.

Proof. The branch matrix $B(n_j)$ is unimodular over $\mathcal{O}_K$, so the ideal (r_{j-1}, r_j) is independent of j, proving (II.9.15). The element λ is a unit. Reducing

$$r_{j-1} + n_j \lambda r_j + r_{j+1} = 0$$

modulo $\mathfrak{p}$ gives the three local transition rules, while the ultrametric inequality gives (2). Finally, $(r_b, r_{b-1})^t = G_{a,b}(r_a, r_{a-1})^t$. If $\mathfrak{p} \mid r_a, r_b$, primitivity makes r_{a-1} a unit modulo $\mathfrak{p}$, so the top row modulo $\mathfrak{p}$ forces $(G_{a,b})_{12} = 0$, as claimed. □

Lemma II.9.16 (The time-truncated split modulus)**.** *At transition j, choose a set $\mathcal{B}_{j+1}$ of time-truncated degree-one prime ideals dividing r_{j+1}, with at most one prime above each rational prime. Write $\mathfrak{p} \cap \mathbf{Z} = (\ell_\mathfrak{p})$, and put*

$$M_{j+1} = \prod_{\mathfrak{p} \in \mathcal{B}_{j+1}} \ell_\mathfrak{p}.$$

After identifying each residue field with its prime field, the congruences

$$n_j \equiv -\frac{r_{j-1}}{\lambda r_j} \pmod{\mathfrak{p}} \qquad (\mathfrak{p} \in \mathcal{B}_{j+1})$$

determine one residue class modulo M_{j+1}. The Rosen digit is an integer in that class and satisfies

$$|n_j| \le \frac{C_7}{|x_{j,0}|}.$$

In particular, if $M_{j+1} > 2C_7/|x_{j,0}|$, it is the unique integer in the indicated archimedean interval belonging to that class.

Proof. The birth congruences are Proposition II.9.15(3). Degree one identifies $\mathcal{O}_K/\mathfrak{p}$ with $\mathbf{Z}/\ell_\mathfrak{p}\mathbf{Z}$, and the distinct rational primes are coprime, so the Chinese remainder theorem gives one class modulo M_{j+1}. The nearest-λ rule and $|x_{j+1,0}| \le \lambda_0/2$ give

$$n_j = -\frac{1}{\lambda_0 x_{j,0}} - \frac{x_{j+1,0}}{\lambda_0},$$

which implies the displayed bound. An interval of length less than M_{j+1} contains at most one member of a fixed residue class. □

Corollary II.9.17 (Primitive-divisor reduction)**.** *For a remainder sequence, call $\mathfrak{p} \mid r_j$ a* prime-ideal birth *if $\mathfrak{p} \nmid r_i$ for every $i < j$. If every primitive Rosen remainder sequence in K has only finitely many prime-ideal births, then the full cusp/special-fixed-point classification follows.*

Proof. Every prime ideal in the support of the sequence has a first occurrence. Finitely many births therefore mean finite total prime support, and Theorem II.9.14 applies. □

Corollary II.9.18 (Bounded-norm Rosen orbits close)**.** *Let $x \in K$. Write its exact Rosen orbit as*

$$x_j = \frac{r_j}{r_{j-1}}, \qquad r_{j-1} + n_j \lambda r_j + r_{j+1} = 0, \tag{R}$$

where $r_{-1}, r_0 \in \mathcal{O}_K$ are coprime and n_j is the Rosen digit. If the orbit does not terminate and

$$\sup_{j \ge -1} |\mathrm{N}(r_j)| < \infty,$$

then an exact state repeats. Consequently x has a nonidentity stabilizer in G_7, and hence is a cusp or a fixed point of a special hyperbolic element.

Proof. Only finitely many prime ideals have norm at most the bound in (II.9.18). They contain every prime divisor of every r_j, so Theorem II.9.14 applies. □

Experiment II.9.19 (Exact unit-core transition graph)**.** PARI/GP verifies that λ and $\alpha = \lambda^2 - 1$ generate the full unit group up to sign. The program `rosen_sunit_transition_tests.py` enumerates

$$u = \pm\lambda^r\alpha^s, \qquad x = 1/u \in [-\lambda/2, \lambda/2),$$

selects the Rosen digit exactly, and retains the transition only if $T(x) = -u - n\lambda$ is again a unit. The boxes $|r|, |s| \le 8, 16, 30$ contained respectively $50, 56, 56$ transitions. At bound 30, only 30 transitions lead to another node of this unit graph; its sole closed component is the two-cycle

$$u = \alpha \longleftrightarrow -\alpha, \qquad x = \alpha^{-1} \longleftrightarrow -\alpha^{-1}.$$

All other tested unit transitions leave the all-unit core after at most two edges. This stabilization is not used to claim exhaustive enumeration—Lemma II.9.9 supplies the finiteness proof—but it gives an exact check of the predicted finite graph and its degenerate cases.

Experiment II.9.20 (Degenerate-transition separation audit)**.** The exact program `rosen_sunit_degeneracy_test` constructs 6851 transitions with $z_j, w_j \in \mathbf{Q}^\times$, using rational denominators at most 30 and digits of absolute value at most 8, then follows each orbit for 30 steps. It finds no two degenerate indices at distance less than three, in agreement with Lemma II.9.13; in this sample the first possible return distance is five. This experiment is not used in the proof of the separation lemma, which is algebraic.

Experiment II.9.21 (Adaptive-height audit of the norm core)**.** The program `rosen_sunit_norm_identity_tests.` uses the primitive cusp pair

$$(p, q) = ((0, 60, 30), (1771, 900, 2670))$$

in the basis $(1, \lambda, \lambda^2 - 1)$, perturbs it by $u_m = \lambda^{-3m}\alpha^{-2m}$, and checks (739) at every exact orbit step. For $m = 24, 48, 80$, all $594, 1404, 2335$ tested transitions respectively passed. The maximum-height rows were

m	$h(x_0)$	$\max_j h(x_j)$	$(\log\lvert\mathrm{N}\, r_{j-1}\rvert, \log\lvert\mathrm{N}\, r_j\rvert, \log\lvert\mathrm{N}\, r_{j+1}\rvert)$
24	33.8278	55.1992	(162.944, 164.346, 162.277)
48	60.9138	109.9465	(329.136, 328.878, 327.193)
80	97.0286	182.0621	(539.943, 545.514, 542.862)

Thus the large excursion is norm-dominated: its three logarithmic norms are approximately three times the state height. At the opposite end of the orbit, the $m = 48$ trajectory contains an exact all-unit triple, but its true state height is only $0.5397246\ldots$. Evaluating its 521-digit basis coefficients against conjugates frozen at 1000 digits had produced a spurious height near 249; the adaptive computation, which recomputes the conjugates at twice the coefficient precision plus a margin, removes this cancellation error. Hence this family confirms both parts of the proposed dichotomy: high states carry large norms, while the bounded-norm unit core occurs only after the height has collapsed.

Theorem II.9.14 shows exactly what is still required from the height argument: a nonclosing orbit would have to have unbounded finite support in its exact remainders. It is enough to prove that an orbit capable of producing indefinitely many compounded record excursions has bounded finite support (uniformly bounded norms would be stronger than necessary). Once that implication is obtained, the S-unit theorem supplies the finite-state closing; no 2-adic image calculation enters.

Corollary II.9.22 (No indefinitely repeated sharp triples)**.** *Fix $C > 0$. If a nonterminating Rosen orbit has infinitely many indices j for which*

$$h(x_{k+1}) \geq 2h(x_k) - C \qquad (k = j-1, j, j+1),$$

then the orbit repeats and its initial point is special. Equivalently, on any nonclosing orbit and for every fixed M, all sufficiently late triples of consecutive remainders contain an r_k with $|\,\mathrm{N}(r_k)| > M$.

Proof. Proposition II.9.8 and (II.9.22) imply

$$\frac{1}{3}\log|\,\mathrm{N}(r_k)| \leq C + O_7(1)$$

for $k = j-1, j, j+1$. Hence all three remainders belong, up to units, to a fixed finite set of coefficient types. Dividing their recurrence by the middle remainder and applying Lemma II.9.9 leaves only finitely many possible exact middle states. Infinitely many such indices therefore repeat a state, and Theorem II.9.18 closes the orbit. The equivalent formulation follows by contraposition, using the same finite-state argument for a fixed norm bound M. □

Conjecture II.9.23 (Four-term number-field n-conjecture)**.** *For a fixed number field L and every $\varepsilon > 0$, suppose $x_1 + x_2 + x_3 + x_4 = 0$ and no nonempty proper subsum vanishes. Put*

$$\mathrm{rad}_L([x_i]) = \frac{1}{[L:\mathbf{Q}]} \sum_{\substack{\mathfrak{p} \\ \mathrm{ord}_{\mathfrak{p}}(x_i) \text{ are not all equal}}} \log \mathrm{N}\,\mathfrak{p}.$$

Then

$$h([x_1 : x_2 : x_3 : x_4]) \leq (3+\varepsilon)\,\mathrm{rad}_L([x_i]) + O_{L,\varepsilon}(1). \tag{740}$$

This is the standard strong four-summand generalization of the *abc*-conjecture: the coefficient $3 = 2 \cdot 4 - 5$ is the predicted strong n-conjecture exponent. It is also a special case of the truncated-counting form of Vojta's generalized *abc* conjecture; see [31]. The precise coefficient will not matter below: any uniform constant multiplying the radical is sufficient.

Remark II.9.24 (Uniform n-conjecture versus generic Vojta). There are two different four-term conjectures here, and their constants must not be conflated. The uniform strong n-conjecture used in Conjecture II.9.23 predicts the exponent $2n - 5 = 3$ for $n = 4$. Vojta's truncated-counting conjecture on the hyperplane

$$X = \{x_1 + x_2 + x_3 + x_4 = 0\} \simeq \mathbf{P}^2$$

predicts the sharper generic coefficient $1 + \varepsilon$, because the four coordinate lines form a divisor D with

$$K_X + D = -3H + 4H = H.$$

The latter statement holds only outside a proper Zariski-closed subset, which is essential in four variables. It therefore does not by itself give the fixed proper vanishing subsum used in Proposition II.9.25. One may use the coefficient-one form only after controlling its exceptional curves; the uniform coefficient-three form avoids that extra obligation. Vojta states this distinction explicitly in [31, pp. 6–7].

Proposition II.9.25 (Generalized *abc* supplies the moving unit gap)**.** *Assume Conjecture II.9.23. Let $\Gamma \subset L^\times$ be finitely generated. Suppose*

$$\sum_{i=1}^{4} a_{i,\nu} u_{i,\nu} = 0, \qquad u_{i,\nu} \in \Gamma,$$

and put

$$Q_\nu = h([u_{1,\nu} : \cdots : u_{4,\nu}]), \qquad A_\nu = 1 + \max_i h(a_{i,\nu}).$$

If $Q_\nu \to \infty$ and $A_\nu = o(Q_\nu)$, then, after passage to an infinite subsequence, a fixed nonempty proper subsum vanishes.

Proof. Choose a fixed finite set S of places for which $\Gamma \subset \mathcal{O}_{L,S}^\times$, and set $x_{i,\nu} = a_{i,\nu} u_{i,\nu}$. Outside S and the prime divisors of the coefficient tuple, multiplication by the $u_{i,\nu}$ changes no valuation. The elementary height–radical estimate therefore gives

$$\mathrm{rad}_L([x_{i,\nu}]) \le C_{L,S} \sum_i h(a_{i,\nu}) + O_{L,S}(1) = O_{L,S}(A_\nu).$$

If infinitely many of the four-term relations were nondegenerate, (740) would imply

$$h([x_{1,\nu} : \cdots : x_{4,\nu}]) = O_{L,S}(A_\nu).$$

Changing the four projective coordinates by the coefficients costs at most $\sum_i h(a_{i,\nu}) + O(1)$, so this gives $Q_\nu = O_{L,S}(A_\nu)$, contrary to $A_\nu = o(Q_\nu)$. Hence all sufficiently large relations have a proper vanishing subsum. There are only finitely many subsets of four terms, so one subset is fixed on an infinite subsequence. □

Thus the slow-coefficient principle needed here is not an unnamed new Roth assertion: it follows from a standard four-term *abc*-type conjecture. It is still not a theorem of Ru–Vojta. Their moving-target Subspace Theorem [21] assumes linear nondegeneracy over the field generated by the moving coefficients, whereas the displayed equation itself makes our point linearly degenerate over that field. Fixed coefficient tuples are covered unconditionally by the ordinary S-unit theorem; Conjecture II.9.23 supplies the missing height-uniformity as the coefficient tuple moves.

Route II.9.26 (Why the effective two-variable unit equation does not replace *abc*)**.** One may hope to replace Proposition II.9.25 by a Baker–Wüstholz height bound linear in the moving coefficient height. The effective theorem does give such bounds for

$$ax + by = c, \qquad x, y \in \mathcal{O}_{K,S}^\times;$$

see [30]. The Galois relation (742), however, is homogeneous in four nonzero terms. After projectivizing it is a unit equation in three affine variables, not the effective two-variable equation. The special pairing does not reduce it: writing

$$z - \tau z = -(w - \tau w)$$

introduces $1 - \tau z/z$ and $1 - \tau w/w$, which are precisely the moving nonunit coefficient factors whose time-truncated divisors are being measured. Thus Baker's effective two-variable theorem closes a fixed support after a two-term subsum has been selected, but it does not supply the selection or a coefficient-linear bound for the nondegenerate four-term relation. The attempted shortcut therefore returns to the same time-truncated-boundary term rather than eliminating Conjecture II.9.23.

Route II.9.27 (Tropical dominance does not prove the moving principle)**.** The four Galois terms have leading logarithms

$$u_k,\ v_k,\ u_{k+1},\ v_{k+1} \qquad (k \pmod 3),$$

where $\sum u_k = \sum v_k = 0$. Cancellation only forces the maximum of these four numbers to occur at least twice for every k. This condition has unbounded rays not arising from one fixed global pairing. For example,

$$u = t(-1, 0, 1), \qquad v = t(1, 1, -2), \qquad t > 0,$$

has top-term pairings

$$(v_0, v_1), \qquad (v_1, u_2), \qquad (u_2, v_0)$$

at the three embeddings. Thus the dominant pair rotates cyclically and no two-term subsum is selected globally. The exact program `rosen_sunit_tropical_tests.py` found 1692 nonzero integral slope vectors and 30 top-pair patterns in the box $[-12, 12]^4$. This disproves the proposed elementary *largest two terms must pair* proof. Establishing Proposition II.9.25 requires the arithmetic radical bound in Conjecture II.9.23; convex geometry alone cannot control these rotating tropical cones.

Proposition II.9.28 (Conditional three-remainder norm gap)**.** *Assume Conjecture II.9.23. There are constants $\eta_7 > 0$ and B_7 such that every exact nonterminal Rosen state*

$$x_j = r_j / r_{j-1}$$

with primitive consecutive remainders satisfies

$$\max_{k=j-1,j,j+1} \log^+ |\,\mathrm{N}(r_k)| \geq \eta_7 h(x_j) - B_7. \tag{741}$$

Proof. For every nonzero $r \in \mathcal{O}_K$, reduction in the logarithmic unit lattice gives a factorization

$$r = a\varepsilon, \qquad \varepsilon \in \mathcal{O}_K^\times,$$

such that

$$\left| \log |\sigma_i(a)| - \frac{1}{3} \log |\,\mathrm{N}(r)| \right| \leq C_K \quad (i = 0, 1, 2).$$

Indeed, translate the logarithmic embedding of r into a fixed fundamental parallelogram for the rank-two unit lattice. In particular

$$h(a) \leq \frac{1}{3} \log^+ |\,\mathrm{N}(r)| + O_K(1).$$

Suppose (741) were false. There would be a sequence of states with $h(x_j) \to \infty$ and

$$D_j := 1 + \max_{k=j-1,j,j+1} \log^+ |\,\mathrm{N}(r_k)| = o(h(x_j)).$$

Use (II.9) to write $r_k = a_k \varepsilon_k$, and set

$$u = \frac{r_{j-1}}{r_j}, \qquad v = \frac{r_{j+1}}{r_j}, \qquad z = \frac{u}{\lambda}, \qquad w = \frac{v}{\lambda}.$$

The Rosen recurrence gives $z + w = -n_j \in \mathbf{Q}^\times$, and hence

$$z + w - \tau z - \tau w = 0, \tag{742}$$

where τ generates $\mathrm{Gal}(K/\mathbf{Q})$. After the balanced factorizations, the four terms in (742) are four units, belonging to the fixed group $\mathcal{O}_K^\times$, multiplied by coefficients of height $O(D_j)$.

Their projective unit height tends to infinity on the scale of $h(x_j)$. To see this, the first unit quotient differs from $u = 1/x_j$ by a coefficient of height $O(D_j)$. Moreover the map

$$\mathcal{O}_K^\times/\{\pm 1\} \longrightarrow \mathcal{O}_K^\times, \qquad \varepsilon \longmapsto \varepsilon/\tau\varepsilon,$$

is an injective linear map of rank-two logarithmic lattices, so its height is bounded below by a positive constant times $h(\varepsilon)$. The four-tuple contains both ε and $\tau\varepsilon$. Consequently its projective height is $\gg_K h(x_j) - O(D_j)$. Proposition II.9.25 applies.

Thus an infinite subsequence of (742) has a proper vanishing subsum. Since all four terms are nonzero, a proper subsum has two terms. Exactly as in Lemma II.9.9, the pairing $z + w = 0$ would give $n_j = 0$, while either remaining pairing gives

$$z, w \in \mathbf{Q}.$$

Write $z = a/b$ in lowest terms. Since

$$\frac{r_{j-1}}{r_j} = \frac{a}{b}\lambda$$

and the consecutive remainder ideals are coprime, valuation at every prime ideal shows

$$(r_{j-1}) = (a), \qquad (r_j) = (b).$$

Here it is essential that λ is a unit: for a rational prime dividing a or b, all prime ideals above it occur with the same valuation. Hence, for one unit ε,

$$r_{j-1} = a\lambda\varepsilon, \qquad r_j = b\varepsilon, \qquad x_j = \frac{b}{a\lambda}.$$

It follows that

$$h(x_j) \le \log\max\{|a|, |b|\} + O_7(1) \le \frac{1}{3}D_j + O_7(1),$$

contradicting (II.9). This proves the asserted uniform gap by the usual contradiction argument. □

Conditionally, Proposition II.9.28 is the global contribution of the moving S-unit equation: at every sufficiently high state, at least one of the three neighboring branches carries a norm slack proportional to the state height. It is stronger than the fixed-bound Corollary II.9.22. By itself it does not yet prove the uniform height excursion bound (754): a scalar sequence can have a proportional saving once in every three steps and still grow. The remaining task is to couple the location of the large norm with the two outer Hodge coordinates, rather than discard that location in the maximum in (741).

The finite-support theorem changes the most economical global target. It is unnecessary to control all relative-unit directions by height. For a primitive remainder pair define the projectively unit-invariant integer energy

$$E_j = |\,\mathrm{N}(r_{j-1})| + |\,\mathrm{N}(r_j)|.$$

Multiplying both remainders by a unit does not change E_j.

There is an exact integral linearization of this energy. For $p, q \in \mathcal{O}_K$, write

$$p\tau(p)\tau^2(q) = b_0 + b_1\lambda + b_2\alpha, \qquad p\tau(q)\tau^2(q) = c_0 + c_1\lambda + c_2\alpha$$

and set

$$\Phi(p,q) = (A, b_0, b_1, b_2, c_0, c_1, c_2, D)^t, \qquad A = \mathrm{N}(p), \quad D = \mathrm{N}(q).$$

Lemma II.9.29 (Integral Asai norm cocycle)**.** *Under the Rosen branch* $(p,q)\mapsto(p',q')=(-n\lambda p-q,p)$*, one has*

$$\Phi(p',q')=\mathcal{A}(n)\Phi(p,q),$$

where

$$\mathcal{A}(n)=\begin{pmatrix} n^3 & 2n^2 & -4n^2 & -n^2 & -n & -5n & -3n & -1\\ -n^2 & 0 & 2n & 2n & 1 & 1 & 1 & 0\\ n^2 & 0 & -n & -2n & 0 & 0 & 1 & 0\\ 0 & n & -n & n & 0 & -1 & -1 & 0\\ 0 & -1 & 0 & -1 & 0 & 0 & 0 & 0\\ -n & 0 & 1 & 1 & 0 & 0 & 0 & 0\\ 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0\\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}\in \mathrm{SL}_8(\mathbf{Z}[n]).$$

In particular the content of the eight integral coordinates is invariant, the two corner coordinates are $(\mathrm{N}(p'),\mathrm{N}(q'))$*, and the first and last rows recover* $n\mid \mathrm{N}(p')+\mathrm{N}(q)$*.*

Proof. Expand the three Galois factors of $(-n\lambda p-q)\tau(-n\lambda p-q)\tau^2(-n\lambda p-q)$ and the two mixed products in the integral basis $(1,\lambda,\alpha)$, using $\lambda^2=1+\alpha$, $\lambda\alpha=\lambda+\alpha$, and $\alpha^2=1+\lambda+\alpha$. This gives the displayed rows. Direct elimination gives $\det\mathcal{A}(n)=1$. The last row is $D'=A$, while the first row is congruent to $-D \pmod n$, proving the remaining claims. □

Lemma II.9.30 (The ideal cube inside the Asai coordinates)**.** *With* A,b,c,D *as above, one has*

$$\mathrm{N}(b)=A^2D,\qquad \mathrm{N}(c)=AD^2,\qquad b\tau^2(c)=AD.$$

For a prime ideal $\mathfrak{p}$*, set*

$$a_i=v_{\mathfrak{p}}(\tau^i p),\qquad d_i=v_{\mathfrak{p}}(\tau^i q)\quad(0\le i<3).$$

If $(p,q)=\mathcal{O}_K$*, then* $\min(a_i,d_i)=0$ *for each* i*, and the valuations of the eight split Asai coordinates are precisely*

$$e_0+e_1+e_2,\qquad e_i\in\{a_i,d_i\}.$$

In particular their minimum is zero. Prime by prime, the twisted Segre point is therefore a three-dimensional binary valuation cube, not an arbitrary vector in K^8*.*

Proof. The first two identities follow by taking the three conjugates of $b=p\tau(p)\tau^2(q)$ and $c=p\tau(q)\tau^2(q)$. Direct multiplication gives

$$b\tau^2(c)=p\tau(p)\tau^2(p)\,q\tau(q)\tau^2(q)=AD.$$

Primitivity applied at the three conjugate primes gives $\min(a_i,d_i)=0$. The split tensor coordinates choose independently one of p,q in each conjugate factor, proving (II.9.30) and its final assertion. □

The chosen coordinates identify Φ with the twisted Segre embedding of $\mathrm{Res}_{K/\mathbf{Q}}\mathbf{P}^1$. Standard functoriality of height gives

$$h_{\mathbf{P}^7}(\Phi(p,q))=3h_{\mathbf{P}^1(K)}([p:q])+O_7(1).\tag{743}$$

Consequently the sextic norm barrier below is exactly a quadratic bound for the two corner coordinates of the integral cocycle $\mathcal{A}(n_j)\cdots\mathcal{A}(n_0)$ on the Rosen-admissible part of the twisted Segre variety. It does not ask for a quadratic bound on the other six coordinates.

Route II.9.31 (A preserved quadratic energy does not exist)**.** The program `rosen_asai_norm_cocycle.py` solves, coefficient by coefficient in n, the linear systems

$$\mathcal{A}(n)^t Q \mathcal{A}(n) = Q.$$

Among symmetric 8×8 matrices the solution space has dimension zero. Among alternating matrices it has dimension one, generated by

$$J = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & -3 & -1 & -2 & 0 \\ 0 & 0 & 0 & 0 & -1 & 2 & -3 & 0 \\ 0 & 0 & 0 & 0 & -2 & 4 & 1 & 0 \\ 0 & 3 & 1 & 2 & 0 & 0 & 0 & 0 \\ 0 & 1 & -2 & -4 & 0 & 0 & 0 & 0 \\ 0 & 2 & 3 & -1 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

This is the expected product of the three standard symplectic forms in the tensor representation. More precisely, direct multiplication gives the useful two-time identity

$$\Phi(p,q)^t J \Phi(r,s) = \mathrm{N}(ps - qr). \tag{744}$$

Thus the pairing vanishes exactly when the two projective field points coincide; for $(r,s) = g(p,q)$ it is the norm of the fixed-point form $Q_g(p,q)$. Since $X^t J X = 0$ for every single vector, however, it cannot be a positive or signed scalar energy controlling the corner coordinates. There is also no nonzero common fixed linear functional. The norm barrier must therefore use the Rosen-admissible cone or the two-time pairing (744); it cannot follow from a globally preserved quadratic form.

Route II.9.32 (The symplectic integer does not force the record to close)**.** For the initial vector $X_0 = \Phi(r_0, r_{-1})$ and a record vector X_j, integrality would prove repetition if $|X_0^t J X_j| < 1$. The adversarial unit ray has the opposite behavior. At $m = 12, 24, 48, 80$, the exact values of $\log|X_0^t J X_j|$ at the maximum norm record are

$$159.869, \qquad 263.238, \qquad 510.964, \qquad 836.532.$$

They differ from the elementary product upper scale $\log E_j + 3h(x_0)$ by only $-1.05, -2.81, -1.49, -0.62$, respectively. Thus the integral fixed-point form is nearly as large as its trivial height bound at the sharp record; it does not become an integer of absolute value below one. The pairing remains an exact recognition certificate for repetition, but a proof of the norm barrier needs a different inequality before that certificate can be invoked.

Proposition II.9.33 (Exact norm–unit height decomposition)**.** *For a nonterminal primitive remainder pair $[r_j : r_{j-1}]$, put*

$$A_j = \log|\,\mathrm{N}\, r_j|, \qquad M_j = \max\{A_{j-1}, A_j\},$$

and

$$\ell_{j,i} = \log\left|\frac{\sigma_i(r_j)}{\sigma_i(r_{j-1})}\right|, \qquad U_j = \sum_{i=0}^{2} |\ell_{j,i}| - |A_j - A_{j-1}|.$$

Then $U_j \geq 0$ and

$$6h([r_j : r_{j-1}]) = 2M_j + U_j. \tag{745}$$

Moreover,

$$U_j = 2\left(\min\{A_{j-1}, A_j\} - \sum_{i=0}^{2} \min\{\log|\sigma_i(r_{j-1})|, \log|\sigma_i(r_j)|\}\right).$$

Thus $U_j/2$ is the archimedean gcd defect of the primitive pair: it measures relative-unit imbalance after the smaller algebraic norm has been removed.

Proof. Primitivity makes the finite-place contribution to the projective height zero. If $a_i = \log|\sigma_i(r_j)|$ and $b_i = \log|\sigma_i(r_{j-1})|$, then

$$3h([r_j : r_{j-1}]) = \sum_i \max\{a_i, b_i\} = \frac{A_j + A_{j-1} + \sum_i |a_i - b_i|}{2}.$$

Now $A_j + A_{j-1} + |A_j - A_{j-1}| = 2M_j$, which proves (745). The formula with the coordinatewise minimum follows from $\sum_i \min(a_i, b_i) = (A_j + A_{j-1} - \sum_i |a_i - b_i|)/2$. □

Proposition II.9.34 (The unit reservoir is a lower rank-two weight)**.** *With the notation of Proposition II.9.33, put*

$$P_j = \sum_i \max\{\ell_{j,i}, 0\}, \qquad N_j = \sum_i \max\{-\ell_{j,i}, 0\}.$$

Then

$$3h([r_j : r_{j-1}]) = A_{j-1} + P_j = A_j + N_j, \qquad \frac{U_j}{2} = \min\{P_j, N_j\}. \tag{746}$$

As functions of $\ell = (\ell_0, \ell_1, \ell_2)$, P and N are the support functions of the cubes $[0,1]^3$ and $[-1,0]^3$, respectively. However, U is neither convex nor concave, even on the Rosen half-space $\ell_0 < 0$. In particular U is not the support function of a polytope and cannot by itself be the local weight of one semipositive toric line metric.

Proof. Write $a_i = \log|\sigma_i(r_j)|$ and $b_i = \log|\sigma_i(r_{j-1})|$. Since $\ell_i = a_i - b_i$,

$$\sum_i \max\{a_i, b_i\} = \sum_i b_i + \sum_i \max\{\ell_i, 0\} = A_{j-1} + P_j.$$

Interchanging the coordinates gives $A_j + N_j$. Also

$$\sum_i |\ell_i| - \left|\sum_i \ell_i\right| = P_j + N_j - |P_j - N_j| = 2\min\{P_j, N_j\}.$$

The support-function assertions follow directly from $P(\ell) = \max_{v\in[0,1]^3}\langle v, \ell\rangle$ and the analogous formula for N.

For convexity take $x = (-5, -5, -3)$ and $y = (-1, 0, 5)$. All relevant first coordinates are negative, but

$$U(x) = 0, \qquad U(y) = 2, \qquad U(x+y) = 4 > U(x) + U(y).$$

A positively homogeneous convex function is subadditive, so U is not convex. For concavity take $x = (-5, -5, -5)$ and $y = (-5, -5, 1)$. Then

$$U(x+y) = 0 < 2 = U(x) + U(y),$$

which violates the superadditivity of a positively homogeneous concave function. □

The first equality in (746) is best read as two equal global ledgers:

$$(\text{denominator norm}) + (\text{real expansion}) = 3h, \qquad (\text{numerator norm}) + (\text{real contraction}) = 3h.$$

The ledger with larger finite part has the smaller archimedean part, namely $U/2$. Thus the rank-two object is not an invented replacement for the Weil height. It is the standard projective height with its two coordinate accounts kept separate long enough to remember a transfer from one account to the other.

Route II.9.35 (Why the global object has rank at least two)**.** The preceding proposition corrects a tempting but false geometric picture: there is no "reservoir hexagon" whose support function is U. The two convex quantities P and N must first be kept as the two weights of a split rank-two norm. At a given archimedean state, sorting those weights gives lower weight $\min(P, N) = U/2$. Only after this sorting should one form a global filtered degree. Calling U a single toric height loses precisely the choice of which coordinate is the smaller one, and it incorrectly assumes semipositivity of a nonconvex function. Thus the user's global-Arakelov objection is not just stylistic: it forces a different geometric object.

Experiment II.9.36 (The raw reservoir is not locally monotone)**.** The program `rosen_reservoir_monotonicity_tests.py` evaluates

$$E_j = \max_{k \le j} M_k + U_j$$

at every step of the adversarial unit rays. The stronger assertion $E_{j+1} \le E_j$ fails often and by an amount which grows along this family:

m	tested steps	$\max(E_{j+1} - E_j)$	$\#\{j : E_{j+1} > E_j\}$
1	90	3.854244730	38
2	117	6.216691558	52
3	141	8.001825491	60
4	143	9.691829087	63
5	148	6.851367033	66
6	162	6.476695328	72
7	203	7.514554787	95
8	179	8.805209769	84
9	275	9.886481857	121
10	264	12.749073689	110
11	226	16.139821056	100
12	439	19.630077164	198

All twelve exact trajectories close. Thus the conjectural boundedness of E_j cannot be proved by making the displayed raw quantity decrease branch by branch. A successful global height must contain either a Hodge-theoretic transfer term whose change absorbs these jumps, or a record-to-record inequality which permits reversible motion inside one fixed global budget.

Experiment II.9.37 (A signed Hodge correction exists but has no lower bound)**.** The program `rosen_reservoir_transfer_tests.py` asks whether the one-step failures in Experiment II.9.36 can be absorbed by a scalar transfer function. First it permits

$$\Phi = a(H_0 + H_1) + b(H_0 + H_2) + cH_0,$$

where $H_i = \log \max\{|\sigma_i(r_j)|, |\sigma_i(r_{j-1})|\}$. On all 2387 nonterminal transitions of the first twelve unit rays, a linear program finds coefficients

$$(a, b, c) = (2.5951484\ldots, 2.1079440\ldots, 17.7580748\ldots)$$

for which $E_j + \Phi_j$ never rises. This does not yield a height argument: H_0 tends to $-\infty$ down a cusp, so Φ has no global lower bound. Its negative divergence merely supplies unlimited credit for later reservoir growth.

There is an exact way to impose the missing lower bound. Put $t = -\ell_0 > 0$, the distinguished cusp depth. Then

$$P \geq 0, \qquad N - t = \sum_{i=1}^{2} \max(-\ell_i, 0) \geq 0, \qquad t \geq 0.$$

Every linear expression in P, N, t which is bounded below throughout the full logarithmic ratio space can, up to an additive constant, be written

$$\Phi_{\mathrm{pos}} = aP + b(N - t) + dt, \qquad a, b, d \geq 0.$$

Indeed, the rays $\ell_1 \to +\infty$, $\ell_1 \to -\infty$, and $t \to +\infty$ force respectively the three nonnegative coefficients in this display. The linear program minimizes the largest step of $E + \Phi_{\mathrm{pos}}$ over this entire cone. Its optimum is still

$$10.0141665552\ldots > 0,$$

attained at $(a, b, d) = (0, 1.1649797\ldots, 0.3842892\ldots)$. Thus even on the finite test set no globally lower-bounded scalar polyhedral correction can make the reservoir branchwise monotone.

This experiment identifies the obstruction more sharply than the raw jump table. The Hodge Busemann coordinate has enough signed drift, but turning that drift into a genuine height destroys scalar monotonicity. At this stage it suggested keeping two ordered slopes and comparing norm-record times. Proposition II.9.46 below shows that suggestion was still too strong: the second slope can recharge without bound in a single exact inverse branch. The scalar-program failure was diagnosing the falsity of the reservoir, not demanding a higher-rank replacement.

Proposition II.9.38 (Exact record-time decomposition)**.** *For a nonterminal Rosen state define*

$$H_{i,j} = \log \max\{|\sigma_i(r_j)|, |\sigma_i(r_{j-1})|\}, \qquad \mathcal{Q}_j = (H_{0,j} + H_{1,j}) + (H_{0,j} + H_{2,j}),$$

and put $R_j = \max_{k \leq j} M_k$. Since the central Rosen ratio has absolute value less than one,

$$\mathcal{Q}_j = H_{0,j} + 3h(x_j), \qquad H_{0,j} = \log |\sigma_0(r_{j-1})|. \tag{747}$$

Consequently the reservoir energy has the exact form

$$E_j := R_j + U_j = 2\mathcal{Q}_j - (2H_{0,j} + M_j) + (R_j - M_j). \tag{748}$$

At a norm-record time, when $M_j = R_j$, the last term vanishes:

$$E_j = 2\mathcal{Q}_j - (2H_{0,j} + M_j). \tag{749}$$

Proof. Primitivity gives $H_{0,j} + H_{1,j} + H_{2,j} = 3h(x_j)$. The Rosen interval gives $|\sigma_0(r_j)| < |\sigma_0(r_{j-1})|$, proving (747). Substitute $U_j = 6h(x_j) - 2M_j$ and $2\mathcal{Q}_j = 2H_{0,j} + 6h(x_j)$ to obtain (748). □

This identity explains why the rejected route looked plausible without prime terminology. The two summands making up $\mathcal{Q}_j$ are the coefficient-max versions of the pairwise quantities in the McMullen descent estimate; on a fixed return box the difference from the actual Hodge norms is uniformly bounded. The term $2H_{0,j} + M_j$ couples central Busemann motion to a finite-place norm; it is the required cusp compensation. Finally $R_j - M_j$ remembers a past record while the current norm has fallen. The attempted proof split was to control (749) at record times and then bound the rebuilding of (748) between records. The second assertion is false. The corrected argument retains only the norm record and works at its first-hitting time.

The extended output of `rosen_reservoir_monotonicity_tests.py` tests the second item. For the first twelve adversarial unit rays it gives

m	1	2	3	4	5	6	7	8	9	10	11	12
# norm records	2	4	2	7	3	4	6	5	5	4	4	11
max recharge	4.50	5.87	7.27	9.27	6.08	6.07	7.16	8.70	8.09	8.07	9.96	6.52

Although individual positive steps become much larger, the total rise after the most recent record remains below 9.96 in this sample. This finite table originally looked like evidence for a record-to-record mechanism. It is included because the balanced inverse family below shows exactly how such finite evidence can be misleading.

Route II.9.39 (Rejected record control and bounded recharge)**.** Let $0 = a_0 < a_1 < a_2 < \cdots$ be the strict norm-record times, so that $M_{a_\nu} = R_{a_\nu}$. There are field constants $B_{\mathrm{rec}}, B_{\mathrm{gap}}$ such that

$$E_{a_\nu} \le 6h(x_0) + B_{\mathrm{rec}}, \tag{750}$$

$$E_j \le E_{a_\nu} + B_{\mathrm{gap}} \qquad (a_\nu \le j < a_{\nu+1}). \tag{751}$$

If there is a final record time, the second inequality is required on the whole remaining tail.

Proposition II.9.40 (The two-stage estimate is sufficient)**.** *The two inequalities displayed in Route II.9.39 imply the reservoir inequality (752), with $B_7 = B_{\mathrm{rec}} + B_{\mathrm{gap}}$, and hence the full classification.*

Proof. Every index lies after a most recent record time. Apply (751) and then (750). The reservoir implication is Proposition II.9.43. □

The point of this reformulation is mathematical, not cosmetic. Equation (750) is where the balanced Hodge norms and the central finite-place compensation must be combined. Equation (751) is a finite-state transition assertion about the two equal height ledgers. It may use the S-unit equation, but it does not need a sum over newly born primes.

Experiment II.9.41 (Recharge across seven coefficient scales)**.** The exact-state program `rosen_record_recharge_tests.py` samples 100 primitive initial states in each coefficient box and follows every orbit to a cusp or special repetition. With random seed 83921 it gives

coefficient bound	max recharge	mean recharge	cusp	special
10	12.1851466	2.3493232	72	28
10^2	12.2515530	4.2322713	82	18
10^3	12.6207524	5.1794733	79	21
10^4	12.9168119	5.2703337	76	24
10^5	12.2902344	5.5366160	73	27
10^6	12.8397657	5.8738657	77	23
10^7	12.9408243	6.0548793	75	25

No trajectory was unresolved. The worst value changes little across six orders of magnitude in the coefficient bound, while the mean grows slowly. A separate 300-point box of bound 10^5 reached $13.1049599\ldots$. These data originally suggested an additive field constant in (751). Proposition II.9.46 below refutes it; the random boxes simply did not sample the relevant balanced-unit inverse branches.

Route II.9.42 (Rejected unit-to-norm reservoir inequality)**.** There is a constant B_7 such that every primitive exact Rosen orbit satisfies

$$\max_{0\le k\le j} M_k + U_j \le 6h([r_0 : r_{-1}]) + B_7 \qquad (j \ge 0). \tag{752}$$

This was proposed as a structured strengthening of the sextic norm barrier. The initial budget $6h(x) = 2M_0 + U_0$ consists of two copies of initial norm mass and one relative-unit reservoir. A sharp cusp excursion may convert that reservoir, or the unused second norm copy, into a new norm record. The running maximum in (752) says that the same mass cannot then be used again to create an independent record. This is exactly the noncompounding mechanism which the cumulative radical and local birth-payment approaches failed to express.

Proposition II.9.43 (Why the rejected strengthening would have finished)**.** *The inequality in Route II.9.42 implies Conjecture II.9.51, and hence the full classification.*

Proof. Since $U_j \ge 0$, equation (752) gives

$$\max_{k\le j} M_k \le 6h(x) + B_7.$$

In particular every $\log|\mathrm{N}\, r_j|$ has this bound. Hence $\log(1 + E_j) \le 6h(x) + B_7 + O(1)$, which is (753). Proposition II.9.52 then applies. □

Experiment II.9.44 (Exact stress test of the unit reservoir)**.** The program `rosen_norm_reservoir_tests.py` evaluates

$$\max_{k\le j} M_k + 6h(x_j) - 2M_j - 6h(x_0),$$

which equals the left side of (752) minus its main term without separately subtracting nearly equal embedding logarithms. On the same 300 random primitive integral inputs used for the sextic test, all 226 cusp and 74 special orbits satisfy the proposed bound; the largest excess is $-5.2399163\ldots$. A further 200 primitive pairs with rational denominators through 2000 again all close, with largest excess $-12.7659856\ldots$.

For the twelve fully computed adversarial unit rays the excesses range from -15.68 to -29.80. Following the sharper rays $m = 16, 24, 48, 80, 120, 160$ for 200 exact steps gives excesses between -26.77 and -31.61, while their norm records approach the sharp sextic coefficient. An exhaustive coefficient box $[-6, 6]^3$, after Rosen normalization, contains 2184 nonzero states; all close, but the largest reservoir excess is $3.33819499\ldots$. Thus the constant-free version is false at small height. No growing excess was found in these samples, but Proposition II.9.46 gives a growing exact family. This is a useful warning: broad random searches and the earlier sharp unit ray both missed a different, balanced direction in the rank-two unit lattice.

Experiment II.9.45 (Backward search for a second reservoir conversion)**.** The program `rosen_preperiod_reservoir_search.py` prepends certified exact inverse Rosen branches to a known special cycle. Every candidate therefore has a proved closing future, while the preperiod is selected

adversarially. One half of the beam maximizes the largest reservoir excess; the other half maximizes the excess occurring only after two strict increases beyond the initial norm record.

At depth 30, digit bound 10,000, and beam width 500, the search inspected 204,257 exact inverse edges. The largest reservoir excess was $3.8903182933\ldots$. This is another finite contribution to the necessary constant B_7, not a growing-height violation. The largest excess after two later norm-record increases was only $2.0542226311\ldots$, attained at depth 14; it did not improve in the remaining sixteen levels. Thus this targeted test found no second sharp conversion of unit mass into norm mass.

The beam is not exhaustive. In fact it missed the elementary inverse family below, because its score and retained base cycles did not explore balanced outer-unit slopes. The negative outcome is therefore recorded as a failed search heuristic, not as evidence for boundedness.

Proposition II.9.46 (Balanced inverse branches refute the reservoir)**.** *Neither the reservoir inequality* (752) *nor the recharge inequality* (751) *holds with a field constant.*

Proof. Put $\alpha = \lambda^2 - 1$ and $\theta = \lambda^{-1}\alpha^{-3}$. Its logarithmic absolute values at the three real embeddings are

$$(-(a+b), a, b) = (-3.0176233539\ldots, 1.4717598469\ldots, 1.5458635070\ldots).$$

If $L = \max(a, b)$, then $2(a+b) - 3L = 1.3976561868\ldots > 0$.

Let $y_m = \theta^m$, choose n_m four times larger than all three $|\sigma_i(y_m)/\sigma_i(\lambda)|$, up to a fixed additive constant, and set

$$q_m = y_m + n_m\lambda, \qquad x_m = -q_m^{-1}.$$

Since the distinguished value of y_m tends to zero, the first Rosen digit of x_m is exactly n_m; its next state is exactly y_m. Every conjugate of q_m has absolute value greater than one and

$$\log|\mathrm{N}\, q_m| = 3Lm + O(1).$$

Thus $6h(x_m) = 2\log|\mathrm{N}\, q_m|$ and $U(x_m) = 0$. At the next state the two adjacent norms are one, the old record remains $\log|\mathrm{N}\, q_m|$, and

$$U(y_m) = 6h(y_m) = 2(a+b)m.$$

The reservoir excess over $6h(x_m)$ is therefore

$$\big(2(a+b) - 3L\big)m + O(1) \to +\infty.$$

The recharge from the preceding record is $2(a+b)m$, which is also unbounded. □

The program `rosen_balanced_unit_inverse_tests.py` verifies every branch exactly. At $m = 4, 8, 12, 16, 20$ the excess is

$$-0.20699\ldots,\ 6.32498\ldots,\ 12.87358\ldots,\ 18.95953\ldots,\ 24.59643\ldots.$$

At the violating next state the running norm record is still $\log|\mathrm{N}\, q_m| = 3h(x_m)$, so the violation is not a violation of the sextic norm barrier. Exact full-orbit audits for $1 \le m \le 20$ satisfy that barrier as well; at $m = 20$ the slack is $-45.592754\ldots$. This is finite evidence, not an asymptotic norm bound for the balanced family.

Route II.9.47 (Correct global formulation: first norm-hitting time)**.** Let M_j be the current adjacent norm maximum and, for $T > 0$, set

$$\tau(T) = \min\{j : M_j \geq T\}.$$

The sextic norm barrier is equivalent to

$$T \leq 6h(x_0) + B_7 \quad \text{whenever } \tau(T) < \infty.$$

This formulation is genuinely global but asks for no inter-record control. The finite contribution at $\tau(T)$ is the sum of all prime-ideal valuations of the primitive remainder, while the two McMullen Hodge metrics and the central cusp metric are the archimedean terms. The precise missing step is to prove their adelic sum at this first-hitting time using the product formula.

Experiment II.9.48 (The Hodge drop and cusp correction cannot be separated)**.** The identity

$$M_j = \mathcal{Q}_j - H_{0,j} - \frac{1}{2}U_j$$

suggests bounding $-H_{0,j} - U_j/2$ independently and using Hodge descent only for $\mathcal{Q}_j$. The exact program `rosen_record_hodge_compensation_tests.py` balances each initial pair by one unit and tests this at strict norm records. With seed 8127, 500 inputs of coefficient size at most 10^5 produced 1146 records. The largest value of

$$-H_{0,j} - \frac{1}{2}U_j - 2h(x_0)$$

was $165.0638492\ldots$. At the same record the sextic norm slack was

$$M_j - 6h(x_0) = -17.8291040\ldots.$$

Thus the cusp correction can be enormous while the desired norm inequality remains strong: $\mathcal{Q}_j$ has fallen by the compensating amount. A valid Arakelov proof must sum these terms before estimating; separate upper bounds destroy the cancellation.

Route II.9.49 (Eight stopping-time chambers)**.** Put

$$b_i = \log|\sigma_i(r_{j-1})|, \qquad \ell_i = \log\left|\frac{\sigma_i(r_j)}{\sigma_i(r_{j-1})}\right|, \qquad B = \sum_i b_i, \quad s = \sum_i \ell_i.$$

The central Rosen condition gives $\ell_0 < 0$, and direct substitution gives

$$M_j = B + \max(s, 0), \qquad \mathcal{Q}_j = B + b_0 + \max(\ell_1, 0) + \max(\ell_2, 0).$$

Thus a first norm record lies in one of only eight chambers: two choices for the larger adjacent norm and four outer sign patterns. On each chamber the norm divisor and the pairwise Hodge expression are linear in the logarithms. The proposed Arakelov construction should use the actual larger remainder as its integral section and prove one placewise-summed inequality on each chamber. This avoids asking for one convex scalar metric which realizes every maximum, while retaining the essential Hodge–cusp cancellation inside each case.

Experiment II.9.50 (Prefix height is not the missing global bound)**.** Writing $r_j = a_j x + b_j$, the elementary linear-form estimate would give the sextic norm barrier if the projective height of $[a_j : b_j]$ were $h(x) + O_7(1)$ at norm records. The exact program `rosen_record_prefix_height_tests.py` tracks these coefficient pairs. On the same 500 inputs and 1146 records, the largest prefix height was $127.7233656\ldots$ for an initial point of height only $7.7914473\ldots$, a ratio of $16.3927652\ldots$. Nevertheless the norm slack at that record was $M_j - 6h(x_0) = -17.8291040\ldots$.

Thus the prefix and the small central linear form must also be retained together. Bounding the prefix first and then applying a generic height inequality discards the same cancellation as the separated cusp estimate.

Conjecture II.9.51 (Sextic norm-excursion barrier)**.** *There is a constant B_7 such that every primitive exact Rosen orbit starting at $x = [r_0 : r_{-1}]$ satisfies*

$$\log(1 + E_j) \le 6h(x) + B_7 \qquad (j \ge 0). \tag{753}$$

The coefficient six is forced by degree three together with the sharp height exponent two: always $E_j \le 2\exp(3h(x_j))$, so the conjectural quadratic height barrier would imply (753). The converse is not required and need not hold; E_j deliberately forgets the relative-unit directions, which Theorem II.9.14 already controls.

Proposition II.9.52 (The norm barrier proves the full classification)**.** *Conjecture II.9.51 implies the full cusp/special-fixed- point classification.*

Proof. For a fixed initial x, equation (753) bounds $|\,\mathrm{N}(r_j)|$ uniformly in j. Corollary II.9.18 says that the exact Rosen orbit terminates or repeats. Termination gives a cusp, while repetition gives a special hyperbolic fixer as in Theorem II.9.14. □

Experiment II.9.53 (The sextic exponent on the adversarial unit ray)**.** The adaptive program `rosen_sunit_norm_identity_tests.py` also records the maximum of E_j. On the primitive perturbations above it gives

m	$h(x_0)$	$\max_j \log E_j$	$\max_j \log E_j / h(x_0)$
24	33.827753	164.565782	4.86536
48	60.913846	329.708620	5.41306
80	97.028637	546.068700	5.62894

Longer exact norm-only runs give ratios $5.73\ldots$ and $5.80\ldots$ at $m = 120, 160$, while $\max_j \log E_j - 6h(x_0)$ remains approximately -38. Thus the same family which approaches the height exponent two approaches the norm coefficient six. This does not prove (753), but it selects the sharp arithmetic statement to attack: prevent two record norm excursions from compounding, using the exact divisibility (739) and the finite-support S-unit closing theorem.

Experiment II.9.54 (Random exact stress test of the sextic barrier)**.** The program `rosen_norm_barrier_random_tests.py` sampled 300 primitive integral inputs with coefficient bound 10^5, normalized them to the Rosen interval, and followed each exact orbit for up to 5000 steps. It found 226 cusps, 74 special cycles, and no unresolved trajectory. The largest observed value of

$$\max_j \log E_j - 6h(x_0)$$

was $-14.0783753325\ldots$, attained by $p = (-3535, -26446, 22781)$ at step 22 of a 100-step cusp orbit. Thus the broad random sample gives no violation of (753); the unit ray, not random coefficient growth, remains the sharp test.

Experiment II.9.55 (A positive low-height sextic surplus)**.** The all-state program `rosen_outer_incidence_dynamics.py` exhausts reduced algebraic integers in a coefficient box, rather than sampling only large random inputs or selected record states. The searches of radii 10 and 20 found the same largest surplus. It occurs for

$$P = -8 + 6\lambda - \alpha, \qquad P^3 + 20P^2 + 61P - 41 = 0, \qquad h(P) = 1.4283756410\ldots.$$

The exact orbit is eventually two-periodic. At step 7 its dominant remainder has norm 1 893 949, so

$$\log(1\,893\,949) - 6h(P) = 5.8839207790\ldots.$$

Consequently the zero-constant inequality is false: the constant B_7 in (753) cannot be suppressed. The stable maximum from radius 10 to radius 20 is evidence for a bounded low-height correction, not a proof of it.

The integral chart is not the whole projective problem. The same program was therefore extended to primitive pairs (p, q), with primitivity checked exactly by $\gcd(|\,\mathrm{N}\,p|, q) = 1$ when q is rational. Among 10 196 reduced inputs with numerator coefficient bound 4 and denominator bound 20, the larger extremum was

$$x = \frac{-2 + 4\lambda + 3\alpha}{20}, \qquad h(x) = \log 20,$$

whose primitive polynomial is

$$8000X^3 - 1600X^2 - 1620X - 167.$$

At step 38 the dominant remainder has exact norm 35 101 000 831, giving

$$M_{38} - 6h(x) = 6.3071018393\ldots.$$

The exact orbit terminates at step 105. For each of the three numerators producing the largest surpluses, all coprime rational denominators through 200 were tested. In each family the surplus reached an interior maximum and then decreased; there was no growth along the denominator ray. A further 3000 random primitive-pair tests with numerator and denominator bounds 10^5 had largest surplus $-17.4927469549\ldots$.

This extension also caught a computational error before it entered the conjecture. The rational branch of the height routine had written the polynomial of a/b in reverse coefficient order, assigning height zero to $-1/18$. The corrected polynomial is $bX - a$, so the height is $\log\max(|a|, |b|)$. The spurious surplus disappeared. The cubic examples above and all earlier degree-three computations are unchanged.

The same audit also corrects a misleading sign pattern in the earlier record-only tests. For this example at step 7,

$$D_{\mathrm{out}} = -5.7412617337\ldots, \qquad \lambda_0 - 2H_0(w) = -0.1426590453\ldots.$$

Thus neither separate defect has a useful nonnegative sign. Their coupled identity gives exactly the positive norm surplus above. A second exact input $P = -1398 - 108\lambda + 709\alpha$ has $D_{\mathrm{out}} = -0.1747935824\ldots$ at a strict norm record while the central slack is $0.3688549609\ldots$; in that case the central term pays for the negative outer term. These two inputs force the proof to control the sum of all archimedean incidence terms, with the finite places entering through primitivity and the product formula. A prime-by-prime time truncation estimate cannot supply this coupling.

Experiment II.9.56 (The record norms are radical-dominated)**.** The program `rosen_norm_radical_tests.py` factors both norms at the record-energy state for $m = 4, 8, 12, 16$. The larger norm was squarefree in all four cases, with respectively $19, 25, 44, 48$ decimal digits. At $m = 12$, for example, the two exact factorizations are

$$7 \cdot 7699 \cdot 1229603873 \cdot 24707767473626700050988422 3,$$
$$43 \cdot 293 \cdot 14741 \cdot 125383 \cdot 407567 \cdot 24638152522380089663920 33.$$

Both are squarefree. At $m = 16$ the smaller record norm is a square, but the larger 48-digit norm is again squarefree. Thus a generalized *abc* inequality bounding height by radical is compatible with, rather than contradictory to, the sharp norm excursion: the radical already pays essentially the full cost. Conjecture II.9.23 remains useful for separating unit directions from coefficient directions, but it cannot by itself prove (753). A successful radical argument must exploit the *time ordering* and non-repetition of the finite support, not just its total radical.

Experiment II.9.57 (Record support turnover at split places)**.** The PARI-backed program `rosen_sunit_prime_dynamics.py` factors the record norms into actual prime ideals, rather than identifying only their rational norms. For $m = 4, 8, 12, 16$, every one of the 15 prime ideals in the larger record norm appears for the first time at that record. The program checks all local rules in Proposition II.9.15 and all 227 returns of the tracked ideals against (II.9.15); no check fails. For the dominant record, let M be the product of the underlying rational primes and let x_{in} be the state whose next remainder is the record. The exact factorizations and high-precision real evaluation give

m	n_{in}	$\log M$	$\log(M\lvert x_{\mathrm{in},0}\rvert)$
4	562	43.739	36.820
8	-2	55.818	54.789
12	-60	99.861	95.184
16	-2115480726754043	110.319	74.442

Thus these are not marginal applications of Lemma II.9.16: the time-truncated modulus exceeds the archimedean digit scale by e^{36} to e^{95} in the tested cases.

There is also a structural reason that the dominant squarefree factors are split primes. In this cyclic cubic field an unramified rational prime is either inert of residue degree three or splits completely. If ℓ occurs to exponent one in the norm of an algebraic integer, the prime ideal above it has residue degree one, so ℓ must split completely. Hence the radical-dominated sharp excursion is supported on time-truncated degree-one prime ideals. The difficult branch is therefore not the recycling of a fixed S, which the two-step unit equation controls, but repeated creation of new split primes through (3).

Experiment II.9.58 (Most time-truncated mass is disposable, not recycled)**.** A possible way to exploit the return-continuant theorem is to wait until a uniformly positive portion of a large first-hit ideal occurs again. The exact PARI-backed program `rosen_prime_return_delay_tests.py` tests this directly. It factors every remainder into actual prime ideals, distinguishes the three ideals above a split rational prime, records the residue degree, and weights an ideal by $\log \mathrm{N}\,\mathfrak{p}$. The first percentage below is the portion of the logarithmic norm of all post-initial births which ever returns before the orbit closes. The second is the returned portion of the birth mass at the dominant norm

record:

m	$\log \mathrm{N}$(births)	$\log \mathrm{N}$(returned)	returned (%)	$\log \mathrm{N}$(record births)	record returned (%)
1	853.29	61.11	7.16	24.60	21.13
2	1518.61	46.37	3.05	24.97	0
3	2020.17	95.04	4.70	32.20	0
4	2446.24	84.59	3.46	43.74	9.75
5	2479.12	135.83	5.48	40.03	4.86
6	3795.10	94.40	2.49	45.03	12.93
7	4448.15	147.12	3.31	50.50	0
8	3637.16	131.76	3.62	55.82	6.65
9	7858.37	187.41	2.38	62.58	0
10	8385.25	267.13	3.19	65.98	0
11	7407.53	194.12	2.62	67.86	10.61
12	22027.27	364.41	1.65	99.86	9.45

For $m = 12$, only 1.15% of all birth mass returns within 128 steps, and only 1.65% ever returns before the special orbit repeats. At its dominant record the two largest components, of logarithmic norms 56.16 and 12.92, never return. For $m = 9, 10$, no component of the dominant-record birth mass returns.

These examples eventually close, so they do not by themselves decide an infinite-path assertion. They do refute the proposed proof mechanism: return compression sees only a thin recycled residue, whereas the dominant finite-place cost is repeatedly created and discarded without ever becoming a continuant divisor. The main bridge must therefore control a prime at first occurrence (or control an unbounded-delay first-intersection counting function). Merely strengthening the continuant estimate or waiting for a bulk return does not address the sharp branch.

Experiment II.9.59 (The norm-polynomial discriminant does not catch time-truncated records)**.** At transition j, put

$$P_j(N) = \mathrm{N}(r_{j-1} + N\lambda r_j) \in \mathbf{Z}[N].$$

The Rosen recurrence gives $|P_j(n_j)| = |\mathrm{N}(r_{j+1})|$. A possible time-ordered refinement of the norm-form divisibility would be that a large time-truncated record shares prime factors with $\operatorname{Disc} P_j$, making its residue root ramified and hence exceptional. The exact program `rosen_norm_polynomial_discriminant_tes` disproves this mechanism on the adversarial unit ray:

m	record step	$\log \lvert \mathrm{N}\, r_{j+1} \rvert$	$\log \lvert \operatorname{Disc} P_j \rvert$	$\log \gcd(\lvert \mathrm{N}\, r_{j+1} \rvert, \lvert \operatorname{Disc} P_j \rvert)$
1	0	24.599	89.925	0
2	2	31.252	117.987	0
4	11	43.739	149.787	0
8	5	55.818	209.268	0
12	36	99.861	388.523	0
16	4	110.319	306.255	0

Thus every dominant record in the test is coprime to the discriminant. Its time-truncated split primes are simple roots of P_j modulo those primes, not repeated or ramified roots. Neither the discriminant nor a discriminant–value gcd supplies the missing bound; the restriction must use the sequence of Rosen-selected roots rather than one norm polynomial.

Route II.9.60 (The exact remaining S-unit obligation)**.** The S-unit route now separates into two rigorously distinct pieces. Theorem II.9.14 and Lemma II.9.12 completely control every fixed

prime support, including its rational scaling. Proposition II.9.15 controls recycled primes by return continuants. To finish globally by this route it is enough to prove that a Rosen-admissible orbit of a fixed $x \in K$ cannot satisfy the birth congruence (3) for time-truncated split prime ideals of unbounded norm at infinitely many record times. Equivalently, one needs a time-ordered or moving-target estimate for the primitive ideal parts of the remainder sequence. A size-only *abc* estimate does not supply it, because the computed primitive parts are already squarefree. Lemma II.9.16 gives an equivalent sieve formulation on the sharp branch: a modulus comparable to the time-truncated record norm has a residue class whose Rosen-selected representative is archimedeanly tiny. This is the precise input on which a Hodge outer-gap estimate or a deterministic large-sieve/moving-target theorem would have to act. The small lift alone is not contradictory: the elementary degree-three product-formula bound $\log M_{j+1} \le \log|\mathrm{N}(r_{j+1})| \le 3h(r_{j+1}) + O(1)$ is essentially saturated at the sharp records. Any sieve estimate must therefore use the fact that the residue targets arise from one Rosen-admissible Hodge cocycle, not treat them as arbitrary congruences.

This is an exact description of what an S-unit-only proof would owe, not a recommendation to continue prime by prime. Time truncation depends on the chosen time ordering and disappears from the usual projective height. The experiments above show that most time-truncated mass is disposable and that even record-prime charges do not telescope. The more natural strategy is therefore to recover this moving-target restriction as the finite-place shadow of one global Arakelov–Hodge inequality. The fixed-support S-unit theorem should enter only at the final closing step, after that global inequality has bounded the remainder norms.

Route II.9.61 (Logarithmic geometry of the S-unit obstruction)**.** Over a splitting field, the twisted Segre threefold is

$$X = (\mathbf{P}^1)^3.$$

Let H_i be the pullback of a point from its i-th factor, and let

$$D = \sum_{i=0}^{2} \big(\{p_i = 0\} + \{q_i = 0\}\big)$$

be the six-component toric boundary. Then

$$K_X = -2(H_0 + H_1 + H_2), \qquad [D] = 2(H_0 + H_1 + H_2), \qquad K_X + D = 0.$$

The complement $X \setminus D$ is the three-dimensional torus in which the S-unit equations live. A fixed prime support means that the Asai orbit is S-integral relative to D; Theorem II.9.14 proves finiteness after imposing the two successive additive Rosen relations. First appearance of a divisor is a new finite-place intersection with D, and the valuation cube in Lemma II.9.30 records its boundary stratum.

The identity $K_X + D = 0$ explains why an untwisted Vojta or *abc* argument supplies no coercive global height on the ambient Asai threefold. Its truncated counting term can be saturated by time-truncated squarefree boundary intersections, precisely as in the record-prime experiment. Therefore the missing global inequality must add a positive archimedean class. In the manuscript's language that class is the Hodge outer-gap contribution. The arithmetic role of the S-unit theorem is to close the orbit once the time-truncated boundary counting is stopped; the Hodge role is to make infinitely repeated time-truncated boundary intersections incompatible with the archimedean return budget.

Proposition II.9.62 (The exponent two is necessary)**.** *There are a cusp $r \in G_7\infty$, units $u_m \in \mathcal{O}_K^\times$, and times n_m such that, for $x_m = r + u_m$,*

$$h(T^{n_m} x_m) = 2h(x_m) + O(1).$$

Consequently no estimate

$$h(T^n x) \le (2 - \varepsilon)h(x) + O_\varepsilon(1)$$

can hold uniformly.

Proof. Take

$$u_m = \lambda^{-3m}\alpha^{-2m}.$$

At the distinguished embedding $u_{m,0} \to 0$, while at both other real embeddings $|u_{m,1}|, |u_{m,2}| \to \infty$. For the second outer embedding this uses

$$|\lambda_2|^{-3}|\alpha_2|^{-2} > 1.$$

Since u_m is a unit,

$$-\log|u_{m,0}| = \log|u_{m,1}| + \log|u_{m,2}|, \qquad h(u_m) = -\frac{1}{3}\log|u_{m,0}|.$$

Choose, for example, the cusp r represented by the finite Rosen word $(-30, -1, 30, 1)$, and let g be its prefix matrix, oriented so that $g(r) = 0$. All entries needed below are nonzero at the three embeddings. Put $y_m = g(x_m)$. A fixed projective transformation changes height by $O(1)$, so

$$h(y_m) = h(x_m) + O(1) = h(u_m) + O(1).$$

At the distinguished place, the nonvanishing derivative of g at r gives

$$y_{m,0} = c_0 u_{m,0}(1 + o(1)), \qquad c_0 \neq 0.$$

At the two outer places $x_{m,i} \to \infty$, and hence

$$y_{m,i} \longrightarrow g_i(\infty) \in \mathbf{R}^\times.$$

Let N_m be the next Rosen digit. Nearest-λ reduction gives

$$\log|N_m| = -\log|u_{m,0}| + O(1).$$

For the branch

$$B_{N_m} = \begin{pmatrix} -N_m\lambda & -1 \\ 1 & 0 \end{pmatrix},$$

the distinguished local max-norm ratio is $|y_{m,0}|$, because $B_{N_m}(y_{m,0}, 1)^t = y_{m,0}(Ty_{m,0}, 1)^t$ and $Ty_{m,0}$ lies in the compact Rosen interval. At each outer embedding the ratio is $|N_m|\,\Theta(1)$, since $y_{m,i}$ tends to a finite nonzero number. The branch matrix is integral unimodular, so all finite-place height factors cancel. Therefore

$$\begin{aligned} h(B_{N_m} y_m) - h(y_m) &= \frac{1}{3}\left(\log|u_{m,0}| + 2\log|N_m|\right) + O(1) \\ &= -\frac{1}{3}\log|u_{m,0}| + O(1) = h(u_m) + O(1). \end{aligned}$$

Combining this with the preceding height identities proves the claim, with n_m equal to the fixed prefix length plus one. □

Lemma II.9.63 (Exact Rosen branch selection)**.** *Every branch decision in the current programs is certified by rational arithmetic.*

Proof. The distinguished root λ_0 is isolated in $[9/5, 181/100]$. Since

$$\lambda^{-1} = 1 + \lambda - \alpha,$$

the number whose nearest integer is required is an exact quadratic polynomial in λ_0. On a rational isolating interval its extrema occur at the endpoints or at its single rational critical point. Bisect the isolating interval using $X^3 - X^2 - 2X + 1$ until the resulting value interval lies in one nearest-integer cell. The only possible exact half-integer case is detected directly from the rational basis coefficients. This procedure is implemented in `curt10_rosen_tests.py`; it replaces the unsafe fixed-precision selector. A 1000-point regression run with coefficient bound 5000 returned 866 cusp and 134 special certificates, none unresolved. □

Conjecture II.9.64 (Sharp quadratic height excursion)**.** *There is a constant B_7 such that every $x \in \mathbf{P}^1(K)$ and every defined Rosen iterate satisfy*

$$h(T^n x) \le 2h(x) + B_7. \tag{754}$$

Equivalently, for multiplicative height, $H(T^n x) \le e^{B_7} H(x)^2$.

A more local statement would imply this conjecture. It is enough to prove that there is a B such that whenever

$$h(T^{j+1}x) > h(T^j x) + B,$$

one has

$$h(T^j x) \le h(x) + B \quad \text{and} \quad \sup_{m \ge j+1} h(T^m x) \le h(T^{j+1}x) + B. \tag{single-excursion}$$

Indeed, Proposition II.9.7 then bounds the unique large record jump by $2h(x) + O(1)$. Hodge-theoretically, a large jump uses both outer expanding directions at once; the two pairwise decreasing Hodge coordinates are the natural budgets which should prevent a second such jump, except on an exact unit cycle.

Experiment II.9.65 (Backward search from known special cycles)**.** The program `rosen_preperiod_height_search` starts at each of the seven explicit $q = 7$ special cycles, prepends exact inverse Rosen branches, and records the complete known future height profile. A targeted beam search through approximately 44,000 exact preperiod edges, including digits up to 10^5, found no positive value of

$$\max_m h(T^m x) - 2h(x).$$

Random backward walks of depth 80 gave the same result. Since every constructed point is preperiodic by construction, this is a direct test for compounded excursions rather than a finite stopping-time test. It supports the single-excursion formulation, but does not replace its uniform proof.

Experiment II.9.66 (Backward norm search for a second excursion)**.** The companion program `rosen_preperiod_norm_search.py` carries primitive remainder pairs, not just scalar states, while prepending exact inverse branches to the standard special cycle. Thus every tested path has a

certified closing future. A targeted beam through 50,535 exact inverse edges (depth 25, digits through 10^4) found

$$\max\{\log(1+E_{\max})-6h(x)\}=-0.5205615\ldots.$$

The path maximizing a genuine second increase of the running log-energy had first and second record increments only

$$0.5877867\ldots,\qquad 0.5108256\ldots,$$

and sextic surplus $-4.8431\ldots$. This is a more targeted diagnostic than a forward stopping-time search: every future state is known exactly, and the beam separately retains paths with a large second record jump. It found no compounded sharp excursion. The calculation supports the noncompounding target but, because finite-preperiod blindness permits arbitrary local blocks, cannot replace an initial-to-late-state proof.

Proposition II.9.67 (Bounded excursion proves the full classification)**.** *Conjecture II.9.64 implies*

$$\mathbf{P}^1(K)=G_7\infty\ \sqcup \bigcup_{g\ \text{special hyperbolic}} G_7\,\mathrm{Fix}(g).$$

Proof. By Northcott's theorem, only finitely many points of the fixed field K have height at most $2h(x)+B_7$. Hence the Rosen orbit of x either terminates or repeats. Termination is equivalent to membership in the cusp orbit. If $T^a x = T^b x$ with $a<b$, the exact prefix matrices give $g=P_a^{-1}P_b\in G_7$ fixing x. Its eigenvalue on $(x,1)^t$ is $cx+d\in K$, so the nonparabolic fixer is special hyperbolic. Disjointness is the usual parabolic/hyperbolic stabilizer dichotomy. □

This reduces the global theorem to a quantitative height barrier, with no need to construct a moving Asai target. The exponent $2=[K:\mathbf{Q}]-1$ is also the one suggested by the two nonprincipal Hodge summands. The unit ray in Experiment II.9.6 shows that a proof must retain the relative-unit lattice; a pointwise scalar descent argument cannot establish (754).

Route II.9.68 (Certificate that would finish)**.** It is enough to prove (754). A finite-state version would attach to each induced Rosen block a polyhedral region in the two relative-unit coordinates and certify

$$h(T^m x)\le 2h(x)+B_7$$

on that region. Regions on which the ordinary cubic conorm gives descent are already covered by (737); the residual regions must be paired across the unit lattice. Exact equality cycles are checked by the fixed-point form. Unlike the discarded mediant route, this certificate acts on the actual scalar Rosen state and does not require an unconstructed cusp approximant.

II.10 What a standard number-theory conjecture actually gives

There is a clean conditional theorem, but it covers the uniformly supercritical part rather than the parabolic-ray tail. The relevant standard conjecture is Schmidt's symmetric Roth conjecture, strengthened only to allow bounded-degree target fields to vary.

Conjecture II.10.1 (Bounded-degree symmetric Roth)**.** *Fix q and $\varepsilon>0$. Apart from finitely many pairs (L,Y), with $[\mathbf{Q}(L):\mathbf{Q}]\le q$ and $Y\in\mathbf{P}^r(\mathbf{Q})$,*

$$-\log d_\infty(Y,L)<(2+\varepsilon)\max\{h(Y),qh(L)\}+O_{q,r}(1).$$

The fixed-field, one-dimensional version is Schmidt's symmetric Roth conjecture. Allowing the bounded-degree fields to vary is an additional uniformity, so it is stated explicitly rather than silently attributed to the standard conjecture.

Proposition II.10.2 (Conditional uniformly supercritical closing)**.** *Assume bounded-degree symmetric Roth. Suppose an infinite return sequence satisfies* (655)*, the targets have degree at most q, and for some fixed $\eta > 0$,*

$$\chi_n \le (1-\eta)m_n, \qquad qh(L_n) \le (1-\eta)m_n. \tag{755}$$

Then the sequence cannot consist of distinct nonclosing pairs. In the Hodge reduction, repetition or rationality of a pair gives a cusp or a special hyperbolic fixed point.

Proof. Choose $\varepsilon > 0$ so small that $(2+\varepsilon)(1-\eta) < 2$. Equations (655) and (755) give

$$-\log d_\infty(Y_n, L_n) \ge 2m_n - O(1) > (2+\varepsilon)\max\{h(Y_n), qh(L_n)\} + O(1)$$

for all sufficiently large n. Symmetric Roth makes the set of such pairs finite. A repeated return pair is a repeated natural-extension state; its intervening word fixes the endpoint. If the return terminates, the endpoint is a cusp; otherwise the fixed return is hyperbolic, and its K-rational eigenline makes it special. □

The first inequality in (755) is equivalent to a uniform relative surplus $\delta_n/m_n \ge \eta$. It fails on $A_2A_1^N$, where $\delta_N/m_N \to 0$. Therefore standard Schmidt, Roth, *abc*-type irrationality measures, and the usual Vojta conjecture do not conditionally prove the full classification. Adding the assertion that they work with $\varepsilon_n \asymp 1/h(Y_n)$ would be a new boundary approximation conjecture, and in an unrestricted form it would conflict with continued-fraction examples. The mediant-recognition route above is the structurally credible replacement.

II.11 Exact ordinary target height on $A_2A_1^N$

For

$$g_N = A_2A_1^N = \begin{pmatrix} \lambda & N\lambda^2+\alpha \\ 1 & (N+1)\lambda \end{pmatrix},$$

the two fixed slopes over the embedding $\lambda_i = \sigma_i(\lambda)$ are

$$z_{i,\pm} = \frac{-N\lambda_i \pm \sqrt{(N+2)^2\lambda_i^2 - 4}}{2}. \tag{756}$$

The mixed eigenline is the Segre tensor of one selected root in each factor. Since every factor vector is $(z_{i,\pm}, 1)$, a coordinate is identically 1, so there is no finite-place contribution to its projective height. Averaging the eight sign choices gives the exact archimedean formula

$$h(L_N) = \frac{1}{2}\sum_{i=0}^{2}\sum_{\epsilon=\pm}\log\max\{1, |z_{i,\epsilon}|\}. \tag{757}$$

Equivalently, if z_N denotes either root over K, then $h(L_N) = 3h(z_N)$.

Using $\prod_i |\lambda_i| = 1$ and $\sum_i \log\max(1, |\lambda_i|) = \log(\lambda^2 - 1)$, formula (757) yields

$$h(L_N) = \frac{3}{2}\log N + \frac{1}{2}\log(\lambda^2-1) + o(1), \qquad \frac{h(L_N)}{\chi_N} \longrightarrow \frac{3}{2}. \tag{758}$$

The eight Asai eigenvalues have the reciprocal characteristic polynomial

$$\begin{aligned} P_k(X) = X^8 + k^3X^7 + (6k^4 - 10k^2 + 4)X^6 + (5k^5 - 5k^3)X^5 \\ + (k^6 + 13k^4 - 20k^2 + 6)X^4 + (5k^5 - 5k^3)X^3 \\ + (6k^4 - 10k^2 + 4)X^2 + k^3X + 1, \qquad k = N + 2. \end{aligned} \tag{759}$$

PARI/GP reports P_{N+2} irreducible for every $1 \le N \le 1000$. Hence the mixed eigenvalue, and therefore its simple eigenline, has degree 8 in every tested case. This degree test is finite evidence; formula (757) does not depend on an irreducibility assertion.

Experiment II.11.1 (Target-height table)**.** The following values were computed directly from (756)–(757).

N	χ	δ	κ	$h(L_N)$	$h(L_N)/\chi$
3	0.087856	0.382363	10.704348	2.140363	24.3623
10	1.267767	0.370489	2.584475	3.937065	3.1055
100	3.446730	0.368171	2.213635	7.322508	2.1245
1000	5.732023	0.368139	2.128450	10.767448	1.8785
10000	8.032815	0.368138	2.091659	14.220406	1.7703
10^6	12.637787	0.368138	2.058260	21.128060	1.6718

For degree $q = 8$, the ratio $8h(L_N)/(2\delta_N)$ is already 42.51 at $N = 10$ and 229.57 at $N = 10^6$; it diverges. Moreover $8h(L_N)/m_N \to 12$, so the ordinary height also fails the best-case symmetric budget (702).

Route II.11.2 (Verdict)**.** Neither an unconditional optimization of Vojta's theorem nor Schmidt's stronger symmetric conjecture can use the ordinary Weil height of the mixed eigenline. A successful Roth proof must establish a genuinely different adelic relative height and must verify the distance estimate in that same metric. The exact numerical targets are $8h_{\mathrm{rel}}(L) < 2\delta$ for the additive route or at least $8h_{\mathrm{rel}}(L) < m$ for the symmetric route.

II.12 Constructing the alleged rational approximant

The two estimates in (655) must hold for the same rational point. There is a natural exact construction which reveals a tradeoff that is hidden when the estimates are written separately.

Let $v = (p, q)^t \in \mathcal{O}_K^2$ represent the initial endpoint and put $w = gv$. Polarize the rational twisted-Segre map:

$$\Phi(w + tv) = Y^{(0)} + tY^{(1)} + t^2Y^{(2)} + t^3Y^{(3)}. \tag{760}$$

Every coefficient is rational in the eight-dimensional Asai representation. In split coordinates,

$$Y^{(0)} = w_0 \otimes w_1 \otimes w_2, \tag{761}$$

$$Y^{(1)} = v_0 \otimes w_1 \otimes w_2 + w_0 \otimes v_1 \otimes w_2 + w_0 \otimes w_1 \otimes v_2. \tag{762}$$

Lemma II.12.1 (The close Asai approximant is not a scalar point)**.** *Suppose that v_i, w_i are linearly independent for $i = 0, 1, 2$. Then $Y^{(1)}$ is not decomposable. In particular it is not on the rational twisted-Segre variety and does not correspond to a point of $\mathbf{P}^1(K)$, let alone to a cusp in $G_7\infty$.*

Proof. Choose bases with $w_i = e_0$ and $v_i = e_1$. Then

$$Y^{(1)} = e_{100} + e_{010} + e_{001}.$$

Its flattening with respect to the first factor is

$$e_1 \otimes e_{00}e_0 \otimes (e_{10} + e_{01}),$$

which has matrix rank two. Every decomposable tensor has rank-one flattenings, so $Y^{(1)}$ is not decomposable. The twisted-Segre identification of $\mathbf{P}^1(K)$ consists exactly of the rational decomposable tensors. □

This lemma blocks a direct use of Rosen recognition. The only canonical G_7-cusps supplied by a group word g are its two column endpoints, $g\infty$ and $g0$ (up to the convention of using g^{-1}). If g is a Rosen prefix, these are already the principal endpoint and its adjacent endpoint of the corresponding Rosen cylinder. Declaring either one to be a principal or mediant convergent therefore recovers information which was built into the prefix; it does not create the new prefix overlap needed for closing. The decomposable evaluations $\Phi(w+tv)$, $t \in K$, do correspond to scalar K-points, but $w + tv = (g + tI)v$ has no reason to be in the cusp orbit. Thus the spectral constant $e^{-2\delta_7} < L_7^{\mathrm{med}}$ cannot be fed into a Legendre theorem until an additional, genuinely group-theoretic cusp construction is supplied. For a genuine long return, v_0 is exponentially close to the contracting line of $\sigma_0(g)$, while w_1, w_2 are exponentially close to the two expanding lines. Therefore the first summand in (762) points toward the desired mixed line. The other two summands are smaller by e^{-a-b} and e^{-a-c}, respectively. Uniformly in a nondegenerate box this gives

$$-\log d_\infty(Y^{(1)}, L_g) \geq 2m - O(1). \tag{763}$$

But the norm of its dominant summand is

$$\|v_0\| \, \|w_1\| \, \|w_2\| \asymp e^{b+c} \prod_i \|v_i\|,$$

and exact primitive-coordinate computations give

$$h(Y^{(1)}) = b + c + 3h([p : q]) + O(1), \tag{764}$$

not $\chi + O(1)$. Conversely, $Y^{(0)} = \Phi(gv)$ has the central height contribution $-a + b + c = \chi$, but its distinguished factor w_0 is in the terminal compact direction and is not close to the contracting eigenline. It does not satisfy (763).

Experiment II.12.2 (Exact 169-step Rosen orbit)**.** Take

$$x = \frac{-41 + 193\lambda + 94(\lambda^2 - 1)}{97}.$$

Exact Rosen iteration terminates after 169 steps. For each prefix matrix g_s, the coordinates of $Y_s^{(1)}$ were constructed in the rational Asai basis by differentiating $\Phi(g_s v + tv)$, cleared to primitive integers, and compared numerically with the exact eigenlines.

s	$\chi(g_s)$	$2m(g_s)$	$h(Y_s^{(1)})$	$-\log d(Y_s^{(1)}, L_{g_s})$
5	1.808713	8.374003	26.902743	9.661659
10	1.854663	13.030322	32.176047	13.511140
20	4.163850	24.872111	53.593193	25.068300
40	2.377560	39.321876	78.272201	39.556680
60	3.970908	58.334200	111.156429	59.524590

The distance estimate works to a bounded additive error. The asserted height estimate fails by a quantity growing linearly with the return length. The primitive coordinate content was 1 at every displayed prefix, so the discrepancy is not removed by clearing a common factor.

For comparison, the obvious decomposable output $Y^{(0)} = \Phi(g_s v)$ had the smaller central height but its inner product with the distinguished contracting eigenvector ranged from 0.0066 to 0.77 in the sampled prefixes; its projective distance was of order 1, not e^{-2m}.

Experiment II.12.3 (Search through all four polarizations)**.** At prefix $s = 10$, all primitive integer combinations

$$c_0 Y^{(0)} + c_1 Y^{(1)} + c_2 Y^{(2)} + c_3 Y^{(3)}, \qquad |c_i| \le 15,$$

were enumerated, for 855,936 projective coefficient vectors. The four basis points themselves had the following $(h, -\log d)$ pairs:

	$Y^{(0)}$	$Y^{(1)}$	$Y^{(2)}$	$Y^{(3)}$
h	15.3141	32.1760	25.2077	15.4026
$-\log d$	0.0442	13.5111	0.5702	0.2363.

Among combinations with $-\log d \ge 8$, the smallest height found was 32.1760456, attained by a vector projectively equivalent to the tangent scale; no height saving occurred. The combination $(-1, 2, -4, 8)$ had height 28.6077 and proximity 6.3038. It is exactly $-\Phi(w - 2v)$, since these are the coefficients of the cubic at $t = -2$. This illustrates the decomposable tradeoff: adding enough of v_0 to make the distinguished factor contracting introduces first-order errors e^{-b} and e^{-c} in the conjugate factors, losing half of the desired proximity exponent.

Route II.12.4 (Verdict)**.** The polarization calculation identifies where the two halves of (655) come from, but they belong to two different rational vectors. Any proof using Roth must either construct a third rational vector that has both properties or prove a new adelic inequality that turns (764) into the required central height without altering (763). Moreover the close vector is not decomposable, so the mediant-Rosen theorem does not apply to it. The manuscript currently supplies neither the required vector nor a noncircular scalar cusp approximant.

II.13 Route IV: twisting by the Hodge transfer section

The proposed section

$$\Xi = \frac{\theta_1 \theta_2}{\theta_0} \in H^0\big(\Omega^1_{\log} \otimes \mathcal{M}^{-2}\big)$$

has divisor three times the order-seven orbifold point. In particular it has no algebraic zero or pole at the cusp. Multiplication of an auxiliary form by a power of Ξ can change its Hodge norm, but it does not change its vanishing order at L_g or its base locus there.

There are two consistency checks which any such twist must pass.

1. The height of the coefficients, the height of the rational approximant, and the local proximity must all be computed in the same metrized arithmetic vector bundle. An unbounded change of metric changes all three quantities.

2. If the automorphy factor removes a power of the cusp parameter N from the archimedean coefficient norm, the integral model must show where the corresponding denominator or evaluation factor goes. For $g_N = A_2 A_1^N$, the fixed-point form

$$Q_{g_N}(p, q) = p^2 + N\lambda pq - \big(N\lambda^2 + \lambda^2 - 1\big)q^2$$

is primitive because its p^2-coefficient is 1. Its norm form therefore has no common N-power that can be divided out integrally.

On a fixed compact return box, two fixed continuous adelic metrics differ by $O(1)$, and hence cannot change the $\frac{3}{2}\log N$ term in (758). A g-dependent unbounded metric is permitted in Vojta's arithmetic-vector- sheaf formalism, but then (655) must be rederived in that metric. At present the twist is a proposed bookkeeping mechanism, not a verified adelic cancellation.

Route II.13.1 (Verdict)**.** No tested Hodge twist has yet passed the same-metric and integrality checks. In particular, the transfer section by itself does not repair the explicit height of the tangent rational approximant (764).

II.14 Comparison with `curt11.tex`: nearest-wall Hodge charge and the S-unit boundary

The new draft `curt11.tex` and the S-unit route address different halves of one adelic inequality.

1. The S-unit argument is finite-place and dynamical. It studies the actual Rosen remainders, proves that fixed prime support closes, and reduces a nonclosing orbit to infinitely many time-truncated split-prime births. Its natural compactification is the twisted Segre threefold with toric boundary D, for which $K_X + D = 0$.

2. The nearest-wall construction is archimedean and Hodge-theoretic. If $b \le c$, its rank-two pencil
$$\mathcal{E}_1 = \mathcal{L}_0^{-1} \otimes \mathcal{V}_1 \otimes \mathcal{L}_2$$
has mean return weight $-\delta$ and $\det \mathcal{E}_1 = \mathcal{N}_2^{-2}$; the other ordering is symmetric. For the critical degree-d jet module of rank $k = d - 2r + 1$, its dual determinant has weight $kd\delta$.

Thus the approaches are not two descriptions of the same lemma. The S-unit geometry explains why the untwisted finite-place calculation has no positivity, while the nearest-wall determinant proposes the missing positive archimedean class. Their common home is the integral Asai lattice: Lemma II.9.30 describes its local valuation strata and Theorem II.14.1 below describes the two archimedean weights nearest the central line.

Proposition II.14.1 (Nearest-wall calculation from `curt11`)**.** *Suppose $b \le c$. The two return weights of $\mathcal{E}_1$ are*
$$\chi = -a + b + c, \qquad \chi - 2b = -a - b + c,$$
and their mean is $-a + c = -\delta$. Moreover
$$\det \mathcal{E}_1 = (\mathcal{L}_0^{-1}\mathcal{L}_2)^2 = \mathcal{N}_2^{-2}.$$
For a split rank-two pencil $E = L \oplus M$, the module
$$J_{d,r} = X^r Y^r \operatorname{Sym}^{d-2r}\langle X, Y\rangle$$
has rank $k = d - 2r + 1$ and
$$\det J_{d,r} \simeq (\det E^\vee)^{kd/2}.$$
Consequently its average dual weight is $d\delta$, independently of r. The case $c \le b$ interchanges indices $1, 2$.

Proof. The fixed factors $\mathcal{L}_0^{-1}\mathcal{L}_2$ have weight $-a+c$, and the two eigenlines of $\mathcal{V}_1$ have weights $\pm b$. This gives the two weights and the determinant formula. Writing $n = d - 2r$, one has

$$J_{d,r} \simeq (\det E^\vee)^r \otimes \operatorname{Sym}^n E^\vee.$$

Since $\det \operatorname{Sym}^n E^\vee = (\det E^\vee)^{n(n+1)/2}$, the total exponent is

$$rk + \frac{n(n+1)}{2} = (n+1)\left(r + \frac{n}{2}\right) = \frac{kd}{2}.$$

□

Route II.14.2 (Correction to the schematic critical inequality)**.** The schematic inequality displayed in `curt11.tex` contains an extra $+kd\chi$. As written, at $r = d/2$ its right side is $-kd\delta + kd\chi + \mathfrak{s}_g$, which does not have the claimed negative critical slope. The weight

$$-d\chi + 2mj$$

of a dual monomial shows that, after the central term has been absorbed in the nearest-wall determinant, the corrected schematic form is

$$\log|\,\mathrm{N}\,F(Y)| \le -kd\delta + \mathfrak{s}_g(d,r) + 2km\left(\frac{d}{2} - r\right) + o(kd\delta). \tag{765}$$

Indeed the last explicit term is the cost of moving the vanishing index from $d/2$ to r. If

$$\frac{1}{2} - \frac{r}{d} = o(\delta/m), \qquad \mathfrak{s}_g(d,r) = o(kd\delta),$$

then the right side is negative. This is algebraically consistent with the claimed integrality argument; the version containing $+kd\chi$ is not.

Route II.14.3 (Finite-place saturation: what is proved and what must be checked)**.** The primitive-polynomial lemma in `curt11.tex` is correct. If R is a DVR and $q \in R[X,Y]$ is primitive, then

$$q^r R[X,Y]_{d-2r} \subset R[X,Y]_d$$

is saturated: reducing $\pi F = q^r G$ modulo the uniformizer first forces $\overline{G} = 0$, after which division by π gives $F = q^r G_1$.

This lemma alone does not yet prove that the finite-place defect of the *tensor-induced* jet lattice is zero. One must still verify:

1. tensor induction and descent from the three conjugate factors preserve the claimed lattice rather than a finite-index sublattice;
2. the diagonal Veronese and the Galois-invariant projection are saturated, especially at the ramified prime above 7;
3. the normalization by the Higgs-ratio section Ψ_i has an arithmetic integral model and introduces no unrecorded content ideal.

The integral matrix $\mathcal{A}(n) \in \mathrm{SL}_8(\mathbf{Z})$, its invariant content, and the valuation cube (II.9.30) provide a direct way to perform this audit without changing bases through the discriminant-49 trace lattice. Smith normal forms of the actual Asai jet maps, not only of multiplication by q^r, are the required certificate.

Proposition II.14.4 (The abstract pointwise defect bound is determinant-only)**.** *Let E be a rank-two real or complex Hermitian space with a split boundary, and equip*

$$J_{d,r} = X^r Y^r \operatorname{Sym}^{d-2r} E^\vee, \qquad k = d - 2r + 1,$$

with the induced Hilbert quotient norm. If two metrics on E have comparison map A, then their determinant norms on $J_{d,r}$ differ by

$$|\det A|^{-kd/2}.$$

Consequently, after their determinant metrics are identified, the determinant defect on $J_{d,r}$ is exactly zero. Replacing induced Hilbert norms by the usual coefficient/sup norms changes the logarithmic determinant by at most

$$Ck \log(d+2) + Ck$$

with an absolute C.

Proof. The determinant character is Proposition II.14.1:

$$\det J_{d,r} = (\det E^\vee)^{kd/2}.$$

Every change of metric therefore acts on this one-dimensional line by the stated power of $|\det A|$; the individual singular values and the condition number cancel. For binary degree-d forms, the Hilbert (Bombieri) and sup/coefficient norms differ by at most a fixed power of $d+1$. On a k-dimensional determinant their logarithmic ratio is therefore $O(k \log(d+2))$. □

Corollary II.14.5 (Exact construction target for the realization defect)**.** *The pointwise estimate proposed in* `curt11.tex` *follows once one constructs an isomorphism between its tensor-induced integral jet lattice and the saturated module*

$$Q_g^r \operatorname{Sym}^{d-2r}(\mathcal{O}_K^2)^\vee$$

such that the determinant metric corresponds to the $\Psi_i^{kd/2}$-normalized nearest-wall determinant. After this identification, Proposition II.14.4 gives the required $O(k \log d)$ error.

Lemma II.14.6 (The primitive quadratic jet has no parabolic base component)**.** *Let q be a nonzero homogeneous quadratic form over a field, let $r > 0$, and let $d \geq 2r$. The common zero locus on $\mathbf{P}^1$ of*

$$q^r \operatorname{Sym}^{d-2r}\langle X, Y\rangle$$

is exactly $V(q)$. In particular, for

$$Q_N(X,Y) = X^2 + N\lambda XY - (N\lambda^2 + \lambda^2 - 1)Y^2$$

one has $Q_N(\lambda, 1) = 1$, so $[\lambda : 1]$ is not in this jet base locus.

Proof. Every member vanishes on $V(q)$. If $q(v) \neq 0$, choose a degree $d-2r$ monomial (or linear-form power) nonzero at v; its product with q^r is nonzero at v. The final assertion follows by direct substitution. □

Route II.14.7 (The parabolic equality term is not supplied by saturation)**.** Lemma II.14.6 shows that the primitive quadratic saturation lemma and the critical determinant calculation do not produce the parabolic alternative claimed in `curt11.tex`. On the extremal cusp family, one of two things must happen:

1. the actual tensor-induced realization defect has an additional boundary contribution of the same critical size, permitting the nonzero unit evaluation $Q_N(\lambda, 1) = 1$; or
2. the auxiliary system differs from $Q_g^r \operatorname{Sym}^{d-2r}$ and has extra, explicitly constructed parabolic-incidence equations in its common base locus.

Until the tensor-induced evaluation map is written down, these possibilities cannot be distinguished. In particular, finite-place saturation alone does not justify a parabolic common-zero component.

Proposition II.14.8 (The extremal cusp forces a full-size boundary defect)**.** *Suppose the corrected critical evaluation inequality is applied to the rank-one critical module with $d = 2r$, generated by Q_N^r, at $v = (\lambda, 1)$. Then its realization defect satisfies*

$$\mathfrak{s}_{g_N}(d, r) \geq d\delta_N - o(d\delta_N).$$

In particular it is not $o(d\delta_N)$.

Proof. Here $k = 1$ and $Q_N(v)^r = 1$, so the left side of the norm-evaluation inequality is zero. The index-error term vanishes because $r = d/2$. Thus

$$0 \leq -d\delta_N + \mathfrak{s}_{g_N}(d, r) + o(d\delta_N),$$

which is the assertion. □

This proposition makes the equality mechanism precise. If the integral jet is the primitive Q_N^r-module, the parabolic family is not an extra conclusion under the small-defect hypothesis; it is an equality-sized failure of that hypothesis. A single optimal theorem can still cover both branches, but it must state the dichotomy at the level of the defect:

$$\mathfrak{s}_g = o(kd\delta) \quad \text{or} \quad \mathfrak{s}_g = kd\delta + o(kd\delta) \text{ on an explicit parabolic incidence.}$$

Conjecture II.14.9 (Corrected nearest-wall defect alternative)**.** *For every recurrent sequence with $\chi_g \to \infty$, choose critical indices $r/d = 1/2 - o(\delta/m)$. After passage to an infinite subsequence, exactly one of the following occurs:*

1. *$\mathfrak{s}_g(d, r) = o(kd\delta)$, and the corrected tensor-induced GIT inequality forces the fixed-eigenline equations;*
2. *the returns belong to one fixed finite positive multi-parabolic connector type,*

$$\mathfrak{s}_g(d, r) = kd\delta + o(kd\delta),$$

and the corresponding fixed-point remainders have bounded norm.

Corollary II.14.10. *Conjecture II.14.9 implies the full classification.*

Proof. The first branch gives a special hyperbolic fixer. The second gives a cusp by Theorem II.14.11. □

Theorem II.14.11 (All finite multi-parabolic bounded-norm families)**.** *Fix $s \geq 1$ and $U_0, \dots, U_s \in G_7$. Put*

$$g_{\mathbf{n}} = U_0 T^{n_1} U_1 T^{n_2} \cdots T^{n_s} U_s, \qquad \mathbf{n} \in \mathbf{Z}^s_{>0}.$$

Let $v \in K^2 \setminus \{0\}$. If $|\operatorname{N} Q_{g_{\mathbf{n}}}(v)|$ is bounded along an infinite sequence of tuples $\mathbf{n}$ escaping every finite set, then $[v] \in G_7 \infty$.

Proof. Pass to a subsequence on which each n_i is either constant or tends to infinity. Absorb every constant power, together with its neighboring fixed connectors, into new fixed matrices. We obtain

$$g_{\mathbf{m}} = W_0 T^{m_1} W_1 \cdots T^{m_t} W_t, \qquad m_i \longrightarrow \infty.$$

Write $E = e_1 e_2^t$, so $T^m = I + m\lambda E$. If an internal connector satisfies

$$e_2^t W_i e_1 = 0 \qquad (1 \le i < t),$$

then W_i fixes ∞. The cusp stabilizer in G_7 is generated projectively by T, so $W_i = \pm T^{\ell}$; it can be absorbed by merging the two adjacent positive parabolic blocks. Repeat until every internal scalar $e_2^t W_i e_1$ is nonzero.

The coefficient matrix of the unique full monomial $m_1 \cdots m_t$ in $g_{\mathbf{m}}$ is

$$\lambda^t W_0 E W_1 E \cdots E W_t = \lambda^t \left(\prod_{i=1}^{t-1} e_2^t W_i e_1 \right) W_0 e_1 e_2^t W_t.$$

Since $Q_g(v) = \det(v, gv)$ is linear in the entries of g, the coefficient of $m_1 \cdots m_t$ in $Q_{g_{\mathbf{m}}}(v)$ is

$$C(v) = \lambda^t \left(\prod_{i=1}^{t-1} e_2^t W_i e_1 \right) (e_2^t W_t v) \det(v, W_0 e_1). \tag{766}$$

Every lower monomial omits at least one m_i. Therefore

$$\frac{Q_{g_{\mathbf{m}}}(v)}{m_1 \cdots m_t} \longrightarrow C(v)$$

at all three embeddings. If $C(v) \neq 0$, then

$$\frac{|\,\mathrm{N}\, Q_{g_{\mathbf{m}}}(v)|}{(m_1 \cdots m_t)^3} \longrightarrow |\,\mathrm{N}\, C(v)| > 0,$$

contradicting boundedness. All internal factors in (766) are nonzero, so either $e_2^t W_t v = 0$, which gives $[v] = W_t^{-1}\infty$, or $\det(v, W_0 e_1) = 0$, which gives $[v] = W_0 \infty$. Both are cusps. □

Theorem II.14.12 (A moving dominant parabolic block forces a cusp)**.** *Fix a primitive $v \in \mathcal{O}_K^2$. Let $U_\nu, V_\nu \in G_7$, $N_\nu \to +\infty$, and put*

$$g_\nu = U_\nu T^{N_\nu} V_\nu, \qquad Q_{g_\nu}(v) = B_\nu + N_\nu A_\nu,$$

where

$$A_\nu = \lambda (e_2^t V_\nu v) \det(v, U_\nu e_1), \qquad B_\nu = Q_{U_\nu V_\nu}(v).$$

Suppose

$$h([A_\nu : B_\nu]) = o(\log N_\nu) \tag{767}$$

and, for some fixed $\varepsilon > 0$,

$$\log |\,\mathrm{N}\, Q_{g_\nu}(v)| \le (3 - \varepsilon) \log N_\nu \tag{768}$$

along an infinite subsequence. Then $A_\nu = 0$ for some ν, and therefore

$$[v] = U_\nu \infty \quad \text{or} \quad [v] = V_\nu^{-1} \infty.$$

In particular, $[v]$ is a cusp.

Proof. Suppose $A_\nu \neq 0$ throughout an infinite subsequence and set $z_\nu = B_\nu/A_\nu$. Since K has fixed degree three, (767) gives, at every real embedding,

$$\log^+ |\sigma_i(z_\nu)| \le 3h(z_\nu) = o(\log N_\nu).$$

Thus $\sigma_i(z_\nu) = N_\nu^{o(1)}$, uniformly for $i = 0, 1, 2$, and

$$|N_\nu + \sigma_i(z_\nu)| = N_\nu^{1+o(1)}.$$

Both A_ν and B_ν are algebraic integers. Hence $|\,\mathrm{N}\,A_\nu| \ge 1$, and

$$\log |\,\mathrm{N}(B_\nu + N_\nu A_\nu)| = \log |\,\mathrm{N}\,A_\nu| + \sum_{i=0}^{2} \log |N_\nu + \sigma_i(z_\nu)| \ge (3 - o(1)) \log N_\nu.$$

This contradicts (768). Therefore $A_\nu = 0$ for at least one index. Formula (II.14.12) then gives $e_2^t V_\nu v = 0$ or $\det(v, U_\nu e_1) = 0$, which are exactly the two displayed cusp incidences. □

Corollary II.14.13 (Hodge closing for one overwhelmingly dominant digit)**.** *In the setting of Theorem II.14.12, suppose the return estimate gives*

$$\log |\,\mathrm{N}\,Q_{g_\nu}(v)| \le \chi_\nu + O(1), \qquad \chi_\nu = (1 + o(1)) \log N_\nu,$$

and the connector height satisfies (767). Then $[v]$ is a cusp.

Proof. Take any fixed $\varepsilon < 2$. The Hodge estimate implies (768) for all sufficiently large ν, and apply the theorem. □

Theorem II.14.14 (Moving finite multi-block subcubic closing)**.** *Fix $s \ge 1$ and a primitive $v \in \mathcal{O}_K^2$. Let*

$$g_\nu = U_{0,\nu} T^{N_{1,\nu}} U_{1,\nu} \cdots T^{N_{s,\nu}} U_{s,\nu}, \qquad \min_i N_{i,\nu} \longrightarrow \infty,$$

with $U_{i,\nu} \in G_7$. Expand

$$Q_{g_\nu}(v) = \sum_{I \subset \{1,\dots,s\}} c_{I,\nu} \prod_{i \in I} N_{i,\nu},$$

and let $C_\nu = c_{\{1,\dots,s\},\nu}$ be the full-monomial coefficient. Assume the expression is reduced, in the sense that $e_2^t U_{i,\nu} e_1 \neq 0$ for $1 \le i < s$; parabolic internal connectors can always be merged to achieve this. Suppose, on every subsequence on which $C_\nu \neq 0$,

$$\max_I h([C_\nu : c_{I,\nu}]) = o\left(\min_i \log N_{i,\nu} \right). \tag{769}$$

If for some fixed $\varepsilon > 0$,

$$\log |\,\mathrm{N}\,Q_{g_\nu}(v)| \le (3 - \varepsilon) \sum_{i=1}^{s} \log N_{i,\nu}, \tag{770}$$

along an infinite subsequence, then $[v]$ is a cusp.

Proof. Assume $C_\nu \neq 0$. Divide the expansion by $C_\nu \prod_i N_{i,\nu}$. Fixed degree and (769) imply, at each embedding,

$$\left| \sigma_k(c_{I,\nu}/C_\nu) \prod_{i \notin I} N_{i,\nu}^{-1} \right| \le \left(\min_i N_{i,\nu} \right)^{-1+o(1)}$$

for every proper I. There are only $2^s - 1$ such terms, hence

$$\sigma_k(Q_{g_\nu}(v)) = \sigma_k(C_\nu) \prod_i N_{i,\nu}\,(1 + o(1))$$

simultaneously for $k = 0, 1, 2$. Since C_ν is a nonzero algebraic integer,

$$\log|\,\mathrm{N}\,Q_{g_\nu}(v)| \geq (3 - o(1)) \sum_i \log N_{i,\nu},$$

contrary to (770). Thus $C_\nu = 0$ at some index. Formula (766) factors C_ν as nonzero connector scalars times

$$(e_2^t U_{s,\nu} v)\det(v, U_{0,\nu} e_1).$$

Its vanishing makes $[v] = U_{s,\nu}^{-1}\infty$ or $[v] = U_{0,\nu}\infty$, so $[v]$ is a cusp. □

Route II.14.15 (What remains after dominant-block removal)**.** Theorem II.14.12 is stronger than the fixed connector theorem on its natural branch: both connectors may vary, and the norm need only be subcubic rather than bounded. It handles every near-critical return in which one parabolic logarithm dominates the height of everything on its two sides. The unresolved parabolic tail is therefore a balanced many-block regime: every large digit has connector height comparable to its own logarithm. Theorem II.14.14 also removes every fixed number of long blocks whose coefficient heights are small relative to the shortest block. Equivalently, in the surviving regime thick connectors are created at a rate which can be charged by the segmentwise Hodge inequality (785). This is the branch on which the time-ordered S-unit equation, rather than a stronger single-target Roth theorem, is genuinely needed.

Experiment II.14.16 (The adversarial ray lies on the balanced-block boundary)**.** The program `parabolic_dominant_block_tests.py` applies the preceding decomposition to the largest digit in each exact adversarial orbit. For $g = T^{-n_j} S P_j$, it computes the exact relative coefficient height $h(B/A)$ and the norm-growth exponent $\log|\,\mathrm{N}\,Q_g(v)|/\log|n_j|$:

m	$\log\lvert n_j\rvert$	$h(B/A)$	$h(B/A)/\log\lvert n_j\rvert$	$\log\lvert\,\mathrm{N}\,Q_g(v)\rvert/\log\lvert n_j\rvert$
4	6.332	24.174	3.818	11.046
8	8.202	25.447	3.103	11.094
12	21.745	34.472	1.585	5.375
16	35.288	43.500	1.233	4.024
24	62.374	61.557	0.987	3.059

Thus the sharp family does not satisfy $h(B/A) = o(\log|n_j|)$ at the tested scale; instead its coefficient height becomes comparable to the parabolic logarithm, while the norm exponent approaches the cubic boundary. This is exactly the balanced-block regime left by Theorem II.14.12, not a counterexample to it.

Proposition II.14.17 (Exact balanced asymptotics of the adversarial unit ray)**.** *Let*

$$u_m = \lambda^{-3m}\alpha^{-2m}, \qquad \beta = -\log|\sigma_0(u_1)| > 0,$$

and use the cusp perturbation and the fixed four-digit prefix $(-30, -1, 30, 1)$ *from the experiments above. Let* n_m *be the next Rosen digit, write*

$$g_m = T^{-n_m} S P, \qquad Q_{g_m}(v_m) = B_m - n_m A_m,$$

and take A_m, B_m with the sign convention of (II.14.12)*. Then*

$$\frac{\log|n_m|}{m} \longrightarrow \beta, \qquad \frac{h([A_m : B_m])}{m} \longrightarrow \frac{2\beta}{3}, \qquad \frac{\log|\mathrm{N}\, Q_{g_m}(v_m)|}{m} \longrightarrow 2\beta. \tag{771}$$

In particular,

$$\frac{h([A_m : B_m])}{\log|n_m|} \longrightarrow \frac{2}{3}, \qquad \frac{\log|\mathrm{N}\, Q_{g_m}(v_m)|}{\log|n_m|} \longrightarrow 2.$$

Proof. Write the perturbed homogeneous vector as $v_m = v_c + u_m w$, where v_c is the fixed cusp vector. The fixed prefix satisfies

$$Pv_c = (0, a)^t, \qquad Pw = (b, c)^t$$

with $a, b \neq 0$ at all three embeddings; this is checked directly from the displayed four matrices. Thus the state before the large digit is

$$x_m = \frac{u_m b}{a + u_m c}.$$

At the distinguished embedding $u_{m,0} \to 0$, and nearest-λ reduction gives

$$n_m = -\frac{a_0}{\lambda_0 b_0 u_{m,0}} + O(1).$$

This proves the first limit.

The coefficient A_m is a fixed nonzero factor times u_m, while

$$B_m = d_0 + d_1 u_m + d_2 u_m^2$$

with $d_0 d_2 \neq 0$ at every embedding. Hence

$$\log|\sigma_0(B_m/A_m)| = \beta m + O(1),$$

whereas, for $i = 1, 2$,

$$\log|\sigma_i(B_m/A_m)| = \log|\sigma_i(u_m)| + O(1).$$

The ideals of A_m range in one fixed ideal class because u_m is a unit, so finite-place cancellation contributes $O(1)$. The product formula gives

$$\log|\sigma_1(u_m)| + \log|\sigma_2(u_m)| = \beta m,$$

and the height limit is $2\beta/3$.

It remains to justify that nearest-integer cancellation cannot make the distinguished conjugate exponentially smaller than $u_{m,0}$. Put

$$\theta = u_1^{-1}, \qquad C = -\frac{a}{\lambda b}, \qquad D = -\frac{c}{\lambda b}, \qquad \delta_* = \frac{p_c}{\lambda q},$$

where $v_c = (p_c, q)^t$. Then n_m is the nearest integer to $\sigma_0(C\theta^m + D)$, and, if $\delta_m = \sigma_0(C\theta^m + D) - n_m$, direct substitution in the determinant gives

$$\sigma_0(Q_{g_m}(v_m)) = \sigma_0(u_m b)\big(\sigma_0(p_c + u_m q) - \sigma_0(\lambda q)\delta_m\big).$$

Now θ is a Pisot unit: its three conjugates have absolute values

$$29.540\ldots, \qquad 0.05668\ldots, \qquad 0.59717\ldots.$$

Choose D_0 with $D_0C, D_0D \in \mathcal{O}_K$. The residues of $\mathrm{Tr}(C\theta^m + D)$ modulo $\mathbf{Z}$ therefore lie in the finite group $D_0^{-1}\mathbf{Z}/\mathbf{Z}$, while the two small conjugates of $C\theta^m$ tend exponentially to zero. Consequently the centered fractional parts δ_m approach a finite periodic set. None of its limits is a boundary value $\pm 1/2$: such an equality would say that D differs from a rational number by zero at one real embedding, hence in K, whereas the λ- and α- coordinates of D are nonzero. Likewise no limit equals $\sigma_0(\delta_*)$, because this would force $D - \delta_* \in \mathbf{Q}$, whereas exactly

$$(D - \delta_*)_\lambda = \frac{700811100}{607570951}, \qquad (D - \delta_*)_\alpha = \frac{-1382829330}{607570951}.$$

Thus the last factor in (II.14) is bounded above and below away from zero for all sufficiently large m, and

$$\log|\sigma_0(Q_{g_m}(v_m))| = -\beta m + O(1).$$

At each outer embedding the term $n_m A_m$ dominates: its logarithm is $\beta m + \log|\sigma_i(u_m)| + O(1)$. Summing the three logarithms gives

$$-\beta m + \big(\beta m + \log|\sigma_1(u_m)|\big) + \big(\beta m + \log|\sigma_2(u_m)|\big) = 2\beta m + O(1),$$

which proves the final limit. □

Proposition II.14.18 (Pisot-cone rounding is uniformly noncancelling)**.** *Let $K/\mathbf{Q}$ be totally real cubic, with embeddings $\sigma_0, \sigma_1, \sigma_2$. Fix $C, D, E \in K$, with $C \neq 0$ and*

$$D - E \notin \mathbf{Q}.$$

Let $\theta_\nu \in \mathcal{O}_K^\times$ satisfy

$$|\sigma_0(\theta_\nu)| \longrightarrow \infty, \qquad |\sigma_1(\theta_\nu)|, |\sigma_2(\theta_\nu)| \longrightarrow 0,$$

and let n_ν be a nearest integer to $\sigma_0(C\theta_\nu + D)$. Then

$$\liminf_{\nu\to\infty} |\sigma_0(C\theta_\nu + D) - n_\nu - \sigma_0(E)| > 0. \tag{772}$$

No rank-one or power-sequence hypothesis on the units is required.

Proof. Choose an integer $M > 0$ such that $MC, MD \in \mathcal{O}_K$. Since θ_ν is an algebraic unit,

$$\mathrm{Tr}(C\theta_\nu + D) \pmod{\mathbf{Z}}$$

takes values in the finite group $M^{-1}\mathbf{Z}/\mathbf{Z}$. On a subsequence where this residue is fixed, the identity

$$\sigma_0(C\theta_\nu + D) = \mathrm{Tr}(C\theta_\nu + D) - \sigma_1(C\theta_\nu + D) - \sigma_2(C\theta_\nu + D)$$

shows, modulo $\mathbf{Z}$, that the left side approaches $\sigma_0(D) + r$ for one $r \in \mathbf{Q}$. There are only finitely many such limiting residues.

If one of their centered representatives were $\sigma_0(E)$, then for some $m \in \mathbf{Z}$,

$$\sigma_0(D - E + r - m) = 0.$$

An embedding of a field is injective, so $D - E = m - r \in \mathbf{Q}$, contrary to the hypothesis. The finite limiting set is therefore a positive distance from $\sigma_0(E)$. This proves (772), including the harmless possibility that a limiting residue lies on a nearest-integer boundary. □

Corollary II.14.19 (Every nonexceptional cusp prefix has a full two-dimensional sharp cone)**.** *Let*

$$P = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \in G_7, \qquad A \neq 0,$$

be a finite admissible Rosen prefix, and perturb its cusp endpoint by

$$x_u = -\frac{B}{A} + u, \qquad u \in \mathcal{O}_K^{\times}.$$

Suppose

$$\frac{B - C}{\lambda A} \notin \mathbf{Q}. \tag{773}$$

For every unit sequence with

$$|\sigma_0(u)| \to 0, \qquad |\sigma_1(u)|, |\sigma_2(u)| \to \infty,$$

the rounding in the first branch after P cannot create an additional exponential cancellation in the fixed-point determinant.

Proof. Represent x_u by

$$v_u = (-B + uA, A)^t = (-B, A)^t + u(A, 0)^t.$$

The determinant-one identity gives

$$P(-B, A)^t = (0, 1)^t, \qquad P(A, 0)^t = (A^2, AC)^t.$$

Thus, immediately before the next digit,

$$Pv_u = (uA^2, 1 + uAC)^t$$

and that digit is the nearest integer to

$$\sigma_0\left(-\frac{u^{-1}}{\lambda A^2} - \frac{C}{\lambda A}\right).$$

Writing its centered error as δ_u, direct evaluation of the next-branch determinant gives a fixed nonzero factor times

$$\sigma_0(u)\left(-\sigma_0(B) + \sigma_0(uA) - \sigma_0(\lambda A)\delta_u\right).$$

The sole limiting cancellation value is

$$E = -\frac{B}{\lambda A},$$

whereas the constant term in the rounding target is $-C/(\lambda A)$. Their difference is exactly $(B - C)/(\lambda A)$. Apply Proposition II.14.18 with $\theta = u^{-1}$. □

Proposition II.14.20 (Inhomogeneous S-unit rounding in every direction)**.** *Let $K/\mathbf{Q}$ be cyclic totally real cubic, fix its distinguished embedding, and let $C, F \in K^{\times}$ with $F \notin \mathbf{Q}$. For every $\varepsilon > 0$, there are only finitely many $(\theta, n) \in \mathcal{O}_K^{\times} \times \mathbf{Z}$, with $|\sigma_0(\theta)|$ unbounded, for which*

$$0 < |\sigma_0(C\theta + F) - n| < H(\theta)^{-\varepsilon}. \tag{774}$$

Proof. Apply the theorem on algebraic approximation to linear combinations of S-units of Nair–Kumar–Rout [35, Theorem 1.1] to the two S-units $(\theta, 1)$, with coefficients (C, F) and $q = 1$. If (774) had infinitely many solutions, then, after deleting finitely many solutions, the tuple $(C\theta, F)$ would be pseudo-Pisot.

We show that this is impossible. Write

$$\theta_i = \sigma_i(\theta), \quad C_i = \sigma_i(C), \quad F_i = \sigma_i(F).$$

Coincidences between two conjugates of $C\theta$, or between a conjugate of $C\theta$ and one of F, occur only finitely often. Indeed, their quotient gives a fixed value of $\sigma_i(\theta)/\sigma_j(\theta)$; the kernel on $\mathcal{O}_K^\times$ consists only of the rational units, because K has prime degree. We may therefore treat all six conjugates as distinct.

The pseudo-Pisot condition says in particular that $|C_1\theta_1|, |C_2\theta_2| < 1$. Since $|\theta_0| \to \infty$ and $|\theta_0\theta_1\theta_2| = 1$, pass to a subsequence on which $\theta_2 \to 0$. The sum of the six conjugates is an integer. Comparing that integer with the nearest integer n, and passing once more to a subsequence, gives a fixed $\ell \in \mathbf{Z}$ such that

$$C_1\theta_1 + C_2\theta_2 + F_1 + F_2 = \ell + O(H(\theta)^{-\varepsilon}). \tag{775}$$

If also $\theta_1 \to 0$, then $F_1 + F_2 = \ell$. Since $\mathrm{Tr}(F) - F_0 = \ell$, the number F_0, and hence the field element F, is rational, a contradiction.

Otherwise θ_1 stays bounded above and below. Equation (775) makes it approach the fixed nonzero algebraic number

$$L = \frac{\ell - F_1 - F_2}{C_1}$$

with error

$$O(|\theta_2| + H(\theta)^{-\varepsilon}) = \exp(-c\,h(\theta))$$

for some $c > 0$, after decreasing c if necessary. Express θ in a fixed basis of the rank-two unit group. Its exponent norm is comparable to $h(\theta)$. The Baker–Wüstholz lower bound for a nonzero linear form in logarithms [55] gives instead

$$|\theta_1 - L| \geq (2 + \|\exp(\theta)\|)^{-C_{K,C,F}},$$

unless $\theta_1 = L$; the latter equality determines at most one field element θ. The polynomial lower bound contradicts the exponential upper bound. Thus the pseudo-Pisot alternative cannot occur, which proves the proposition. □

Corollary II.14.21 (No exponential central cancellation for any unit direction)**.** *In Corollary II.14.19, remove both outer conditions on u, retaining only $|\sigma_0(u)| \to 0$. For every $\varepsilon > 0$, the centered rounding factor in the fixed-point determinant is either zero at one of finitely many indices or has size at least*

$$H(u)^{-\varepsilon}$$

for all sufficiently large indices.

Proof. Use Proposition II.14.20 with

$$\theta = u^{-1}, \qquad C = -\frac{1}{\lambda A^2}, \qquad F = \frac{B - C}{\lambda A}.$$

Hypothesis (773) says exactly that $F \notin \mathbf{Q}$. □

Route II.14.22 (The strongest available pseudo-Pisot theorem and its exact scope)**.** Corvaja–Zannier prove that exponentially small fractional parts of α^n force a power of α to be Pisot. The Kulkarni–Mavraki–Nguyen theorem treats sums $\sum q_i \alpha_i^n$ with $h(q_i) = o(n)$ and concludes that every infinite exponentially accurate family is a pseudo-Pisot trace family [32, 33]. Kumar–Thangadurai give a simultaneous version in which the multiplicative variable ranges through a fixed finitely generated group [34]. The 2025 Nair–Kumar–Rout theorem finally treats linear combinations of arbitrary elements of such a group [35]; Proposition II.14.20 is its inhomogeneous cubic consequence.

The application uses the family of tuples $(\theta, 1)$. The noncollision hypothesis in their Theorem 1.1 is satisfied after deleting finitely many repetitions, because the quotient of the two coordinates is exactly θ; the ambient unit group is enlarged to a Galois-stable finitely generated group. The denominator parameter in their theorem is $q = 1$, so its factor $|q|^{-md-\varepsilon}$ introduces no hidden loss here. This audit is important: the theorem would not permit arbitrary moving algebraic coefficients to be absorbed into the u_i's unless those coefficients belonged to one fixed finitely generated multiplicative group.

These are Subspace-Theorem results, stronger and more relevant here than a C/q^2 Roth estimate: they identify the exceptional trace mechanism, not merely a convergent. Proposition II.14.18 settles the interior two-small-conjugate cone even more elementarily, while Corollary II.14.21 covers its boundary. All of these results still require fixed coefficients, or coefficient height sublinear in the main unit height. Hence they do not by themselves control the remaining balanced tail in which connector and unit heights are comparable.

Route II.14.23 (Moving gcd theorems stop at the same boundary)**.** Grieve–Wang's moving-target gcd theorem [36, Theorem 1.2] allows the polynomials themselves to move, but assumes

$$\max\{h(f_\nu), h(g_\nu)\} = o\Big(\max_i h(u_{i,\nu})\Big)$$

for a fixed set of S-unit variables. Its alternative is a small-gcd subsequence or containment in translates of proper algebraic subgroups. This is exactly strong enough for the dominant-block regime already covered by Theorems II.14.12 and II.14.14. At the computed record times the time-truncated squarefree coefficient part has height comparable to the unit height and has support which grows with the return. Thus neither this moving gcd theorem nor the Nair–Kumar–Rout theorem controls those records without a new Hodge payment. The agreement of their hypotheses with the experimentally observed boundary is useful: it locates the missing uniformity in the moving coefficient support, not in the one-place rational approximation exponent.

Experiment II.14.24 (Exact Pisot rounding audit)**.** The program `rosen_pisot_rounding_tests.py` multiplies the four prefix matrices exactly and obtains

$$Pv_c = (0, 1)^t,$$

as well as the exact vectors Pw, C, D, and δ_* used above. It verifies that $\theta = 3 + 6\lambda + 7\alpha$ has conjugate absolute values $29.54048\ldots, 0.056687\ldots, 0.597171\ldots$, and checks the two nonrational coordinates of $D - \delta_*$ displayed in the proof. For the sample $m = 8, 16, 24, 32, 48, 64$, the exact next digits range from -3648 to an 86-digit integer; their high-precision rounding errors remain separated from $\sigma_0(\delta_*) = 0.0103713\ldots$, in agreement with the finite-residue proof.

Experiment II.14.25 (Full cone and rational-prefix exception audit)**.** The exact program `rosen_pisot_cone_tests.py` tests all 1264 unit directions in the box of fundamental-unit exponents $[-36, 36]^2$ which are small at σ_0 and large at both outer embeddings. For the prefix $(-30, -1, 30, 1)$,

the smallest observed distance from the unique cancellation residue is $0.0052557399\ldots$; the proof of Proposition II.14.18, rather than this decimal, supplies the uniform positive lower bound.

The same program extracted 23948 distinct blocks of length at most ten from exact Rosen orbits. Whenever $(B-C)/(\lambda A)$ was rational, it found only

rational value	matrix condition	number of blocks
0	$B=C$	34
-1	$\operatorname{Tr} P=+2$	1
1	$\operatorname{Tr} P=-2$	1.

Thus the observed exceptional locus consists of time-symmetric words and the two parabolic boundary words. The endpoint interpretation is exact:

$$-\frac{B}{A}=P^{-1}0, \qquad -\frac{C}{A}=-\frac{1}{P\infty}.$$

What is not yet inferred from the experiment is the global group lemma that these are the only rational differences for every admissible word. An additional exhaustive pass enumerated all 349524 words of length at most nine in the alphabet $\{-2,-1,1,2\}$. Of these, 293842 satisfy the necessary Rosen condition excluding three consecutive 1's or three consecutive (-1)'s. Every rational obstruction in that larger set was again $0,\pm1$: there were 344 symmetric zero cases and one case of each parabolic sign, with no nonintegral rational example. Repeating through length seven with digits of absolute value at most three checked 323006 short-run words and gave the same result.

This prospective group lemma is no longer being treated as a prerequisite for a norm upper bound. If the obstruction is rational, the Nair–Kumar–Rout alternative says that every exponentially accurate infinite rounding family must be pseudo-Pisot. Its two omitted conjugates are then bounded by one; in the exact determinant formula

$$Q=A+n\lambda uA^3+uA^2(C-B)+u^2A^3$$

this makes the distinguished factor smaller, while the two outer factors remain bounded by the corresponding tropical monomials. Thus rational shears describe the equality mechanism to be classified, rather than a source of a larger determinant. Proving their $0,\pm1$ classification would still identify that equality locus as symmetric or parabolic, but it does not remove the moving-divisor obstruction.

Route II.14.26 (Independent Rosen computations and their exact scope)**.** Panti's quadratic decreasing-block theorem is the source of (737); his paper also runs the analogous cubic $q=7$ map and exhibits a special hyperbolic fixed point, explaining why strict scalar-height termination is the wrong target in degree three [28]. The correct replacement is termination *or* preperiodicity.

Winsor's 2026 computations give an independent large-scale check of that replacement [64]. In his $q=7$ searches all tested Rosen expansions were finite or preperiodic, and all but one tested special eigenvalue belonged, up to sign, to the cyclic group generated by $u_7=\lambda^2+\lambda$. His Lemma 2.2 proves uniqueness up to cyclic permutation for a class of cyclic Rosen words by reducing them in $\langle S,U\mid S^2=U^q=1\rangle$. It does not prove the rational shear lemma above: equality $(B-C)/(\lambda A)\in\mathbf{Q}$ first produces a rational unipotent shear, not an equality or conjugacy of two elements of G_7. The usable contribution here is therefore the reduced-word coding and independent finite/preperiodic evidence, not a silent appeal to word uniqueness.

Experiment II.14.27 (Exact multi-parabolic coefficient audit)**.** The program `parabolic_multiblock_tests.py` constructs exact random connectors in G_7, expands $Q_{U_0T^{n_1}\cdots T^{n_s}U_s}(v)$ symbolically over $\mathbf{Z}[\lambda]$, and compares the coefficient of $n_1\cdots n_s$ with (766). All 200 random samples with one through five parabolic blocks agree exactly.

Corollary II.14.28 (The connector assumption in `curt11` is unnecessary)**.** *For every fixed finite connector type, the bounded-norm parabolic alternative in* ***`curt11.tex`*** *already forces a cusp. No separate multi-parabolic conjecture is needed.*

Corollary II.14.29 (Closing from the corrected nearest-wall alternative)**.** *Assume the nearest-wall defect estimate and tensor-induced base-locus alternative of* ***`curt11.tex`****, with the numerical inequality corrected as in* (765)*. Then the full cusp/special-fixed-point classification follows; no additional bounded-norm connector hypothesis is required.*

Proof. The corrected critical inequality forces the auxiliary evaluation to vanish. On the eigen-incidence component this gives a special hyperbolic fixer. On every fixed parabolic connector component, bounded norm gives a cusp by Theorem II.14.11. □

Proposition II.14.30 (Relative-gap synthesis criterion)**.** *Assume Conjecture II.2.12. Suppose the evaluation-dependent integral Asai construction has the following per-induced-block estimate on returns with $\nu(g)\to\infty$:*

$$\mathfrak{s}_g(d,r)\le kd\left(\frac{\delta_7}{2}-\eta\right)\nu(g)+o(kd\,\nu(g)) \tag{776}$$

for one $\eta>0$, at the corrected critical indices, and suppose its common-zero locus is the fixed-eigenline locus together with the parabolic boundary locus of ***`curt11.tex`****. Suppose also that the boundary alternative supplies bounded fixed-point norm. Then the full cusp/special-fixed-point classification follows.*

Proof. If $\nu(g)\to\infty$, equation (664) gives

$$kd\delta(g)\ge kd\frac{\delta_7}{2}\nu(g).$$

The strict $\eta kd\nu(g)$ margin in (776) makes the corrected fundamental inequality integral and forces common vanishing. The eigenline component gives a special hyperbolic fixer. On the boundary component the assumed bounded norm closes by Theorem II.14.11.

If $\nu(g)$ is bounded, pass to a subsequence with fixed induced complexity and fixed intervening letters from $\{A_2,A_3,A_4,A_5\}$. The return is then a fixed finite multi-parabolic connector family; its bounded-norm boundary alternative again closes by Theorem II.14.11. □

Route II.14.31 (How this synthesizes `curt10` and `curt11`)**.** The two drafts separate at exactly the quantity $\nu(g)$. The parabolic calculation in `curt11.tex` handles bounded ν, because only finitely many connector letters remain and the full monomial has the cusp-incidence factorization (766). The fusion/Hodge structure in `curt10.tex` is needed when ν grows: it turns each new connector crossing into additional outer surplus. The S-unit theorem then removes every boundary tail whose finite-place support stabilizes. The language of a finite divisor "truncated to time-truncated births" used in an earlier version is not intrinsic: it depends on the preceding orbit and is not the divisor of the current evaluation section. Consequently the joint obligation is not (776) in that primewise form. It is the stopped all-place degree estimate formulated in Section II.16. Fusion supplies an archimedean contraction inside that degree; the finite metric retains the whole remainder ideal. This is also why neither a scalar nearest-wall metric twist nor a fixed-ε Roth theorem resolves the issue.

Theorem II.14.32 (Conditional Hodge–S-unit synthesis)**.** *Assume the following two precise statements for recurrent returns of a Rosen orbit.*

1. *The tensor-induced integral Asai jet lattice satisfies* (765)*, with* $\mathfrak{s}_g = o(kd\delta)$*, and its common-zero locus is the union of the fixed-eigenline locus and a parabolic boundary locus.*
2. *Every orbit lying in the latter boundary locus either is an explicit cusp incidence, or its exact Rosen remainder sequence has only finitely many prime-ideal births.*

Then the full cusp/special-fixed-point classification follows.

Proof. Choose r/d at the stated critical scale. The corrected inequality and integrality force the tensor-induced auxiliary evaluations to vanish. On the fixed-eigenline component, $Q_g(v) = 0$, so g fixes the K-rational point and is special hyperbolic. An explicit parabolic incidence is a cusp. In the remaining boundary case, finitely many prime-ideal births imply finite total prime support; Corollary II.9.17 and Theorem II.9.14 close the orbit. □

Route II.14.33 (What the synthesis saves, and what it does not)**.** The one-block calculation in `curt11.tex`,

$$Q_{UT^NV}(v) = Q_{UV}(v) + N\lambda(e_2^t V v)\det(v, Ue_1),$$

already proves that a bounded nonzero norm along one parabolic parameter forces one of the two cusp incidences. Theorem II.14.11 extends it to every fixed number of positive blocks: the unique full monomial prevents cancellation once bounded parameters have been absorbed. The S-unit theorem is still useful beyond the explicitly parameterized parabolic locus. It replaces a complete geometric enumeration of any residual boundary component by the weaker boundary-to-finite-birth implication in part (2) of Theorem II.14.32. That implication is the exact bridge between the drafts; it must not be silently identified with primitive-polynomial saturation.

Proposition II.14.34 (Tensor induction does not linearize Maculan's quotient)**.** *Let k be a field of characteristic zero, let $C, E_1, \dots, E_q$ be k-vector spaces with $q > 1$, and let*

$$\eta_i : C \longrightarrow E_i$$

be linear maps. The diagonal tensor evaluation

$$\nu(c) = \eta_1(c) \otimes \cdots \otimes \eta_q(c) \tag{777}$$

is a homogeneous polynomial law of degree q. If it is nonzero, no linear map $C \to \bigotimes_i E_i$ agrees with it. Its canonical linearization is instead a map

$$\widetilde{\nu} : \operatorname{Sym}^q C \longrightarrow \bigotimes_i E_i.$$

This linearization need not be injective even when the direct-sum map $c \mapsto (\eta_1(c), \dots, \eta_q(c))$ is injective, and its kernel can contain no diagonal tensors while containing many non-diagonal tensors.

Proof. For $a \in k$, one has $\nu(ac) = a^q \nu(c)$. A linear map would give $a\nu(c)$, which is incompatible for a suitable a whenever $\nu(c) \neq 0$. Polarization gives the stated map on $\operatorname{Sym}^q C$.

For a concrete example take $k = \mathbf{Q}$, $C = \mathbf{Q}u \oplus \mathbf{Q}v$, and evaluate $uX + v$ at the two conjugate points $X = \sqrt{2}$ and $X = -\sqrt{2}$. The direct-sum evaluation is injective, whereas its diagonal tensor norm is

$$(u\sqrt{2} + v)(-u\sqrt{2} + v) = v^2 - 2u^2.$$

As a linear map on $\operatorname{Sym}^2 C = \langle u^2, uv, v^2 \rangle$, this has rank one and a two-dimensional kernel. On the other hand $v^2 - 2u^2 = 0$ has no nonzero rational solution. Thus injectivity on rational diagonal points does not imply injectivity of the linearized tensor representation. □

Route II.14.35 (Exact audit of the proposed Maculan replacement)**.** The preceding proposition applies at the exact point where `curt11.tex` proposes tensor induction. Maculan's GIT point is not built from the orthogonal sum of target jets. It is the Grassmannian point of the single K-subspace

$$K_{q,\mathbf{r}}(\mathbf{a}, t_a) \otimes_K \overline{K} = \bigcap_{\sigma} K_{\mathbf{r}}(\mathbf{a}^{(\sigma)}, t_a) \subset \Gamma(P, \mathcal{O}_P(\mathbf{r}));$$

see [15, §3.2.2]. That intersection and its Plücker point are what the higher-dimensional Dyson lemma proves semistable. The orthogonal direct sum $E = \bigoplus_\sigma E^{(\sigma)}$ appears only later in [15, §4.2.3–4.2.5]: the quotient

$$C = \Gamma(P, \mathcal{O}_P(\mathbf{r}))/K_{q,\mathbf{r}}$$

embeds *linearly* into E, so the Arakelov slope inequality applies. Its direct-sum operator norm is the maximum of the conjugate operator norms.

Replacing this injection by $F \mapsto \bigotimes_\sigma \eta^{(\sigma)}(F)$ does turn a maximum into a product of norms, but Proposition II.14.34 shows that it is not a linear injection of C. Consequently the slope inequality used in Maculan's equation (4.2.3) cannot be applied to it. Passing to $\operatorname{Sym}^q C$ restores linearity but creates a generally large kernel and changes both the Grassmannian point and the Dyson semistability problem. The Faltings–Totaro theorem cited in `curt11.tex` says that tensor products of semistable *filtered vector spaces* remain semistable [16, Theorem 0.2]; it does not assert that Maculan's Plücker point remains semistable after imposing the diagonal Veronese, nor that the resulting linearized norm map is injective.

Thus the tensor paragraph of `curt11.tex` is not yet a proof of the nearest-wall estimate. The precise missing statement is a *norm-form slope theorem*: an arithmetic height inequality for the degree-q polynomial law (777), restricted to its rational diagonal locus, with the nearest-wall determinant as its archimedean contribution. Ordinary linear slope theory does not supply this theorem.

Route II.14.36 (Four viable repairs of tensor induction)**.** The audit leaves four mathematically distinct repairs.

1. *Direct Plücker height.* Do not change Maculan's GIT point or evaluation map. In the quadratic one-variable case the conjugate-jet intersection is explicitly $Q^r \operatorname{Sym}^{d-2r}$, so compute the height of its Plücker line directly from Proposition II.14.4. This bypasses the direct-sum estimate instead of replacing the direct sum. The remaining issue is the arithmetic descent of the nearest-wall Hodge metric to that same line.
2. *Norm-form arithmetic GIT.* Work on the Veronese image $\nu_q(\mathbf{P}C) \subset \mathbf{P}(\operatorname{Sym}^q C)$, rather than on the ambient symmetric power. Prove a lower-height formula only for rational points of that image. The quadratic example in Proposition II.14.34 shows why rational anisotropy can remove diagonal zeroes even though the ambient linearized map has a kernel. The required estimate is nonlinear and must include the parabolic equality divisor.
3. *Restriction of scalars and determinant norm.* Retain the linear injection $C \hookrightarrow \operatorname{Res}_{L/K} E_L$, but replace the direct-sum metric estimate by the determinant of multiplication on the finite K-algebra. This produces a genuine norm form on a line bundle. One must prove that the determinant height, rather than its largest singular value, enters the fundamental formula; the nearest-wall calculation predicts exactly that determinant.
4. *Multilinear auxiliary system.* Allow independent conjugate polynomials $F_1, \dots, F_q$ so that the tensor evaluation is linear in each factor, use tensor semistability there, and then prove a separate rank-one descent theorem forcing a useful tensor to be decomposable and Galois diagonal. The

equations are the 2×2 flattening minors. This approach makes the missing zero estimate explicit rather than hiding it in the word "Veronese."

The first repair is the minimum change to Maculan's proof and should be attempted first. The third is the closest synthesis with the S-unit method: the determinant norm is sensitive to the product over embeddings, while finite-place prime births are already encoded by the integral Asai lattice. It should be attempted before a new general nonlinear GIT theorem.

Route II.14.37 (Independent conjugate variables do not by themselves preserve the nearest-wall gain)**.** One can restore literal linearity by abandoning the Galois-diagonal coefficient space. On $(\mathbf{P}^1)^3$, take independent multihomogeneous forms in the three conjugate variables. Their tensor evaluation at the rational Asai point is linear in the full tensor coefficient space, and multilinear in the three individual factors.

The nearest-wall determinant, however, belongs to a rank-two pencil in which two Asai factors are fixed eigenlines and only one outer factor varies. A general Asai point does not lie on that pencil. Evaluating a pencil section at it requires contraction by the two moving eigen-covectors. The heights of those covectors are precisely the target heights which the direct-sum Roth estimate records. Thus independent variables solve the polynomial-law defect but do not remove the moving height; a separate rank-one descent/flattening-minor theorem is still needed to recover a Galois-diagonal tensor without paying that cost.

The related decomposable-form shortcut is also insufficient. Before a high-index Dyson construction, the norm of the basic evaluation has archimedean exponent

$$-a + b + c = \chi.$$

This is exactly the direct fixed-point norm tested in the adversarial records, where it can grow at the trivial height scale. A parametric Subspace Theorem for the bare norm form therefore restates the failed direct norm estimate. It becomes relevant only after an auxiliary system removes one outer weight; constructing that system with controlled moving covectors is the original nearest-wall problem.

Proposition II.14.38 (Exact Plücker bypass in the quadratic case)**.** *Let $Q \in K[X, Y]$ be a separable quadratic form with conjugate zeroes $a, a^\iota \in \mathbf{P}^1(\overline{K})$. For $d \geq 2r$, the degree-d forms having order at least r at both zeroes form the K-space*

$$K_{2,d}(a, r/d) = Q^r \operatorname{Sym}^{d-2r}(K^2)^\vee. \tag{778}$$

It has rank $k = d - 2r + 1$, and its Plücker line is

$$\det K_{2,d}(a, r/d) \simeq (\det(K^2)^\vee)^{kd/2}$$

after the two boundary covectors are used as the split frame. Therefore Maculan's original semistable Grassmannian point admits a direct target height calculation; neither tensor induction nor a diagonal Veronese is needed. If the integral module is saturated and its determinant metric is the nearest-wall Hodge determinant metric, its realization error is $O(k \log(d + 2))$.

Proof. Over the splitting field, divisibility by the two coprime linear factors to order r is equivalent to divisibility by Q^r. Galois invariance descends the resulting space to K, proving (778). The rank and determinant are Proposition II.14.1. Maculan's semistability theorem applies to the Grassmannian point of this very intersection, not to the later direct-sum target quotient. Hence one may insert its exact Plücker height into the fundamental formula and omit the coarse direct-sum bound. Saturation removes finite-place index terms, and Proposition II.14.4 gives the stated comparison of archimedean determinant norms. □

Proposition II.14.39 (Metric-gauge cancellation in Maculan's polarization)**.** *Let V be a rank-two Hermitian space, $W = \operatorname{Sym}^d V^\vee$, and*

$$X = \mathrm{Gr}_u(W) \times \mathrm{Gr}_v(W), \qquad \mathcal{L} = \mathcal{O}(1) \boxtimes \mathcal{O}(1),$$

with the diagonal action of $G = \mathrm{SL}(V)$. Change the metric on V by an element $A \in \mathrm{GL}(V)$, and use the induced metrics on both Plücker factors and on the GIT quotient. Then Maculan's fundamental difference

$$h_{\overline{\mathcal{L}}}(P) + \frac{1}{[K:\mathbf{Q}]} \sum_w \iota_w(P) - h_{\min}(X, \overline{\mathcal{L}})//G \tag{779}$$

is unchanged at that archimedean place.

More particularly, the scalar part of A multiplies the norm on the Segre–Plücker line by

$$|\det A|^{-d(u+v)/2}.$$

It shifts the point height and the quotient minimum by the same amount. The determinant-one part is absorbed by the infimum over $\mathrm{SL}(V)$ in the instability measure. The same conclusion holds if one merely rescales the metric of one Plücker factor by a line: its contribution occurs with the same coefficient on both sides of the fundamental formula.

Proof. After discarding a unitary phase, write $A = sB$, where $s = |\det A|^{1/2}$ and $B \in \mathrm{SL}(V)$. The scalar s acts on $W = \operatorname{Sym}^d V^\vee$ by s^{-d}, hence on $\bigwedge^u W \otimes \bigwedge^v W$ by $s^{-d(u+v)}$. This is the displayed determinant factor. Because it is a constant rescaling of the polarized line on all of X, the quotient metric defined by the fiberwise supremum is rescaled by exactly the same constant. Thus the two height shifts cancel in (779).

Changing by B identifies the new norm of a point with the old norm of its B-translate. The quotient is G-invariant, and the definition of ι_w already takes the infimum over all $G(K_w)$; translating the variable of that infimum by B changes nothing. Finally, a scalar rescaling of only one projective factor is again a constant rescaling of $\mathcal{L}$, so the identical argument applies. □

Corollary II.14.40 (The direct Plücker bypass is structural, not a free Roth gain)**.** *The identity $K_{2,d}(a, r/d) = Q^r \operatorname{Sym}^{d-2r}$ permits an exact calculation of Maculan's original target Plücker height and removes the nonlinear tensor-induction error. It does not, by itself, replace the ordinary target height by $kd\delta$. If the nearest-wall determinant is inserted only by changing or twisting the target metric, its entire contribution is canceled by the corresponding shift of the GIT quotient minimum.*

Proof. The first statement is Proposition II.14.38. The nearest-wall mean is precisely the determinant, hence scalar, part of the rank-two metric. Proposition II.14.39 applies to that part. An uncanceled term can therefore arise only from a new arithmetic covariant or a new relative-height inequality tying the source and target points together, not from a choice of Hermitian gauge. □

Proposition II.14.41 (Critical parabolic Plücker height is of size $r \log N$)**.** *Let*

$$Q_N = X^2 + N\lambda XY - (N\lambda^2 + \lambda^2 - 1)Y^2$$

be the primitive quadratic form above, and put $d = 2r$. For the standard adelic coefficient metric, Maculan's target Grassmannian point has rank one and

$$h([K_{2,2r}]) = h([Q_N^r]) = r \log N + O_7(r \log(r+2)). \tag{780}$$

On the other hand,

$$d\delta_N = 2r\delta_7 + o_N(r).$$

Consequently the comparison between the standard target metric and the nearest-wall normalized target metric cannot be $O(k\log(d+2))$, uniformly in the return g, even when $k=1$.

Proof. At $d=2r$, equation (778) is the one-dimensional space $K_{2,2r} = KQ_N^r$. The coefficient point of

$$Q_N = X^2 + N\lambda XY - (N\lambda^2 + \lambda^2 - 1)Y^2$$

has height

$$h([Q_N]) = \log N + O_7(1) :$$

at each of the three real embeddings the largest coefficient is N times a fixed nonzero constant up to $O(1)$, and primitivity removes a finite-place content term. The power map $[Q] \mapsto [Q^r]$ is the degree-r Veronese map followed by the multiplication map from ternary coefficient monomials to binary coefficients. The binomial coefficients change its archimedean height by $O(r\log(r+2))$, while Gauss's lemma again gives no finite content. This proves (780). The second assertion is the already computed asymptotic $\delta_N \to \delta_7 > 0$. Since here $k=1$, the difference grows like $r\log N$ for fixed r and cannot be $O(\log(d+2))$. □

Route II.14.42 (Correct interpretation of the determinant-only lemma)**.** Proposition II.14.4 remains a correct statement about two metrics *after their determinant lines have been identified.* Proposition II.14.41 computes the cost of that identification in the actual arithmetic problem: it is not bounded uniformly in the moving target. Treating the identification as a free use of $\Psi_i^{kd/2}$ deletes precisely the moving height that Maculan's fundamental formula retains. Proposition II.14.39 then shows that inserting the same cost as a scalar metric twist merely moves it into the quotient minimum. Thus the direct Plücker repair is valid only as a replacement for the nonlinear tensor map; it does not prove the nearest-wall inequality.

Proposition II.14.43 (Two unconditional pairwise Hodge descents)**.** *Let the heptagonal modular embedding be*

$$t \longmapsto (t, F_1(t), F_2(t)) \in \mathbb{H}^3,$$

where each $F_i : \mathbb{H} \to \mathbb{H}$ is a strict holomorphic contraction. Let $\gamma(s)$ be the unit-speed geodesic directed to x_0, and fix $(a,b) \in K^2$ with $a + bx = 0$. Put

$$\eta_0(s) = \log \frac{|a_0 + b_0\gamma(s)|^2}{\operatorname{Im}\gamma(s)}, \qquad \eta_i(s) = \log \frac{|a_i + b_i F_i(\gamma(s))|^2}{\operatorname{Im} F_i(\gamma(s))} \quad (i=1,2).$$

Then

$$(\eta_0 + \eta_i)'(s) \le -1 + \|DF_i(\gamma(s))\|_{\mathrm{hyp}} \le 0. \tag{781}$$

On the inverse image of a fixed compact part of the Teichmuller curve there are $\alpha_i > 0$ for which the right side is at most $-\alpha_i$. Hence, if $T(s)$ denotes total compact-part time,

$$\eta_0(s) + \eta_i(s) \le \eta_0(0) + \eta_i(0) - \alpha_i T(s) \qquad (i=1,2). \tag{782}$$

Proof. McMullen's Busemann calculation [17, §§4–5] gives $\eta_0' = -1$ because γ points at the zero of $a_0 + b_0 t$, and $\eta_i' \le \|DF_i\|_{\mathrm{hyp}}$ by Schwarz–Pick. Each norm is strictly less than one. It is bounded by $1 - \alpha_i$ over a compact quotient set, and integration proves the last assertion. □

For a quadratic trace field, the single pairwise sum in (781) is the full unit-invariant Hodge height, which is why McMullen's descent forces every field point into a cusp. For the cubic field there are two decreasing pairwise sums but neither is a height on $\mathbf{P}^1(K)$; the following proposition identifies the exact reason.

Lemma II.14.44 (One global unit balances all three real sizes)**.** *There is a constant C_K, depending only on the cubic field K, with the following property. For every primitive pair $(p,q) \in \mathcal{O}_K^2$, some unit $u \in \mathcal{O}_K^\times$ satisfies*

$$|\log \max\{|\sigma_i(up)|, |\sigma_i(uq)|\} - h([p:q])| \le C_K \qquad (i = 0, 1, 2). \tag{783}$$

The same unit is used at all times after transporting the remainder recurrence. It does not change any $|\,\mathrm{N}\, r_j|$, M_j, or U_j.

Proof. Put

$$H_i = \log \max\{|\sigma_i(p)|, |\sigma_i(q)|\}.$$

Primitivity gives $\sum_i H_i = 3h([p:q])$. The logarithms of absolute values of units form a full lattice

$$\Lambda_K = \{(\log|\sigma_i(u)|)_i : u \in \mathcal{O}_K^\times\} \subset \{(v_0, v_1, v_2) : v_0 + v_1 + v_2 = 0\}$$

by Dirichlet's unit theorem. Choose a bounded fundamental parallelogram for this lattice. Translating $(H_i - h([p:q]))_i$ into that parallelogram by an element of Λ_K proves (783). Since $|\,\mathrm{N}\, u| = 1$, all absolute remainder norms are unchanged; since a common scaling cancels in every ratio r_j/r_{j-1}, so is U_j. □

This elementary lemma is the legitimate use of the unit lattice. It does not analyze primes and it does not rebalance the orbit at every step. It fixes one adelic gauge at the beginning, with a uniform error, and then lets Hodge contraction transport that same gauge.

Proposition II.14.45 (Segmentwise Hodge–norm inequality for Rosen remainders)**.** *Let $x = [r_0 : r_{-1}] \in \mathbf{P}^1(K)$, with a primitive integral pair, and let γ be the geodesic directed to $\sigma_0(x)$. Suppose s_j is a sequence of Rosen return times and $P_j \in G_7$ is oriented so that*

$$P_j[r_0 : r_{-1}] = [r_j : r_{j-1}]$$

and $P_j\gamma(s_j)$ lies in one fixed compact return box. Let T_j be the total thick-part time before s_j. Then there are constants $\alpha_1, \alpha_2 > 0$ and C_7, depending only on the box and the heptagonal field, such that, after the single balancing of Lemma II.14.44,

$$\log P_{0,j} + \log P_{i,j} \le 2h(x) + C_7 - \alpha_i T_j, \qquad P_{i,j} = \max\{|\sigma_i(r_j)|, |\sigma_i(r_{j-1})|\}, \quad i = 1, 2. \tag{784}$$

Since $|\sigma_0(r_j/r_{j-1})| < 1$, this implies the local Hodge–norm slack

$$\log|\,\mathrm{N}(r_{j-1})| + \log|\sigma_0(r_{j-1})| \le 4h(x) + 2C_7 - (\alpha_1 + \alpha_2)T_j. \tag{785}$$

Proof. Let $\ell = (-r_0, r_{-1})$, so the affine linear form $-r_0 + r_{-1}t$ vanishes at x. The elementary naturality identity for the Hodge norm is

$$\frac{|a + b(gt)|^2}{\Im(gt)} = \frac{|a' + b't|^2}{\Im t},$$

where (a', b') is the contragredient transform of (a, b). Equivalently, if a group element sends the endpoint x to $[r_j : r_{j-1}]$, the transported covector is $(-r_j, r_{j-1})$, up to the common fixed initial ideal. Unimodularity of every Rosen branch preserves primitivity, so there is no varying finite-place factor.

Apply this identity at all three embeddings at time s_j. On the fixed compact return box, $\Im z$, $\Im F_i(z)$, and the condition numbers of evaluation at $z, F_i(z)$ are bounded above and below. Therefore the corresponding Hodge norms of the transported covector are uniformly comparable to $P_{i,j}^2$. Integrating (782), applying Lemma II.14.44 to the initial pair, and absorbing the fixed comparison constants gives (784).

The Rosen interval satisfies $|\sigma_0(r_j/r_{j-1})| \le \lambda_0/2 < 1$, hence $P_{0,j} = |\sigma_0(r_{j-1})|$. Also $P_{i,j} \ge |\sigma_i(r_{j-1})|$. Adding the two inequalities in (784) and using

$$P_{0,j}^2 P_{1,j} P_{2,j} \ge |\sigma_0(r_{j-1})|^2 |\sigma_1(r_{j-1})\sigma_2(r_{j-1})| = |\sigma_0(r_{j-1})|\,|\,\mathrm{N}(r_{j-1})|$$

proves (785). □

Proposition II.14.46 (Two Hodge compressions equal one algebraic norm slack)**.** *Let $\xi, \eta \in G_7\infty$ be distinct cusps, let g_ξ, g_η be their positive primitive parabolic generators, and put*

$$A(\xi, \eta) = -\operatorname{Tr}\big((g_\xi - I)(g_\eta - I)\big) \in \mathcal{O}_K.$$

Write $A_i = \sigma_i(A)$. For the two contractive modular embeddings, let $r_i(\xi, \eta)$ be McMullen's cusp-compression ratio. Then

$$A(\xi, \eta) = \lambda^2 \det(v_\xi, v_\eta)^2, \qquad \mathrm{N}\,A(\xi, \eta) = \mathrm{N}(\det(v_\xi, v_\eta))^2, \tag{786}$$

where v_ξ, v_η are the primitive first columns of elements of G_7 carrying ∞ to the two cusps. Moreover,

$$r_i(\xi, \eta) = \frac{A_i}{A_0} \quad (i = 1, 2), \qquad r_1(\xi, \eta) r_2(\xi, \eta) = \frac{\mathrm{N}\,A(\xi, \eta)}{A_0^3}. \tag{787}$$

If the central geodesic from ξ to η spends time T in a fixed thick part and

$$\kappa_i = \inf_{\text{thick}} (1 - \|DF_i\|_{\mathrm{hyp}}) > 0,$$

then

$$0 \le \log|\,\mathrm{N}\,A| \le 3 \log A_0 - (\kappa_1 + \kappa_2) T. \tag{788}$$

Proof. Write $T = I + \lambda E_{12}$. If $v = (a, c)^t$ is the first column of $g \in G_7$, then

$$gTg^{-1} - I = \lambda \begin{pmatrix} -ac & a^2 \\ -c^2 & ac \end{pmatrix} = \lambda N_v.$$

The rank-one calculation $-\operatorname{Tr}(N_v N_w) = \det(v, w)^2$ proves the first formula in (786). Since $\mathrm{N}(\lambda)^2 = 1$, the norm formula follows.

McMullen's parabolic pairing is $\langle g, h \rangle = \operatorname{Tr}((g - I)(h - I))$. His naturality theorem sends the positive primitive generator at a cusp to the positive primitive generator for each contractive twist. Consequently his definition of compression gives

$$r_i = \frac{-\langle \sigma_i(g_\xi), \sigma_i(g_\eta) \rangle}{-\langle g_\xi, g_\eta \rangle} = \frac{\sigma_i(A)}{A_0}.$$

All three A_i are positive. Multiplication over $i = 1, 2$ proves (787). Corollary 8.5 of [17], applied separately to F_1, F_2, gives $r_i \le e^{-\kappa_i T}$. Taking logarithms in (787) yields the upper bound. The lower bound follows because A is a nonzero algebraic integer. □

The preceding proposition has a useful two-coordinate refinement which is obscured by multiplying $r_1 r_2$. On McMullen's modular-symbol category put

$$\mathbf{r} = (r_1, r_2) : \mathscr{S}(G_7) \longrightarrow [0,1]^2.$$

Each coordinate is a continuous multiplicative functor, hence so is $\mathbf{r}$, and $\mathbf{r}(\sigma_n) \to (0,0)$ when the thick-part length of σ_n tends to infinity. If $d = \det(v_\xi, v_\eta)$, then the two exact formulas are

$$\log r_i = 2\left(\log|\sigma_i d| - \log|\sigma_0 d|\right) + 2\log\frac{|\lambda_i|}{\lambda_0}, \qquad i = 1,2. \tag{789}$$

Thus joint cusp compression is precisely the projective shape of the three-component Asai lattice. The signed pair is injective up to rational scaling: equality for d, d' says that all three embeddings of d/d' are equal. This is the full-group explanation for why both outer representations, rather than their scalar product, are required.

There is a corresponding properness statement. If $|\operatorname{N} A(\xi,\eta)| \le B$ and $r_1, r_2 \ge \delta > 0$, then

$$A_0^3 = \frac{|\operatorname{N} A|}{r_1 r_2} \le \frac{B}{\delta^2}.$$

The logarithm of A_0 is the renormalized length of the associated cusp-to-cusp geodesic. Properness of the relative metric on $\langle T\rangle\backslash G_7/\langle T\rangle$ leaves only finitely many modular symbols. Hence an infinite family of distinct symbols with bounded radial norm must escape through $r_1 r_2 \to 0$; it cannot wander in a third unrecorded unit direction. This is an unconditional full-group compactness result.

The same construction gives a particularly concrete description of what thinness contributes. Let $\mathcal{E}_7$ denote the nonzero matrix entries of all elements of G_7. If $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, then

$$c = \det(e, ge), \quad -a = \det(f, ge), \quad d = \det(e, gf), \quad -b = \det(f, gf), \qquad e = (1,0)^t,\ f = (0,1)^t.$$

Every determinant on the right is a determinant of two primitive cusp columns. Consequently, for every $m \in \mathcal{E}_7$,

$$\frac{|\sigma_i m|}{|\sigma_0 m|} = \frac{\lambda_0}{|\lambda_i|}\sqrt{r_i(\xi,\eta)} \le \frac{\lambda_0}{|\lambda_i|}, \qquad i = 1,2, \tag{790}$$

for a suitable oriented cusp pair (ξ, η).

McMullen proves that the modular-symbol category, after adding its one point at infinity, is compact and countable, and that compression extends continuously with value zero at infinity. Applying this to both modular embeddings at once proves the following statement.

Proposition II.14.47 (Compact joint ratio set)**.** *The closure of*

$$\left\{ \left(\frac{|\sigma_1 m|}{|\sigma_0 m|}, \frac{|\sigma_2 m|}{|\sigma_0 m|} \right) : m \in \mathcal{E}_7 \right\}$$

is compact and countable. It is the coordinatewise square-root image, up to the two fixed factors $\lambda_0/|\lambda_i|$ and finitely many signs, of the joint compression image. A sequence represented by modular symbols of unbounded thick-part length tends to $(0,0)$.

This is a property of the entire thin Galois image. It would be false for $\mathrm{SL}_2(\mathcal{O}_K)$: diagonal unit matrices already make the two ratios unbounded. It is also much stronger than Zariski density, which supplies no inequality. The central Fuchsian lattice, the abstract presentation $C_2 * C_7$, and equidistribution of its cusps do not by themselves see this compact joint Galois shape.

There is a two-anchor consequence for the actual remainder incidence. Fix a primitive $P = (p, q)^t$, let $v = (a, c)^t$ be a primitive cusp column representing $\xi \neq 0, \infty$, and set

$$R_i(v) = \max\{r_i(\infty, \xi), r_i(0, \xi)\}, \qquad H_0(v) = \max\{|\sigma_0 a|, |\sigma_0 c|\}.$$

Since $c = \det(e, v)$ and $-a = \det(f, v)$,

$$\max\{|\sigma_i a|, |\sigma_i c|\} \leq \frac{\lambda_0}{|\lambda_i|} H_0(v) \sqrt{R_i(v)}, \tag{791}$$

$$|\,\mathrm{N}(pc - qa)| \leq C_P |\sigma_0(pc - qa)| H_0(v)^2 \sqrt{R_1(v) R_2(v)}. \tag{792}$$

The second inequality follows by bounding $\sigma_i(pc - qa)$ with the first and multiplying over the three places. Its left side is either zero or a positive integer. We have therefore proved an unconditional full-group closing test:

$$C_P |\sigma_0(pc - qa)| H_0(v)^2 \sqrt{R_1(v) R_2(v)} < 1 \quad \Longrightarrow \quad [P] = [v]. \tag{793}$$

For a geometric convergent one normally has $|\sigma_0(pc - qa)| \ll_P H_0(v)^{-1}$. On such a subsequence the sufficient condition becomes

$$H_0(v) \sqrt{R_1(v) R_2(v)} \longrightarrow 0. \tag{794}$$

McMullen's thick-part estimate proves this whenever joint compression beats central cusp growth. A special periodic orbit lies on the equality wall: its central growth is exactly compensated by neutral unit holonomy. Thus a hypothetical third behavior is no longer an arbitrary failure of height descent. It must be a nonperiodic sequence of modular symbols for which the product in (794) stays bounded away from zero.

The topology of modular symbols separates this remaining behavior into two cases. A sequence of bounded symbolic degree can degenerate only by inserting finitely many increasingly long parabolic blocks; these are the finite multi-parabolic families treated above. A sequence of unbounded degree has $(R_1, R_2) \to (0, 0)$, but topology alone gives no rate relative to $H_0(v)$. The balanced many-block obstruction is exactly the claim that the latter can grow quickly enough to cancel the former. This is why the actual group structure makes a decisive reduction but does not, by itself, finish the arithmetic argument.

One possible shortcut would be to hope that the two maxima in (792) always come from the same entry. The full group disproves this at its shortest scale. With $\alpha = \lambda^2 - 1$, the first column of

$$A_3 = \begin{pmatrix} \alpha & \alpha \\ \lambda & \alpha \end{pmatrix}$$

satisfies

$$\log \frac{|\sigma_1 \alpha|/|\sigma_0 \alpha|}{|\sigma_1 \lambda|/|\sigma_0 \lambda|} = 0.3681382954786467\ldots > 0,$$

whereas

$$\log \frac{|\sigma_2 \alpha|/|\sigma_0 \alpha|}{|\sigma_2 \lambda|/|\sigma_0 \lambda|} = -1.0303112263277456\ldots < 0.$$

These signs also follow exactly from the three conjugate values of λ. Thus α supplies the first outer maximum and λ the second. The hoped-for total ordering of the joint ratio set is false before any iteration begins.

The correct replacement is a four-component integral tensor. Let τ generate $\mathrm{Gal}(K/\mathbf{Q})$ and put, for $v = (a, c)^t$,

$$Z_{12}(v) = \tau(v) \otimes \tau^2(v) = \big(\tau(a)\tau^2(a), \tau(a)\tau^2(c), \tau(c)\tau^2(a), \tau(c)\tau^2(c)\big). \tag{795}$$

If $(a, c) = \mathcal{O}_K$, the content ideal of this vector is the unit ideal. Writing $H_i(v) = \max\{|\sigma_i a|, |\sigma_i c|\}$, one has exactly

$$\|\sigma_0 Z_{12}(v)\|_\infty = H_1(v)H_2(v), \qquad h([Z_{12}(v)]) = 2h([v]). \tag{796}$$

Indeed, the logarithmic sup norms at the three embeddings are $H_1 + H_2, H_2 + H_0, H_0 + H_1$, with logarithms understood on the right.

For fixed $P = (p, q)^t$, define the tensor linear form

$$L_P = (\tau(p)Y - \tau(q)X) \otimes (\tau^2(p)Y - \tau^2(q)X).$$

Then

$$L_P(Z_{12}(v)) = \tau(pc - qa)\tau^2(pc - qa), \qquad (pc - qa)L_P(Z_{12}(v)) = \mathrm{N}(pc - qa). \tag{797}$$

Thus the asynchronous product of the two outer incidences is one fixed linear evaluation on a primitive rank-four integral orbit. This is an exact, group-compatible linearization, not a choice of the larger entry at each place.

The tensor calculation has a group-wide version which is more useful than the scalar identity. Let $V = \mathcal{O}_K^2$ and define

$$\mathbb{W}_i = V \otimes {}^{\tau^i}V, \qquad \rho_i(g) = g \otimes \tau^i(g), \qquad w_i(P) = P \otimes \tau^i(P), \quad i = 1, 2.$$

These are integral representations of the full group. They satisfy $\rho_i(g)w_i(P) = w_i(gP)$, and equality of the two projective tensors forces equality of their first factor lines. Hence

$$[\rho_i(g)w_i(P)] = [w_i(P)] \quad \Longleftrightarrow \quad g[P] = [P]. \tag{798}$$

This is the state-separation property which was missing from every scalar remainder.

The canonical alternating pairing shows what the fixed representation does not supply. If $r_g = \det(P, gP)$, then

$$B_i\big(w_i(P), \rho_i(g)w_i(P)\big) = r_g\tau^i(r_g), \qquad B_1B_2 = r_g\,\mathrm{N}(r_g). \tag{799}$$

The fixed contracting tensor therefore has a zero of positive order on the incidence divisor. The norm-record line needs the fractional ideal $(r_g)^{-1}$, a pole on that divisor. Reversing this finite sign requires division by the moving evaluation; it cannot be achieved by choosing a different smooth Hodge metric on the fixed tensor.

The division becomes canonical when the complete state is retained. For a primitive $X = (u, v)^t$, $v \neq 0$, set

$$\mathbf{U}(X) = \frac{X^{\otimes 2} \otimes \tau(X) \otimes \tau^2(X)}{v} \in \mathrm{Sym}^2 V \otimes {}^{\tau}V \otimes {}^{\tau^2}V. \tag{800}$$

Its content ideal is exactly $(v)^{-1}$. If $M_i(X) = \max\{|\sigma_i u|, |\sigma_i v|\}$, then

$$\log \|\sigma_0\mathbf{U}(X)\|_\infty = 2\log M_0 + \log M_1 + \log M_2 - \log|\sigma_0 v| = \mathscr{E}(X). \tag{801}$$

Thus the tensor has exactly the reciprocal finite divisor and exactly the pairwise Arakelov metric, with no unspecified comparison map.

If $X' = AX = (u', v')^t$, then the full-group representation gives the exact formula (with $\rho_{\mathbb{W}}(A) = \mathrm{Sym}^2(A) \otimes \tau(A) \otimes \tau^2(A)$)

$$\rho_{\mathbb{W}}(A)\mathbf{U}(X) = \frac{v'}{v}\mathbf{U}(X'). \tag{802}$$

Projective equality of two $\mathbf{U}$-tensors implies equality of their $\mathrm{Sym}^2 V$-factor lines, hence equality of the original states. For a special return $X' = \epsilon X$, one has $\mathbf{U}(X') = \mathrm{N}(\epsilon)\mathbf{U}(X) = \pm\mathbf{U}(X)$.

The numerator has worst distinguished Hodge weight $-2a_0 + a_1 + a_2$ and therefore contracts. Formula (802) reduces the entire missing transport problem to the scalar cocycle v'/v. A norm record controls its field norm, but not its distinguished absolute value; the outer two absolute values can still absorb it. The stopped maximal-slope theorem is exactly the assertion that this absorption cannot keep defeating the Hodge contraction unless the normalized projective tensor repeats.

There is a complete audit of higher central tensor powers. For $m \geq 0$ define, with the evident interpretation at $m = 0$,

$$\mathbf{U}_m(X) = \frac{X^{\otimes m} \otimes \tau(X) \otimes \tau^2(X)}{v^{m-1}}. \tag{803}$$

Changing the primitive representative by $k \in K^\times$ multiplies this tensor by $\mathrm{N}(k)$, so its projective state is well defined. A direct calculation gives

$$\log \|\sigma_0 \mathbf{U}_m(X)\|_\infty = \log |\,\mathrm{N}\, v| + m\eta_0 + \eta_1 + \eta_2 = 3h(X) + (m-1)\eta_0, \tag{804}$$

$$\rho_m(A)\mathbf{U}_m(X) = \left(\frac{v'}{v}\right)^{m-1} \mathbf{U}_m(X'). \tag{805}$$

For a unit return $X' = \epsilon X$, every member of the family satisfies $\mathbf{U}_m(X') = \mathrm{N}(\epsilon)\mathbf{U}_m(X) = \pm\mathbf{U}_m(X)$.

This rules out the most obvious representation-theoretic escape. Raising m improves the Hodge weight to $-ma_0 + a_1 + a_2$, but raises the moving incidence pole to the same power $m-1$. Equations (804) and (805) show exact cancellation, not merely an unsuccessful numerical trend. A useful new structure must change Hodge weight without adding another power of v^{-1}. The effective order-seven Higgs divisor is presently the only candidate with that property.

Route II.14.48 (Outcome of the central-power audit)**.** All symmetric-power variants have now been treated at once. They preserve the correct neutral special holonomy but cannot strengthen the record inequality, because analytic contraction and the finite incidence pole increase in lockstep. Further searches over tensor exponents would only reproduce (804).

At the distinguished real place the central factor is the contracting boundary line. The worst possible weights of the two tensors are

$$-a_0 + a_1, \qquad -a_0 + a_2.$$

The two modular maps give $a_i \leq a_0$, with a uniform strict inequality on each fixed compact core. Therefore both real tensors remain bounded at all reduced returns and contract to zero if the ray accumulates unbounded compact-part length. This conclusion uses the full modular embedding and does not compare individual Rosen letters. In particular, arbitrarily many increasingly deep cusp excursions do not erase the contraction already earned at earlier compact returns.

The following degree table records exactly why this is not yet a lattice proof. The three real realizations of each $\mathcal{O}_K$-representation have characters and Hodge degrees

	characters	degrees
$\mathbb{W}_1$	$(-1,1,0),\ (0,-1,1),\ (1,0,-1)$	$-1/7,\ 1/14,\ 1/14$
$\mathbb{W}_2$	$(-1,0,1),\ (1,-1,0),\ (0,1,-1)$	$-1/14,\ 1/7,\ -1/14.$

(806)

Here $(\deg \mathcal{L}_0, \deg \mathcal{L}_1, \deg \mathcal{L}_2) = (5, 1, 3)/28$. Each row sums to zero. The doubled-central component has character $(-2, 1, 1)$, while its conjugates have degrees $-3/14, 3/14, 0$. Thus the negative component cannot descend by itself: Galois descent adds a positive component of equal total degree.

If $gP = \epsilon P$, then

$$\rho_i(g) w_i(P) = \epsilon \tau^i(\epsilon) w_i(P).$$

This explains why the obstruction is also the desired special alternative. One real component contracts, its conjugates compensate by a unit, and the projective tensor repeats. For an arbitrary integral sequence the same compensation could be caused by unrelated units. The remaining theorem must show that compensation inside the thin orbit is either bounded or is the power of one holonomy fixing $[P]$.

Route II.14.49 (Outcome of the full pairwise-representation test)**.** This route makes a genuine advance. It constructs primitive integral states, proves nonpositive contraction for arbitrary group elements, and proves exact state separation in (798). It also shows that the remaining failure has nothing to do with a weak Roth exponent or with sparse cusp time. It is the Galois descent problem displayed in (806): positive conjugate degree must be charged to the reciprocal radial incidence ideal; equation (799) shows why this charge is not already present in the flat representation. A valid completion is an Arakelov maximal-slope bound for these two tensors at first norm records; zero slope then gives special holonomy by state separation.

Route II.14.50 (Zero SAF and the double-heptagon gcd algorithm)**.** There is a natural translation-surface reformulation. In the standard normalization of the double heptagon, a direction has slope in K if and only if its Sah–Arnoux–Fathi invariant vanishes; equivalently it is an algebraically periodic direction in the terminology of Calta–Smillie and Arnoux–Schmidt [6, 7]. This uses the entire homology of the surface, not a chosen generator.

It does not add a missing inequality. Boulanger's gcd algorithm uses the six cones cut out by the staircase diagonals. Termination is equivalent to a parabolic direction, and eventual repetition produces a hyperbolic fixed direction [8]. For the double heptagon, the statement that every zero-SAF direction terminates or repeats is therefore the present classification itself, expressed in the positive cone model.

In Hodge language the reason is transparent. A zero-SAF direction gives the rational Lagrangian $KP \subset K^2$, whose Pluecker data is the Asai tensor attached to P. Vanishing SAF constructs this rational tensor, but it does not bound its two Galois shape coordinates. The pairwise representations above identify precisely what the surface formulation contributes: the integral state is genuine, while the positive conjugate degree in (806) remains to be charged to the radial incidence ideal. Thus SAF is valuable background and a check on the construction, but not an independent descent mechanism.

Route II.14.51 (Outcome of the joint-order test)**.** The scalar ordering fails exactly on A_3. The tensor (795) repairs the algebraic defect and retains all four dominance chambers. It does not by itself create a Subspace-Theorem surplus: equation (796) doubles both the useful outer evaluation degree and the ambient height, while unit orbits occupy the critical slope. A successful use of this tensor must therefore couple its fixed evaluation (797) to the negative Hodge metric at the first global norm record. This recovers the same stopped transport problem in finite rank, but eliminates the false hope that a single entry norm already solves it.

There is also a sharp limitation. Primitive parabolic generators canonically fix the scale at both cusp endpoints. If one endpoint is a general K-rational point $[P]$, the rank-one identity with N_P remains true, but $N_{uP} = u^2 N_P$ and no primitive parabolic removes the unit. The two anchors

nevertheless give the upper bound (792); the earlier claim that no one-rational-endpoint control existed was too strong. What remains absent is an exact multiplicative compression coordinate for $pc - qa$: the two anchor bounds do not retain cancellation in this linear combination, nor do they stop successive radial norm payments from compounding. That is the narrower global modular-symbol/Arakelov problem.

Corollary II.14.52 (The return continuant is the Hodge-compression divisor)**.** *Let $P_b = G_{a,b}P_a$ be Rosen prefix matrices and set*

$$\xi_a = P_a^{-1}0, \qquad \xi_b = P_b^{-1}0.$$

Then

$$A(\xi_a, \xi_b) = \lambda^2 (G_{a,b})_{12}^2, \qquad \mathrm{N}\, A(\xi_a, \xi_b) = \mathrm{N}((G_{a,b})_{12})^2. \tag{807}$$

Consequently every prime ideal which occurs in both r_a and r_b meets the finite divisor of the Hodge-compression covariant for the intervening modular symbol.

Proof. A primitive vector for ξ_j is $w_j = P_j^{-1}e_2$. Since P_b has determinant one,

$$\det(w_a, w_b) = \det(P_b P_a^{-1} e_2, e_2) = \det(G_{a,b}e_2, e_2) = (G_{a,b})_{12}.$$

Now apply (786). The final statement is exactly the return divisibility (II.9.15). □

Corollary II.14.53 (Quantitative Hodge bound for recycled primes)**.** *Let*

$$\mathfrak{R}_{a,b} = \prod_{\mathfrak{p} \mid r_a,\, \mathfrak{p} \mid r_b} \mathfrak{p}$$

be the reduced ideal of primes recycled between two occurrences, and let $T_{a,b}$ be the thick time of the intervening modular symbol. Then

$$\log \mathrm{N}\, \mathfrak{R}_{a,b} \le \log |\, \mathrm{N}(G_{a,b})_{12}| \le 3 \log |\sigma_0(\lambda (G_{a,b})_{12})| - \frac{\kappa_1 + \kappa_2}{2} T_{a,b}. \tag{808}$$

Proof. The first inequality is Proposition II.9.15. Substitute (807) into (788) and divide by two. □

Experiment II.14.54 (Exact compression/continuant audit)**.** The program `hodge_compression_continuant_test` generates random exact Rosen words, chooses two prefix times, constructs the two primitive parabolic generators, and independently verifies

$$-\operatorname{Tr}((g_{\xi_a} - I)(g_{\xi_b} - I)) = \lambda^2 \det(P_a^{-1}e_2, P_b^{-1}e_2)^2 = \lambda^2 (G_{a,b})_{12}^2.$$

All 500 tests with word length 20 and digits of absolute value at most 50 passed exactly.

Route II.14.55 (This is the non-gauge Hodge–S-unit bridge)**.** Unlike the nearest-wall scalar twist, the quantity $A(\xi, \eta)$ is an evaluation-dependent integral covariant. Formula (788) is therefore not canceled by a GIT quotient metric. It is also local in time: for a modular symbol $\gamma_1 * \cdots * \gamma_s$, each r_i is multiplicative, so the logarithmic compression and norm slack add over the pieces and may be discarded after a piece has been traversed.

The identity also pinpoints the remaining extension. For a general $x = [p : q] \in \mathbf{P}^1(K)$, the rank-one nilpotent

$$N_{p,q} = \begin{pmatrix} -pq & p^2 \\ -q^2 & pq \end{pmatrix}$$

satisfies

$$-\operatorname{Tr}(N_{p,q}N_{a,c}) = (pc - qa)^2.$$

Thus its pairing with a cusp vector is the square of the exact approximation determinant, and the product over the three embeddings is its algebraic norm squared. But $N_{up,uq} = u^2 N_{p,q}$: away from a cusp there is no primitive parabolic generator to remove the rank-two unit ambiguity. Proposition II.14.56 is exactly this failure in Hodge language.

The S-unit theorem now has a precise job. It must make the choice of scale for $N_{p,q}$ finite on any tail with fixed prime support; this is already achieved dynamically by Theorem II.9.14. Proposition II.14.45 proves the required one-algebraic- endpoint analogue of (788), with its norm term assigned to the current Rosen remainder. Thus the remaining step is no longer to construct the local bridge: it is to prove that its successive local norm payments cannot compound indefinitely in the balanced many-block regime. Bounded segment norms close by the S-unit theorem; dominant unbounded blocks close by Theorems II.14.12 and II.14.14. The residual noncompounding statement is strictly weaker than a global bound on the cumulative radical, which Experiment II.14.71 disproves.

Proposition II.14.56 (Unit-invariance obstruction to a bare nearest-wall height)**.** *Let K be totally real cubic. For real weights $w = (w_0, w_1, w_2)$, consider a monomial in the three Hodge factors*

$$N_w(C) = \prod_{i=0}^{2} \|C_i\|^{w_i}.$$

If $N_w(uC) = N_w(C)$ for every unit $u \in \mathcal{O}_K^\times$ and every C, then $w_0 = w_1 = w_2$. Consequently neither a pairwise Hodge product nor a nearest-wall signed product descends by itself to a height on $\mathbf{P}^1(K)$.

Proof. The logarithmic change under multiplication by u is

$$\sum_i w_i \log|\sigma_i(u)|.$$

The logarithms of units form a full lattice in the plane $x_0 + x_1 + x_2 = 0$. A linear functional vanishing on this lattice vanishes on the plane, so its coefficient vector is proportional to $(1, 1, 1)$. □

Proposition II.14.57 (Finite-preperiod blindness)**.** *Every nonempty finite Rosen cylinder contains a cusp point. Hence every finite admissible digit block is the initial block of an exact terminating orbit in K.*

Proof. On a cylinder of length k, the map T^k is one fixed element of G_7 and maps the cylinder interior homeomorphically onto a nonempty open interval J in the fundamental interval. The parabolic fixed points $G_7\infty$ of a cofinite Fuchsian group are dense in its boundary. Choose a cusp $y \in J$ away from branch endpoints and pull it back by the cylinder branch. Its preimage lies in K, has the prescribed first k digits, and terminates. □

Route II.14.58 (Consequence for local certificates)**.** No finite list of forbidden blocks can prove the classification, because there are no locally forbidden blocks: every finite admissible block also occurs on a closing K-orbit. A finite-state fusion certificate can establish an inequality, but the use of that inequality must compare an arbitrarily late state with the initial state or invoke an asymptotic Diophantine theorem. In particular, a contradiction attached to one large birth or one finite Hodge window is structurally impossible. Long closing preperiods realize precisely the local behavior such a contradiction would try to exclude.

Route II.14.59 (What the Hodge–S-unit synthesis must actually prove)**.** Proposition II.14.38 removes the claimed need for tensor induction, but Proposition II.14.56 prevents a free replacement of the ordinary determinant metric by the nearest-wall metric. Proposition II.14.39 gives a second, independent obstruction: a line twist pulled back from the metric parameter merely shifts the height of the GIT quotient by the same amount. Consequently a "compensating line" is useful only if it is represented by an *evaluation-dependent rational covariant*, whose horizontal divisor meets the varying integral Rosen state. A fixed Hermitian renormalization, even one canonically written with Ψ_i, is gauge.

This gives a sharper synthesis target. Construct a G-equivariant rational covariant $\mathscr{R}(g;v)$ on the integral Asai state space, not a constant metric twist, such that its associated metrized line satisfies

$$\widehat{\deg}(\det J_{d,r} \otimes \mathscr{R}(g;v)^{kd/2}) = kd\delta + O(k\log(d+2)), \tag{809}$$

and whose finite divisor is the reduced boundary intersection of the actual remainder state. At a single return the universal covariant

$$\langle \Phi(v), J\Phi(gv)\rangle = \mathrm{N}\, Q_g(v)$$

is a natural candidate, but its divisor counts all prime factors of the fixed-point evaluation and does not remember whether a prime is new. Time truncation is time ordered. The canonical finite object is therefore the cumulative reduced boundary ideal

$$\mathfrak{B}_N = \prod_{\substack{\mathfrak{p}:\ \mathfrak{p}\mid r_j \\ \text{for some } -1\le j\le N}} \mathfrak{p}. \tag{810}$$

Its increment from $N-1$ to N is exactly the reduced time-truncated-birth divisor. A genuine bridge must therefore be a pathwise/truncated-counting identity, or must enlarge the algebraic state by the accumulated boundary divisor. It cannot be a line on the one-return Grassmannian alone.

If $\mathrm{N}\,\mathfrak{B}_N$ is bounded, its support stabilizes and the fixed-support S-unit theorem closes the orbit. If it is unbounded, the desired relative-height formula must charge this truncated counting function against the positive Hodge outer line while retaining the parabolic equality divisor. This is the exact boundary-to-finite-birth bridge: it is an evaluation-dependent, time-ordered compensated Plücker identity, not tensor semistability or a change of metric.

Lemma II.14.60 (Canonical normalization of two consecutive S-unit relations)**.** *At a nonterminal Rosen transition put*

$$z_j = \frac{r_{j-1}}{\lambda r_j}, \qquad w_j = \frac{r_{j+1}}{\lambda r_j}, \qquad u_j = -\frac{z_j}{n_j}, \qquad v_j = -\frac{w_j}{n_j}.$$

Then

$$u_j + v_j = 1, \qquad n_j n_{j+1} v_j u_{j+1} = \lambda^{-2}. \tag{811}$$

The normalization is intrinsic to the projective four-term point

$$[z_j : w_j : -\tau z_j : -\tau w_j] \in \{X_1 + X_2 + X_3 + X_4 = 0\} \simeq \mathbf{P}^2 :$$

it is the affine chart in which $X_1 + X_2 = 1$. In particular it removes the rational projective scalar without choosing an archimedean metric.

Proof. The Rosen recurrence divided by λr_j gives $z_j + w_j = -n_j$, hence $u_j + v_j = 1$. Consecutive remainders also give

$$z_{j+1} = \frac{r_j}{\lambda r_{j+1}} = \frac{1}{\lambda^2 w_j}.$$

Substitution of $z_{j+1} = -n_{j+1}u_{j+1}$ and $w_j = -n_j v_j$ proves the second identity. Finally $X_1 + X_2 = -n_j \neq 0$, so division by this linear form is precisely the displayed canonical projective normalization. □

Proposition II.14.61 (The normalized state is a decorated Plücker cross ratio)**.** *Let* $P_j v = (r_j, r_{j-1})^t$*, and put*

$$w_j = P_j^{-1} e_2.$$

Then

$$r_j = \det(v, w_j), \qquad \det(w_{j-1}, w_j) = \det(w_j, w_{j+1}) = -1, \qquad \det(w_{j-1}, w_{j+1}) = n_j\lambda. \tag{812}$$

Moreover the two normalized coordinates are the Plücker cross ratios

$$u_j = \frac{\det(v, w_{j-1})\det(w_j, w_{j+1})}{\det(v, w_j)\det(w_{j-1}, w_{j+1})}, \tag{813}$$

$$v_j = \frac{\det(v, w_{j+1})\det(w_{j-1}, w_j)}{\det(v, w_j)\det(w_{j-1}, w_{j+1})}. \tag{814}$$

Thus $u_j + v_j = 1$ *is exactly the Plücker relation for the four decorated lines* $(v, w_{j-1}, w_j, w_{j+1})$*. The Hodge compression of the two nonadjacent cusp decorations is*

$$A(w_{j-1}, w_{j+1}) = n_j^2\lambda^4, \qquad |\,\mathrm{N}\, A(w_{j-1}, w_{j+1})| = |n_j|^6. \tag{815}$$

Proof. Since $\det P_j = 1$,

$$\det(v, w_j) = \det(P_j v, e_2) = r_j.$$

The branch relation $P_{j+1} = B(n_j)P_j$, with

$$B(n) = \begin{pmatrix} -n\lambda & -1 \\ 1 & 0 \end{pmatrix},$$

gives, in the P_j-frame,

$$w_{j-1} = -e_1, \qquad w_j = e_2, \qquad w_{j+1} = e_1 - n_j\lambda e_2.$$

This proves (812). Substitution in (813) and (814) gives

$$u_j = -\frac{r_{j-1}}{n_j\lambda r_j}, \qquad v_j = -\frac{r_{j+1}}{n_j\lambda r_j}.$$

The four-vector Plücker relation is therefore precisely $u_j + v_j = 1$. Finally apply (786) to $\det(w_{j-1}, w_{j+1}) = n_j\lambda$; since $|\,\mathrm{N}\,\lambda| = 1$, this gives (815). □

Remark II.14.62 (Cluster interpretation and its limitation). The normalized chain is the X-coordinate shadow of the decorated flag, or cluster A-coordinate, configuration carried by the primitive cusp vectors w_j. This explains both the exact scalar cancellation and the local different (824). Passing from the cross ratios to the decorated Plücker coordinates is birational after retaining the one decoration scale, so it cannot by itself add a positive log-canonical class. The required outer line must use the Hodge metric or modular-symbol transport on these decorations; it is not created by another algebraic compactification of the same cluster configuration.

Proposition II.14.63 (Long-range Plücker recurrence)**.** *Put* $C_{p,q} = \det(w_p, w_q) = (G_{p,q})_{12}$. *For every* $a < j < b$,

$$r_a C_{j,b} - r_j C_{a,b} + r_b C_{a,j} = 0. \tag{816}$$

Thus all local normalized S*-unit relations are adjacent shadows of one type-A Plücker frieze, and each coefficient* $C_{p,q}$ *is the square root, up to* λ*, of the Hodge-compression covariant for the modular symbol from* p *to* q.

Proof. By Proposition II.14.61, $r_i = \det(v, w_i)$. The four-vector Plücker relation for (v, w_a, w_j, w_b) is exactly (816). The last statement is Corollary II.14.52. □

Route II.14.64 (Plain *abc* on long-range Plücker relations is critical)**.** Let $\mathfrak{p} \mid r_j$ be time-truncated relative to endpoints a, b, so $\mathfrak{p} \nmid r_a r_b$. Reducing (816) gives only

$$r_a C_{j,b} \equiv -r_b C_{a,j} \pmod{\mathfrak{p}}. \tag{817}$$

It does not force $\mathfrak{p} \mid C_{a,j} C_{j,b}$. In the three-term equation, a simple time-truncated divisor of r_j therefore occurs once in the middle term and once in the reduced radical. The difficult record norms are empirically squarefree-dominated, so an ordinary number-field *abc* estimate on (816) has no coefficient to spare. Taking many triples does not change this local calculation: it repeats the same Grassmannian boundary divisor.

The useful content is instead the identification of the missing inequality. One needs a moving-target truncated Plücker estimate in which the compression defects of $C_{a,j}$ and $C_{j,b}$ are placed on the left before the prime of r_j is truncated. Schematically, the desired local term has the form

$$h(r_j C_{a,b}) - N^{(1)}_{\text{time-truncated}}(r_j) - \Delta_H(C_{a,j}) - \Delta_H(C_{j,b}),$$

with a positive lower bound on thick windows. This is not supplied by the fixed-target S-unit theorem or by bare *abc*; it is the exact arithmetic–Hodge bridge still to construct.

Route II.14.65 (Squaring the Plücker relation loses its apparent gain on a Veronese conic)**.** Write $A = r_a C_{j,b}$, $B = r_b C_{a,j}$, and $C = r_j C_{a,b} = A + B$. Squaring gives the four-term relation

$$C^2 - A^2 - B^2 - 2AB = 0. \tag{818}$$

At first sight a newly appearing divisor of r_j occurs to order two in C^2 but only once in the reduced radical. This is not a genuine surplus. The four-term point is constrained to the fixed conic

$$\phi : \mathbf{P}^1 \longrightarrow \mathbf{P}^3, \qquad [A : B] \longmapsto [(A+B)^2 : -A^2 : -B^2 : -2AB].$$

The reduced inverse image of the four coordinate hyperplanes is the three-point divisor

$$D_{\text{red}} = \{A = 0\} + \{B = 0\} + \{A + B = 0\}.$$

Consequently

$$K_{\mathbf{P}^1} + D_{\text{red}} = \mathcal{O}_{\mathbf{P}^1}(1), \qquad \phi^* \mathcal{O}(1) = \mathcal{O}_{\mathbf{P}^1}(2).$$

The extra apparent copy is removed by adjunction/the different. Generic coefficient-one Vojta for the ambient four-line arrangement is allowed to put this conic in its exceptional set; restricting Vojta to the conic recovers ordinary three-point *abc* at coefficient one. The uniform four-term n-conjecture with coefficient three is weaker still.

Therefore the evaluation-dependent identity $Q_{w_j}(v) = \lambda r_j^2$ is valuable for recognizing the parabolic Hodge covariant, but its nonreduced square cannot by itself pay for time-truncated prime births. This is the same normal-bundle cancellation as in Route II.14.67, expressed on a single Veronese fiber.

Route II.14.66 (Omitting the central Plücker boundary leaves only a digit-boundary class)**.** For the adjacent configuration put

$$[x : y : z : c : d] = [r_{j-1} : r_j : r_{j+1} : -1 : n_j\lambda].$$

After identifying the two adjacent cusp determinants, its projective closure is the quadric cone

$$Y : \quad c(x+z) = yd \subset \mathbf{P}^4. \tag{819}$$

It has $K_Y = -3H$. Formally the four coordinate hyperplanes $x = 0, z = 0, c = 0, d = 0$, which omit the time-truncated coordinate y, have total Cartier class $4H$. This suggests $K_Y + D = H$, but the boundary is not reduced: the sections $c = 0$ and $d = 0$ share the plane

$$A = \{c = d = 0\} \subset Y.$$

Writing the other two components at infinity as $B = \{d = 0, x + z = 0\}$ and $C = \{c = 0, y = 0\}$, one has

$$(c = 0) = A + C, \qquad (d = 0) = A + B.$$

Therefore

$$D_{\mathrm{red}} = D_{\mathrm{Cartier}} - A, \qquad K_Y + D_{\mathrm{red}} = H - A = B. \tag{820}$$

The last equality uses $H = A + B$. It shows that the remaining class is the fixed digit-boundary component, not the ample Plücker class. Equivalently, on the actual chart $c = 1$,

$$Y \setminus D_{\mathrm{red}} \simeq \mathbf{G}_m^3, \qquad (x, z, d) \longmapsto [x : (x + z)/d : z : 1 : d].$$

Thus the log geometry is critical once more. Standard truncated Vojta for the reduced pair can account for the primes of the digit d, but it cannot omit $y = r_j$ while retaining the full height h_H.

The program `rosen_plucker_omitted_boundary_tests.py` deliberately tests the stronger, invalid inequality as a numerical negative control. It computes

$$h_H([r_{j-1} : r_j : r_{j+1} : n_j\lambda : 1]) - \frac{1}{3} \sum_{\mathfrak{p} \mid r_{j-1} r_{j+1} n_j} \log \mathrm{N}\,\mathfrak{p}$$

with exact PARI prime-ideal factorization. At the selected large norm records for unit-ray exponents (m=1,. . . ,8), the values are

$$-6.129, -8.620, -7.918, -3.222, -10.601, -5.211, -10.146, -13.599.$$

Near the eventual cusp or repeated state, the central remainder becomes a unit while the omitted radical collapses, and the same false surplus is (31.1,45.7,55.5,54.3,56.0,65.7,81.8,73.4). Those late values are explained by the fixed-support S-unit theorem; they do not validate an (H)-height Vojta statement, because the genuine log-canonical line is the boundary-supported class in (820).

Route II.14.67 (The chained S-unit compactification is exactly log-critical)**.** Lemma II.14.60 gives a valid pathwise algebraic replacement for the tensor-induced diagonal evaluation proposed in `curt11.tex`. A normalized transition is a point $u_j \in \mathbf{P}^1 \setminus \{0, 1, \infty\}$, with $v_j = 1 - u_j$. A chain of m transitions lies in $X = (\mathbf{P}^1)^m$, and the consecutive-state equations are moving divisors

$$G_j : \quad n_j n_{j+1}(1 - u_j)u_{j+1} = \lambda^{-2}, \qquad [G_j] = H_j + H_{j+1}.$$

Let D be the union of the three coordinate boundary points in every factor. Since $K_{\mathbf{P}^1}+\{0,1,\infty\}=H$,

$$K_X + D = \sum_{j=0}^{m-1} H_j.$$

If $Y = G_0 \cap \cdots \cap G_{m-2}$ is a transverse chain closure, adjunction gives

$$K_Y + D|_Y = \left(2H_0 + 3\sum_{j=1}^{m-2} H_j + 2H_{m-1} \right) |_Y. \tag{821}$$

This coefficient three is not yet a gain. Let $\mathfrak{p}$ be a newly appearing divisor of r_j which does not return at r_{j+2}. Adjacent primitivity and the recurrence give the linked reductions

$$u_{j-1} = 1, \qquad u_j = \infty, \qquad u_{j+1} = 0 \pmod{\mathfrak{p}}. \tag{822}$$

Indeed v_{j-1} contains r_j in its numerator, u_j and v_j have r_j in their common denominator, and u_{j+1} contains r_j in its numerator. If the last reduction fails, then $\mathfrak{p} \mid r_{j+2}$, so Proposition II.9.15 classifies the prime as an immediate return rather than a time-truncated nonreturning incidence.

There are two equivalent ways to see why coefficient three gives no gain. Applying separate componentwise *abc* inequalities charges the prime three times, exactly matching $3H_j$. Alternatively, use the reduced union boundary and count the linked prime once. In that model the left side also loses two copies of the same divisor.

Indeed, near a time-truncated triple with the three adjacent digits $\mathfrak{p}$-adic units, put

$$a = 1 - u_{j-1}, \qquad b = 1/u_j, \qquad c = u_{j+1}.$$

The two glue equations become

$$b = \lambda^2 n_{j-1} n_j a, \qquad b = \lambda^2 n_j n_{j+1}(b-1)c. \tag{823}$$

Since $b-1$ is a unit at the boundary, the restrictions of $a=0, b=0, c=0$ to the chain are the same prime divisor E_j. The ambient boundary restricts as $3E_j$, whereas the reduced boundary of the chain contains E_j once. Ordinary complete- intersection adjunction therefore gives the local different formula

$$(K_X + D_X + \textstyle\sum G_i)|_Y = K_Y + D_Y^{\mathrm{red}} + 2E_j. \tag{824}$$

Thus the change from a componentwise count to the reduced-union count subtracts $2E_j$ from both sides. If $\mathfrak{p} \mid n_j$, a toroidal resolution adds the corresponding digit-boundary component; its valuation accounts for the extra strict transform and does not restore a surplus. This is the precise log-geometric reason that iterating a size-only four-term *abc* or Vojta inequality remains saturated. Chaining the equations removes the projective scalar, but it does not by itself pay for time-truncated births.

It also identifies the precise useful synthesis. Recycled primes are divisors of the return continuant and are bounded by the genuine Hodge compression in Corollary II.14.53. The remaining boundary is the time-truncated stratum (822). The sharp relative fusion gap supplies an additional positive archimedean Hodge line proportional to the number of induced crossings. A closing theorem must put that extra line on the left of the log-critical formula; ordinary boundary height alone has coefficient exactly zero after subtraction.

Experiment II.14.68 (Exact *abc* radical of the normalized S-unit state)**.** The program `rosen_normalized_sunit_a` checks both identities in Lemma II.14.60 exactly at every transition of the adversarial unit rays. PARI/GP factors the fractional principal ideals (u_j) and $(1-u_j)$, and computes the actual normalized number-field radical

$$N_j^{(1)} = \frac{1}{3} \sum_{\mathfrak{p}:\, v_\mathfrak{p}(u_j)\neq 0 \text{ or } v_\mathfrak{p}(1-u_j)\neq 0} \log \mathrm{N}\,\mathfrak{p}.$$

At the transition where $h(u_j)$ is largest, the output is

m	j	$h(u_j)$	$N_j^{(1)}$	$h(u_j) - N_j^{(1)}$
1	0	9.293	24.235	−14.941
2	1	11.326	28.370	−17.044
3	0	12.205	29.681	−17.476
4	11	16.947	42.185	−25.238
5	2	14.736	41.003	−26.267
6	12	17.430	44.132	−26.702
7	4	18.296	43.237	−24.941
8	4	21.685	46.558	−24.872

Thus even after the canonical scalar normalization, the ordinary *abc* boundary is much larger than the point height at the high states. At a second, low-height transition selected to maximize the cheaper rational-prime proxy, the exact values of $h - N^{(1)}$ are

$$0.251,\ 0.026,\ 0.970,\ 0.970,\ 0.662,\ 0.162,\ -0.368,\ 0.039.$$

This is consistent with bounded S-unit exceptional configurations, but supplies no growing surplus. The calculation rules out a direct application of ordinary *abc* to $u_j + (1-u_j) = 1$: the usable positive term must be the extra outer Hodge/fusion line identified in Route II.14.67, not the standard point height.

Experiment II.14.69 (The formal chained union radical measures the different, not a Vojta gain)**.** For three consecutive transitions, formal ambient adjunction, now also retaining the digit variables, gives the class

$$2H_{u_0} + 3H_{u_1} + 2H_{u_2} + H_{n_0} + 2H_{n_1} + H_{n_2}.$$

The program `rosen_chained_log_voyta_tests.py` computes this formal height and factors, in PARI/GP, the union of all prime ideals occurring in $u_i, 1-u_i, n_i$. At the dominant record windows on the adversarial unit rays it gives

m	h_{formal}	$N^{(1)}_{\text{union}}$	$h_{\text{formal}} - N^{(1)}_{\text{union}}$
1	62.235	38.620	23.615
2	77.099	48.070	29.029
3	82.939	48.756	34.183
4	113.473	69.196	44.277
5	101.165	66.691	34.474
6	116.683	72.414	44.270
7	123.820	80.454	43.365
8	132.903	77.314	55.589.

These positive numbers initially looked like a reduced-union Vojta surplus. Equations (823) and (824) show exactly why that interpretation is wrong: the program intentionally uses the formal ambient class and omits the different $-2E_j$. The table therefore measures the size of the correction that a valid reduced log model must subtract. It is a useful negative control and confirms, rather than contradicts, the log-critical calculation.

Remark II.14.70 (Precise comparison with `curt11.tex`). The two approaches have the same intended surplus and different algebraic realizations. The nearest-wall pencil in `curt11.tex` correctly identifies δ as a Hodge slope, but its diagonal tensor evaluation is nonlinear and its scalar metric compensation is canceled by the quotient minimum. The S-unit approach instead uses the exact linear remainder recurrence and the canonical glue (811); no tensor linearization is required. Their viable synthesis is therefore pathwise: the nearest-wall slope is the extra archimedean term, the normalized S-unit chain is the integral boundary model, and the compression/continuant covariant is its evaluation-dependent finite divisor. This is different from the draft's global tensor jet, but it realizes the same Hodge budget without the metric-gauge cancellation.

Experiment II.14.71 (The cumulative birth radical has no sextic initial budget)**.** The exact program `rosen_cumulative_birth_tests.py` factors the rational norms of *every* remainder through termination or repetition on the adversarial unit perturbations. It records the product of all distinct rational primes encountered. This is a lower bound for the norm of $\mathfrak{B}_N$: an inert rational prime contributes an ideal of norm p^3, while a split rational prime contributes at least one ideal of norm p. The output is

m	outcome	$\#\{r_j\}$	$h(x)$	$\log \mathrm{rad}_{\mathbf{Q}}$	$\log \mathrm{rad}_{\mathbf{Q}} - 6h(x)$
1	cusp	93	7.877	788.048	740.784
2	special	119	8.996	1454.350	1400.376
3	cusp	144	10.129	1935.254	1874.478
4	cusp	146	11.255	2389.604	2322.075
5	special	150	12.385	2424.444	2350.132
6	cusp	165	13.513	3723.882	3642.805
7	cusp	206	14.642	4326.193	4238.341
8	cusp	182	15.770	3538.841	3444.220

Thus the tempting strengthening $\log \mathrm{N}\,\mathfrak{B}_N \le 6h(x) + O(1)$ is decisively false, even on orbits which do close. The largest single-remainder radicals in the same rows range only from 24.60 to 55.82; the enormous cumulative value is caused by many successive, disposable births. A successful Hodge–counting formula must therefore be dissipative in time: it may charge and then forget finite births. It cannot store every previous boundary intersection in a single initial-height budget. What must be excluded is infinitely many births along one nonclosing path, not a large finite number along a long preperiod.

Experiment II.14.72 (Record births do not telescope either)**.** One might try to repair Experiment II.14.71 by forgetting every birth away from a new norm record and charging a record only by the increase of the running maximum. The exact program `rosen_record_birth_budget_tests.py` disproves this version as well. It counts a rational prime only at its first occurrence and only if that

occurrence is a new remainder-norm record. On the same unit-ray samples it gives

m	$\log E_{\max}$	$\sum_{\text{record}} \log \operatorname{rad}_{\text{time-truncated}}$	excess	worst increment	time-truncated charge there
1	24.599	65.510	40.911	3.914	24.599
2	31.252	118.423	87.171	2.214	28.067
3	32.203	78.838	46.635	5.792	32.203
4	43.739	263.447	219.708	5.622	43.739
5	40.025	132.215	92.190	1.237	40.025
6	45.029	183.425	138.397	1.876	45.029
7	50.500	297.443	246.943	1.771	50.500
8	55.818	269.048	213.230	0.580	55.818

Thus a later squarefree record can replace almost its entire support by previously unused factors even when the record itself increases only slightly. Neither the cumulative radical nor the sum of record increments is the required energy. A viable induction must apply the Hodge/parabolic alternative at the individual record and then discard its finite divisor; it cannot pay that divisor from a telescoping maximum-norm account.

Experiment II.14.73 (One-step pairwise Hodge drops do not pay births)**.** The segmentwise Hodge inequality suggests testing the two elementary archimedean potentials

$$H_i(j) = \log \max\{|r_{j,0}|, |r_{j-1,0}|\} + \log \max\{|r_{j,i}|, |r_{j-1,i}|\}, \qquad i = 1, 2.$$

The exact program `rosen_hodge_birth_potential_tests.py` factors every remainder and compares the rational new-divisor charge at time j with the positive one-step drops of H_1, H_2. On the same adversarial unit ray it gives

m	time-truncated charge	$(H_1(j-1) - H_1(j))_+$	$(H_2(j-1) - H_2(j))_+$	unpaid excess
1	24.599	1.369	0.371	22.859
2	30.787	0.172	0	30.615
3	32.203	1.386	0.362	30.455
4	43.739	1.397	0.367	41.975
5	38.486	0.385	0.144	37.957
6	43.897	0.204	0	43.693
7	50.500	1.449	0.397	48.654
8	55.818	1.029	0.115	54.674

Thus even the sum of both local Hodge drops is smaller than a single time-truncated divisor by an arbitrarily large-looking margin in the sample. This does not contradict Proposition II.14.45, whose decrease is measured over thick segments and includes a current norm term. It rules out a simpler telescoping proof: the norm term cannot be replaced by the raw one-step drop of either pairwise archimedean potential. The alternative must be applied at the high-current-norm state itself, exactly as predicted by the parabolic full-defect branch of `curt11.tex`.

II.15 Experiments and proof attempts still in progress

The following items are retained as explicit obligations rather than hidden inside the notation.

1. Complete the finite relative-fusion automaton of Experiment II.2.21. Ten parabolic symbols peel exactly and only 22 residual weak-letter transitions remain. The required statement

is now a cyclic complementary-trace lemma for their 30 bad ordered pairs, with the sharp one-cusp subsemigroups discharged by Corollary II.2.43. All exceptional cycles through length three now pass exactly. An alternative is to prove the infinitesimal inequality (671), certified for every word through length seven by Experiment II.2.24.

2. Control the balanced many-block regime left by Theorems II.14.12 and II.14.14. The required statement is a noncompounding lemma: along an infinite nonclosing Rosen path, (785) and the two consecutive four-term S-unit equations cannot repeatedly return to the coefficient-balanced cubic boundary.

3. Make the preceding item quantitative by decomposing every long return at its largest parabolic digits. Either a finite subtree satisfies the small-coefficient hypotheses of Theorem II.14.14, or exhibit a definite amount of thick time and feed it into (785). The adversarial unit ray is the sharp calibration.

4. If arithmetic GIT is retained, construct an *evaluation-dependent* relative covariant, not a tensor-induced linear map or scalar metric twist. Its fundamental formula must include the shift of the quotient minimum from Proposition II.14.39 and must reproduce the standard parabolic height (780).

5. Build the parabolic equality case into that relative theorem. For Q_N^r, Proposition II.14.8 forces a full-size boundary defect. The equality divisor must reduce to the cusp incidences already classified for every fixed number of blocks by Theorem II.14.11.

6. Prove the weaker but sufficient sextic norm-excursion estimate (753). Work with the integral cubic $P_j(X) = \mathrm{N}(r_{j-1} + X\lambda r_j)$, its digit divisibility (739), and record-to-record norm growth. The unit ray must be included and is asymptotically sharp.

7. A proof of the stronger quadratic height estimate $h(T^n x) \le 2h(x) + B_7$ remains valuable but is no longer logically necessary: the finite-support S-unit theorem controls the unit directions discarded by the norm energy.

8. Construct the induced polyhedral cover in the two logarithmic relative-unit coordinates. On every cell, certify either the quadratic barrier directly or passage to a cell of lower record height; exact cycles are then checked by the fixed-point form.

9. Use ordinary or symmetric Roth only on the uniformly supercritical branch. Proposition II.14.39 and II.14.41 show that a nearest-wall metric renormalization does not repair the critical tail.

10. Do not use the $7 + 2$ orbit decomposition modulo 2 as a type classifier: 671 proves that both real outcomes occur over the residual cusp orbit. A congruence calculation can re-enter only if it supplies a new archimedean inequality.

11. Use the canonical normalization (811), not a projective tensor evaluation, for the arithmetic chain. Route II.14.67 and Experiment II.14.68 show that its ordinary *abc* boundary is exactly critical. The missing statement must put the relative Hodge/fusion line on the left while assigning immediate returns to the compression/continuant divisor. Equation (824) rules out a reduced-union shortcut: the apparent saving of two boundary counts is canceled by the local different $2E_j$. Any Hodge line must therefore come from a genuine evaluation-dependent lift, not another birational boundary model.

II.16 The global incidence reformulation

The prime-ideal experiments above were useful diagnostics, but they do not supply the missing bound. A statement that a large remainder has a "time-truncated" prime divisor is local and time dependent. To turn it into a uniform global estimate one would still have to prove that losses cannot migrate among different finite places. That synchronization is exactly what an Arakelov height and the product formula are designed to do. Accordingly, new-divisor control is no longer treated as a proof route. It records the finite-place shadow of the following single incidence calculation.

Fix the primitive initial pair $P = (p, q)$. A Rosen prefix row $w = (a, b)$ is primitive and produces the remainder $r = ap + bq$. At the three real embeddings define

$$H_i(P) = \log \max(|\sigma_i p|, |\sigma_i q|), \qquad H_i(w) = \log \max(|\sigma_i a|, |\sigma_i b|)$$

and

$$\lambda_i(w, P) = H_i(P) + H_i(w) - \log |\sigma_i(ap + bq)|.$$

The last expression is a local projective proximity: it is large exactly when the prefix row nearly annihilates the initial pair at that embedding. Summing gives the exact product-formula identity

$$\log |\,\mathrm{N}(ap + bq)| = 3h(P) + 3h(w) - \lambda_0(w, P) - \lambda_1(w, P) - \lambda_2(w, P). \tag{825}$$

At every finite place the local term is simply the valuation of $ap + bq$; their sum is the norm on the left. Thus (825) controls all primes at once and is unchanged by multiplication by a global unit.

At the distinguished embedding a prefix row is a Rosen convergent. The standard error estimate $|a\sigma_0(p/q) + b| \ll_7 |a|^{-1}$ gives

$$\lambda_0(w, P) \geq 2H_0(w) - C_R. \tag{826}$$

This is exactly the q^{-2} phenomenon discussed in the Roth route: it recognizes a continued-fraction convergent, but supplies no new outer-place information.

Put

$$\chi(w) = H_1(w) + H_2(w) - H_0(w), \qquad \Lambda_{\mathrm{out}} = \lambda_1 + \lambda_2$$

and

$$D_{\mathrm{out}}(w, P) = \Lambda_{\mathrm{out}} - \chi(w) + 3h(P).$$

The two outer prefix heights cancel from this expression:

$$D_{\mathrm{out}}(w, P) = H_0(w) + H_0(P) + 2H_1(P) + 2H_2(P) - \log |\sigma_1(r)\sigma_2(r)|. \tag{827}$$

So the target compares the product of the two outer remainders with one central prefix size and the initial adelic height. No factorization of the remainder occurs. Then the norm slack has the exact decomposition

$$\log |\,\mathrm{N}\, r| - 6h(P) = -D_{\mathrm{out}}(w, P) - \big(\lambda_0(w, P) - 2H_0(w)\big). \tag{828}$$

Consequently the full sextic norm barrier follows from the following stopping-time assertion:

Conjecture II.16.1 (Outer incidence at a norm record)**.** *There is a field constant C_{out} such that, whenever w produces the first remainder whose norm reaches a prescribed threshold,*

$$D_{\mathrm{out}}(w, P) \geq -C_{\mathrm{out}}.$$

This is not merely another name for a prime estimate. Its left side is a single adelic evaluation: the finite stopping condition chooses the section, and the two outer Hodge metrics measure its archimedean proximity. It is also not an application of ordinary Roth. Roth bounds how often an algebraic point can have *too much* proximity to rational targets; Conjecture II.16.1 asks for a *lower* bound on the outer proximity forced by a Rosen prefix at a finite norm record.

There is a dangerous but instructive weakening. Dropping Λ_{out} would ask for

$$\chi(w) \le 3h(P) + O_7(1). \tag{829}$$

This is false. For the exact balanced inverse orbit with $m = 20$, a strict norm record at step 11 has

$$\chi(w) - 3h(P) = 21.3480097446\ldots.$$

At that same record the two outer proximities are

$$\lambda_1 = 33.2589762417\ldots, \qquad \lambda_2 = 33.4021767898\ldots;$$

they more than cancel the false residual excess, and the actual norm slack is

$$\log|\mathrm{N}\, r| - 6h(P) = -45.5927539848\ldots.$$

At $m = 32$ the residual excess rises to $36.5427300858\ldots$, while the smallest outer-incidence defect on the orbit is still $68.0240705934\ldots$. This family is the clearest reason to retain the evaluation section rather than estimate a prefix height by itself.

Experiment II.16.2 (Global incidence audit)**.** The exact program `rosen_record_residual_prefix_tests.py` transports the prefix row along with the Rosen remainders. On 500 random initial inputs it found 1146 strict norm records. The maximum error in (825) was zero to the displayed precision. The smallest observed D_{out} was $9.5421529758\ldots$, and the smallest central slack $\lambda_0 - 2H_0(w)$ was $-0.1491141147\ldots$. The unit-ray tests through exponent 48 had $D_{\mathrm{out}} \ge 21.5504028477\ldots$; the balanced inverse tests through exponent 32 had $D_{\mathrm{out}} \ge 11.1835738167\ldots$. Every tested norm slack was negative in those three samples. Experiment II.9.55 records exact all-state exceptions to both apparent sign patterns and shows why only the coupled defect, up to an additive field constant, is a viable target.

The normalization $q = 1$ makes later records unnecessarily rare. A second random run therefore used 3000 primitive pairs with comparable field numerators and rational denominators at most 1000. It produced 6454 records and had

$$\min D_{\mathrm{out}} = 3.6086038572\ldots.$$

The exhaustive low-height command

```
python3 rosen_outer_incidence_dynamics.py  -grid-bound 5 -max-denominator 20
                                           -maximum 120 -records-only
```

tested 18168 distinct primitive inputs and 28029 records. Here 262 records had negative outer defect. The sharpest was

$$P = (-5\lambda + 4\alpha, 18), \qquad j = 74,$$
$$D_{\mathrm{out}} = -8.8869611755\ldots, \qquad \lambda_0 - 2H_0(w) = 0.3384586668\ldots,$$
$$\log|\mathrm{N}\, r| - 6h(P) = 8.5485025087\ldots.$$

Thus the conjecture must have an additive constant; neither $D_{\mathrm{out}} \geq 0$ nor the sextic bound with zero constant survives at records. The numerator has norm one. With that numerator fixed, a sweep over every $1 \leq N \leq 2000$ found a negative value only at $N = 18$; the next smallest outer defect was $0.0796630973\ldots$. This is evidence that the obstruction is a resonant stopped excursion rather than a ray on which the defect drifts to $-\infty$. Proving precisely that boundary non-drift is now the point of the Arakelov stopping inequality.

Route II.16.3 (Hodge proof of the outer-incidence inequality)**.** The prefix row w is a transported covector. At the i-th embedding, its evaluation on the initial line is exactly $\sigma_i(ap+bq)$, so $\lambda_i(w, P)$ is the logarithmic evaluation term occurring in McMullen's Hodge norm. The two pairwise contraction inequalities (781) should therefore be applied to this same covector and added only at the first norm record. The desired conclusion is

$$\lambda_1(w, P) + \lambda_2(w, P) \geq H_1(w) + H_2(w) - H_0(w) - 3h(P) - C_{\mathrm{out}}.$$

The central term $H_0(w)$ is supplied by the convergent estimate (826); the two outer terms must be retained as evaluation proximities until the product formula is used.

The deep step is to construct the two outer metrics on the Rosen return correspondence and prove that the stopped evaluation section has the displayed lower bound. The elementary steps are the incidence identity, primitivity of the prefix row, and the central exponent-two estimate. This distinction prevents a numerical check of (828) from being mistaken for a proof of the Hodge inequality.

There is now an exact diagnosis of the missing metric. The two established pairwise Hodge inequalities count the distinguished Hodge factor twice, once with each outer factor. Formula (827) allows only one central prefix copy. The full three-factor Hodge energy has the correct multiplicity, and along a central geodesic its derivative is bounded by

$$-1 + \|DF_1\|_{\mathrm{hyp}} + \|DF_2\|_{\mathrm{hyp}}.$$

The two individual strict contractions do not imply that their sum is less than one. More decisively, the zero-balanced hypergeometric expansion gives

$$\tau_k(t) = i(-\log t + c_k) + O(t \log t), \qquad c_k = -2\gamma - \psi(a_k) - \psi(b_k),$$

and therefore

$$\|D(\tau_k \circ \tau_1^{-1})\|_{\mathrm{hyp}} = \frac{-\log t + c_1}{-\log t + c_k} + O\big(t(\log t)^2\big), \qquad k = 2, 3.$$

The displayed three-factor derivative bound tends to 1, rather than to a negative number, at the cusp.

The local cusp model itself is not conjectural. For $N = \log T$, where T is the unipotent cusp monodromy, a strictly adapted metric is

$$h_{\mathrm{ad}}(x + iy) = e^{xN} e^{-\frac{1}{2}(\log y)Y} h_{\mathrm{ad}}(i)$$

for an $\mathfrak{sl}_2$-triple containing (Y, N), and it obeys

$$h_{\mathrm{ad}}(\tau + n) = e^{nN} h_{\mathrm{ad}}(\tau).$$

Schmid's SL_2-orbit theorem makes the Hodge metric uniformly comparable to this model in the cusp; see Filip [20, Appendix A]. Hence an arbitrary parabolic power contributes its transported

prefix through exact automorphy, with only a bounded comparison error. This proves the local normalization which the earlier informal cusp-correction phrase had concealed.

There is an exact product-formula calculation at this interface. It removes the parabolic exponent and every finite prime from the list of separate variables. Let $(u, v) \in \mathcal{O}_K^2$ be primitive, let

$$P_i = \max\{|\sigma_i(u)|, |\sigma_i(v)|\}, \qquad \eta_i = \log P_i - \log |\sigma_i(v)|,$$

and suppose

$$r = u + n\lambda v, \qquad |\sigma_0(r)| \le C_c |\sigma_0(v)|. \tag{830}$$

Then

$$|\,\mathrm{N}\, r| \le C_c' e^{3h([u:v])}\,(1 + \Theta(u, v))\,, \qquad \Theta(u, v) = e^{\eta_0 - \eta_1 - \eta_2} = \frac{P_0 |\sigma_1(v)\sigma_2(v)|}{|\sigma_0(v)| P_1 P_2}. \tag{831}$$

Indeed, (830) gives $|n| \ll P_0/|\sigma_0(v)|$. For $i = 1, 2$, it follows that

$$|\sigma_i(r)| \ll P_i + \frac{P_0}{|\sigma_0(v)|} |\sigma_i(v)|.$$

After multiplication by the central bound, the first three terms are $O(P_0 P_1 P_2)$, while the final term is exactly

$$P_0 P_1 P_2 \Theta(u, v).$$

Primitivity gives $P_0 P_1 P_2 = e^{3h([u:v])}$. Moreover the product formula and $|\,\mathrm{N}\, v| \ge 1$ give

$$\Theta(u, v) \le P_0 P_1 P_2 = e^{3h([u:v])}, \qquad \log |\,\mathrm{N}\, r| \le 6h([u : v]) + O_{C_c, K}(1). \tag{832}$$

For a Rosen step, take $u = r_{j-1}$, $v = r_j$, with the sign changed according to the recurrence. The next ratio belongs to the fixed Rosen interval, so $C_c = \lambda_0/2$. Formula (831) therefore treats every digit uniformly, however large it is and however its prime factors are distributed. The only loss left after the all-place summation is

$$\log \Theta = \eta_0 - \eta_1 - \eta_2. \tag{833}$$

This is one central evaluation deficiency minus the two outer deficiencies. It is the same one-central-copy defect found independently in the Hodge weights of the incidence component. Consequently the remaining cusp calculation is no longer an estimate for n. It is a stopped Hodge estimate showing that a large $\log \Theta$ is paid for by the preceding decrease of the two pairwise metrics.

The factor has a second, completely elementary description. Expanding its definition and using the product formula for v gives

$$3h([u : v]) + \log \Theta(u, v) = \log |\,\mathrm{N}\, v| + 2\eta_0. \tag{834}$$

Hence the transfer estimate is equivalently

$$|\,\mathrm{N}\, r| \ll_7 e^{3h([u:v])} + |\,\mathrm{N}\, v| e^{2\eta_0}. \tag{835}$$

For a deep Rosen cusp, e^{η_0} is comparable to the next digit. Thus the genuinely dangerous monomial is the previous remainder norm times the square of the cusp depth. It is not a contribution from a new prime. The square occurs because the central cancellation removes one power from a cubic norm and leaves the two outer powers.

Experiment II.16.4 (Stopped cusp-transfer audit)**.** The program `rosen_cusp_transfer_tests.py` evaluates every term in (831) immediately before a strict norm record. On 500 random primitive inputs with coefficient bound 10^5, it found 642 such transitions. The largest value of

$$\log|\mathrm{N}\, r| - 3h([u:v]) - \log(1+\Theta)$$

was $0.2023271483\ldots$, consistent with the proved uniform constant (the direct triangle-inequality proof gives the nonoptimized bound $1.1202685136\ldots$). The stronger stopped envelope

$$3h([u:v]) + \log(1+\Theta) - 6h(P)$$

was at most $-10.4855885164\ldots$ on this sample. On the exact low-height exception $P = (-5\lambda + 4\alpha, 18)$, its maximum is instead $10.8197793730\ldots$, at the transition producing the step-74 record. This confirms two points: an additive field constant is necessary, and the explicit Θ-envelope survives the known counterexample to the zero-constant norm bound.

An exhaustive coefficient grid of bound 5, rational denominators at most 20, and 120 steps produced 9862 pre-record transitions. Its largest transfer slack was $0.2261422075\ldots$; in particular the all-place lemma remains numerically sharp and stable on the full grid which discovered the low-height exception.

The balanced inverse-unit family was also tested through exponent 16. Its false unit reservoir grows linearly, but the stopped Θ-envelope remains below the initial sextic budget on every record-producing transition tested. This separates the new scalar from the reservoir already disproved above.

The exact sufficient boundary statement can now be made scalar. Assume that there is $B_\Theta = O_7(1)$ such that every transition producing a strict new norm record satisfies

$$3h([r_{j-1}:r_j]) + \log(1+\Theta(r_{j-1},r_j)) \le 6h([r_0:r_{-1}]) + B_\Theta. \tag{836}$$

Then (831) bounds the new remainder norm by the same initial sextic budget plus a field constant. Every norm larger than the initial two first occurs at such a transition, so (836) proves the global norm-record conjecture and the classification.

This is a strict reduction, not a proof: the stopped inequality remains to be established. Unlike the earlier reservoir, however, its left side is the exact boundary norm of one evaluation section. On each of the eight record chambers it is a fixed linear function of the six adjacent remainder logarithms. Explicitly,

$$\log\Theta = -\ell_0 - \max(-\ell_1, 0) - \max(-\ell_2, 0), \tag{837}$$

where $\ell_i = \log|\sigma_i(r_j/r_{j-1})|$ and $\ell_0 < 0$.

This scalar is the norm of a canonical tautological section. Let $\mathcal{U}_i$ be the line spanned by the state vector $(\sigma_i(u), \sigma_i(v))$, with norm P_i, and let $e_i \in \mathcal{U}_i^*$ select its v-coordinate. Since $\|e_i\| = |\sigma_i(v)|/P_i = e^{-\eta_i}$,

$$\frac{e_1 e_2}{e_0} \in \mathcal{U}_0\mathcal{U}_1^{-1}\mathcal{U}_2^{-1}, \qquad \left\|\frac{e_1 e_2}{e_0}\right\| = \Theta. \tag{838}$$

This line is not automatically $\mathcal{T}^{-1}$. The line $\mathcal{U}_i$ is generated by the transported state vector gP, whereas $\mathcal{T}_i$ is generated by a boundary row of g. Symplectic duality identifies the ambient rank-two systems but not these two point-dependent lines. Their pairing is exactly what the incidence tensor retains. Equation (837) is therefore the form in which the state-tautological metric must be compared, through the incidence section, with the prefix/Hodge metric and the primitive finite metric.

The comparison has an exact scalar identity. Put $X = gP = (u, v)^t$ and $w = e_2^t g$, so $v = wP$. With $H_i(X) = \log \max\{|\sigma_i(u)|, |\sigma_i(v)|\}$, define

$$\Delta_i(g, P) = H_i(X) - H_i(P) - H_i(w).$$

The state deficiency and prefix incidence proximity share the evaluation term $-\log|\sigma_i(v)|$, and hence

$$\eta_i = \lambda_i(w, P) + \Delta_i(g, P). \tag{839}$$

Writing $\chi(X) = H_1(X) + H_2(X) - H_0(X)$, the signed sum is

$$\log \Theta = \lambda_0 - \lambda_1 - \lambda_2 - \chi(X) + \chi(P) + \chi(w), \tag{840}$$

$$D_{\mathrm{out}}(w, P) + \log \Theta = \lambda_0(w, P) - \chi(X) + \chi(P) + 3h(P). \tag{841}$$

Thus the line mismatch is measured by the signed action defect $\Delta_0 - \Delta_1 - \Delta_2$. The cusp-transfer program transports both prefix rows and checks (839) exactly. On 1000 random primitive inputs with denominators at most 1000, the largest numerical error in (840) was $4.98 \cdot 10^{-14}$. At the 1354 pre-record transitions, the signed defect ranged from $-274.1155332854\ldots$ to 0. An exhaustive coefficient-3, denominator-10 grid gave $[-136.3950277262\ldots, 0]$ on 1074 pre-record transitions.

Those samples suggested the tentative bound $\Delta_0 - \Delta_1 - \Delta_2 = O_7(1)$ at a record. It is false. The failure is important because it identifies exactly what random boxes do not test: simultaneous alignment with the two outer conorm directions.

Proposition II.16.5 (Weak approximation defeats the local line comparison)**.** *For every $A > 0$ there is a primitive integral Rosen state for which the second transition is a strict first norm record, but at the preceding state*

$$\Delta_0 - \Delta_1 - \Delta_2 > A.$$

In particular the state and prefix lines cannot be compared with bounded error independently of the initial adelic height, even at a first record.

Proof. Fix t in the interior of the central Rosen interval. Choose positive integers $n \to \infty$ and $m = m(n)$ with $m/n^3 \to \infty$. Prescribe the initial affine slope at the distinguished embedding by

$$y_1 = -\frac{1}{m\lambda_0 + t}, \qquad x_0 = -\frac{1}{n\lambda_0 + y_1}.$$

Its first two digits are n, m, and its central states are x_0, y_1, t. At the two outer embeddings prescribe instead

$$x_i = -\frac{\operatorname{sgn}(\lambda_i)}{1 + n|\lambda_i|}, \qquad i = 1, 2.$$

Starting from the local pair $(x_i, 1)$, put

$$r_{1,i} = -1 - n\lambda_i x_i, \qquad r_{2,i} = -x_i - m\lambda_i r_{1,i}.$$

At the outer places $|r_{1,i}| = |x_i| \asymp n^{-1}$ and $|r_{2,i}| \asymp m/n$. At the central place

$$|x_0| \asymp n^{-1}, \qquad |r_{1,0}|, |r_{2,0}| \asymp (nm)^{-1}.$$

Thus the four successive norm sizes, after normalizing the second initial coordinate to one, are

$$1, \qquad \asymp n^{-3}, \qquad \asymp (mn^3)^{-1}, \qquad \asymp m/n^3.$$

The last is a strict first record.

For the prefix $B(n)$, the initial local sup norm is one, the transported sup norm is $|x_i|(1+o(1))$, and the selected-row norm is $n|\lambda_i|$. Therefore

$$\Delta_i = -2\log(n|\lambda_i|) + o(1)$$

and, since $|\lambda_1\lambda_2| = \lambda_0^{-1}$,

$$\Delta_0 - \Delta_1 - \Delta_2 = 2\log n - 4\log\lambda_0 + o(1) \longrightarrow \infty. \tag{842}$$

The image of K in $\mathbf{R}^3$ is dense: in the basis $1, \lambda, \lambda^2 - 1$, it is the image of $\mathbf{Q}^3$ under an invertible real matrix. Approximate the three prescribed slopes by one $x \in K$. All digit and record inequalities above are strict and hence survive sufficiently close approximation. Clearing denominators and removing the common ideal produces the claimed primitive pair. □

The new `-weak-family` mode of `rosen_cusp_transfer_tests.py` implements that proof by rounding in the Minkowski basis and then rechecking both digits and all four norms exactly. It found:

n	m	$\Delta_0 - \Delta_1 - \Delta_2$	record gap	$h(P)$	stopped excess
5	5,000	1.38334	1.05301	23.02585	−65.31954
10	50,000	2.52956	0.93831	27.63102	−78.67200
20	1,000,000	3.78180	2.76730	35.54409	−101.36087
50	10,000,000	5.52847	1.94549	43.74912	−126.33633

Here the stopped excess is

$$3h([r_1 : r_0]) + \log(1 + \Theta(r_1, r_0)) - 6h(P).$$

It becomes more negative while the local defect grows. The initial height needed for simultaneous weak approximation is precisely the arithmetic cost discarded by the false local comparison.

Route II.16.6 (Verdict)**.** The signed action defect is an exact diagnostic, not a separately bounded invariant. Charged cones and the adapted cusp metric control the prefix line, but the endpoint can still align adversarially at the two outer embeddings. The state line, the prefix/Hodge line, the incidence evaluation, and the primitive finite metric must remain in one Arakelov degree calculation. The surviving scalar target is exactly the stopped Θ-bound, with its $6h(P)$ term; a new-divisor charge or a uniform pointwise line comparison cannot replace it.

The next construction must therefore put the adapted logarithmic metric on the full transported evaluation line and induce across complete cusp excursions. Its boundary term must be the increase of $H_0(w)$ in (827), while its finite and initial-height terms charge weak approximation globally.

Experiment II.16.7 (The moving Asai wedge cancels unit height but not positive-weight leakage)**.** For the exact 169-step orbit used to test (763), form the rational Plücker vector

$$Z_j = \Phi(P) \wedge \Phi(g_j P) \in \bigwedge^2 \mathrm{As}(K^2).$$

The new program `rosen_asai_wedge_tests.py` constructs all 28 coordinates exactly, clears them to a primitive integral vector, and compares its height with the formal nearest-wall wedge weight $2(\chi - m) = -2\delta$. The output is

j	5	10	20	40	60
$2(\chi - m)$	−4.7566	−9.3210	−16.5444	−34.5668	−50.3924
$h(Z_j)$	31.7374	30.7734	33.3045	31.4936	34.0238
$h(\det(P, g_j P))$	10.5142	11.7982	18.8724	27.8440	38.4365

and the coordinate content of Z_j is one at every displayed step.

Two conclusions pull in opposite directions. First, the Plücker construction cancels most of the enormous unit growth in the determinant: the primitive height of Z_j remains nearly constant on this sample. Thus it is false that the wedge automatically pays the full coefficient height of the return word. In fact the elementary height inequality

$$h(Z_j) \le 3h(P) + 3h(g_jP) + O(1)$$

already predicts the scale in the table; it depends on the two projective endpoints, not on a chosen matrix representing the return.

Second, adding the last two entries of the first two rows would naively make a nonzero integral vector smaller than one. The exact Z_j is nonzero. The formal negative weight therefore does not bound the whole vector. It is the weight of the nearest-wall component obtained when the central factor lies exactly in the contracting eigenline. A continued fraction prefix gives only an approximation to that line; its small expanding component is amplified by the return and contributes positive-weight terms to the wedge. The remaining problem is to control that leakage, not merely to bound the primitive Plücker height.

This experiment sharpens the recurring obstruction. One copy of $\Phi(P)$ has no nonzero lower-weight covariant: lower Schur components vanish identically on the twisted Segre variety. A second copy creates the desired wedge but is moving, and the actual vector is not contained in the negative Hodge subbundle. An all-place auxiliary construction must control both its positive-weight leakage and its primitive height.

The program also tests the exact balanced-unit inverse family. In this case $h(Z_g(P))$ grows from 13.5815 at exponent $m = 1$ to 125.2760 at $m = 16$, while the formal negative wedge weight converges to -0.736277. The coordinate content is one, apart from a factor two at $m = 6, 13$. The growth agrees with

$$h(Z_g(P)) \le 3h(P) + 3h(gP) + O(1),$$

because gP is the large balanced-unit state. Those states occur after the old norm record and create no new norm record. Hence an all-state Plücker contraction is false for exactly the same reason as the reservoir inequality. The first-hitting restriction is not an expository convenience; it removes this explicit counterfamily.

Proposition II.16.8 (The Asai wedge is a rational $1 + 27$ problem)**.** *For $W = V_0 \otimes V_1 \otimes V_2$, with each V_i of dimension two, there is an $SL(V_0) \times SL(V_1) \times SL(V_2)$-equivariant decomposition*

$$\begin{aligned}\bigwedge^2 W \simeq {} & \Big(\bigwedge^2 V_0 \otimes \operatorname{Sym}^2 V_1 \otimes \operatorname{Sym}^2 V_2\Big) \\ & \oplus \Big(\operatorname{Sym}^2 V_0 \otimes \bigwedge^2 V_1 \otimes \operatorname{Sym}^2 V_2\Big) \\ & \oplus \Big(\operatorname{Sym}^2 V_0 \otimes \operatorname{Sym}^2 V_1 \otimes \bigwedge^2 V_2\Big) \oplus \bigotimes_{i=0}^{2} \bigwedge^2 V_i.\end{aligned}$$

For the cubic Asai representation, the first three nine-dimensional summands are permuted by Galois and descend together. Hence over $\mathbf{Q}$

$$\bigwedge^2 \operatorname{As}(K^2) = \mathcal{R}_{27} \oplus \mathbf{Q}.$$

If $Z_g(P) = \Phi(P) \wedge \Phi(gP)$, its scalar coordinate is

$$s_g(P) = \mathrm{N}_{K/\mathbf{Q}} \det(P, gP) \in \mathbf{Z}.$$

It vanishes exactly when g fixes $[P]$.

Moreover, after fixing norms, every decomposable bivector $z = s + r$ with $r \in \mathcal{R}_{27}$ satisfies

$$|s| \le C_{\mathrm{Pl}} \|r\|$$

for a fixed constant C_{Pl}.

There is a second application in which the scalar is the remainder norm rather than the fixed-point norm. For a prefix row $w = (a, b)$, put

$$w^\sharp = (-b, a), \qquad \mathcal{Z}(w, P) = \Phi(P) \wedge \Phi(w^\sharp).$$

Then

$$\det(P, w^\sharp) = ap + bq = r, \qquad s(w, P) = \mathrm{N}_{K/\mathbf{Q}}(r).$$

If $\mathcal{Z}_i$ denotes the nine-dimensional component in which V_i is alternating, then

$$\begin{aligned}
\log \frac{\|\mathcal{Z}_0\|}{|s|} &= \lambda_1(w, P) + \lambda_2(w, P) + O(1), \\
\log \frac{\|\mathcal{Z}_1\|}{|s|} &= \lambda_0(w, P) + \lambda_2(w, P) + O(1), \\
\log \frac{\|\mathcal{Z}_2\|}{|s|} &= \lambda_0(w, P) + \lambda_1(w, P) + O(1).
\end{aligned}$$

Proof. In the tensor square of W, an antisymmetric tensor has an odd number of antisymmetric factors; this gives the four displayed summands and the dimension count $9 + 9 + 9 + 1 = 28$. Tensoring the determinant pairings on the three V_i gives a nondegenerate alternating form Ω, and

$$\Omega(\Phi(P), \Phi(Q)) = \prod_i \det(\sigma_i P, \sigma_i Q) = \mathrm{N}_{K/\mathbf{Q}} \det(P, Q),$$

which proves the scalar formula. For the last assertion, projectivized decomposable bivectors form the closed Grassmannian $\mathrm{Gr}(2, W)$. The scalar bivector is inverse to the nondegenerate form Ω and has rank eight, whereas a decomposable bivector has rank two. The scalar point is therefore disjoint from the Grassmannian. Compactness gives the uniform separation. The same inequality follows algebraically from $z \wedge z = 0$.

For the remainder application, the component $\mathcal{Z}_0$ is

$$\det(P_0, w_0^\sharp)(P_1 \odot w_1^\sharp) \otimes (P_2 \odot w_2^\sharp)$$

up to a fixed normalization. Its norm is therefore comparable to

$$|\sigma_0(r)| \, \|P_1\| \|w_1\| \|P_2\| \|w_2\|.$$

Dividing by $|s| = \prod_i |\sigma_i(r)|$ gives the first proximity identity; the other two follow cyclically. □

The Cartan weights explain exactly which part remains difficult. In the ideal frame with central exponent a_0 and outer exponents a_1, a_2, the relevant pure lines in the summands alternating in V_0, V_1, V_2 have weights

$$2(a_1 + a_2), \qquad -2(a_0 - a_2), \qquad -2(a_0 - a_1),$$

respectively. A weight is just the logarithm of the scaling factor on that line. The last two are negative on a thick return and are the two pairwise Hodge contractions. The first is positive:

alternation in the central rank-two factor contributes a determinant, hence weight zero, and loses the central contraction. Yet this first summand is exactly the one whose norm gives $\lambda_1 + \lambda_2$. Thus the $1 + 27$ decomposition does not conceal the earlier multiplicity problem; it locates it in the single component $\mathcal{Z}_0$. The stopped vector also has off-line leakage, so these limiting weights are diagnostic rather than by themselves an estimate.

Lemma II.16.9 (The finite content is exactly the remainder ideal)**.** *For a record $r = ap + bq$, let $I_0(w, P)$ be the $\mathcal{O}_K$-ideal generated by the nine tensor coordinates of $\mathcal{Z}_0(w, P)$. Then*

$$I_0(w, P) = (r).$$

Proof. The explicit formula gives $I_0 \subseteq (r)$. Locally at a prime ideal, a primitive pair is a nonzero binary linear form over the residue field. The product of two nonzero linear forms is a nonzero quadratic, so at least one of its three coefficients is a unit. Therefore the three coefficients of each outer symmetric product generate the unit ideal. The nine tensor-product coefficients after division by r generate the product of those two unit ideals. This proves the reverse inclusion. □

Since $s = \mathrm{N}(r) = r\,\tau(r)\tau^2(r)$, the normalized square has the fractional-ideal divisor

$$\mathrm{div}_f \left(\left(\frac{\mathcal{Z}_0}{s} \right)^{\otimes 2} \right) = 2(r) - 2(s) = -2(\tau r) - 2(\tau^2 r). \tag{843}$$

This gives the exact all-place meaning of the outer label: the two conjugate denominator ideals occur together, with no distinction between new and old primes.

The tautological Θ-section supplies the missing integral twist. For the same selected remainder define

$$\vartheta(r) = \frac{\tau(r)\tau^2(r)}{r}. \tag{844}$$

As a section of the state-tautological line $\mathcal{U}_0\mathcal{U}_1^{-1}\mathcal{U}_2^{-1}$, its distinguished archimedean quotient norm is Θ, while its finite divisor is

$$\mathrm{div}_f\, \vartheta(r) = (\tau r) + (\tau^2 r) - (r). \tag{845}$$

Consequently

$$\mathrm{div}_f \left\{ \left(\frac{\mathcal{Z}_0}{s} \right)^{\otimes 2} \otimes \vartheta(r)^{\otimes 2} \right\} = -2(r). \tag{846}$$

Both outer denominator ideals cancel, leaving exactly the distinguished remainder divisor with multiplicity two. This proves, at every finite place and with no exceptional-prime error, the integral "one central copy" correction on the state-tautological line. It does not identify that line with the prefix-tautological Hodge line. Proposition II.16.5 shows that a bounded identification is false; the remaining bridge must be a stopped incidence-degree estimate which retains the initial height.

The moving tensor has a useful exact state-frame normal form. Write $X = gP = (u, v)^t$, let $w = e_2^t g = (c, d)$, and set $w^\sharp = (-d, c)$. Direct multiplication gives

$$gw^\sharp = -e_1, \qquad \mathrm{As}(g)\mathcal{Z}(w, P) = -Y(X) \wedge Y(e_1). \tag{847}$$

This removes the prefix from the second vector after transport. In the component alternating in V_0, the right side is

$$\det(X_0, e_1)(X_1 \odot e_1) \otimes (X_2 \odot e_1),$$

up to a fixed sign and normalization. Since $s = \mathrm{N}(v)$, its quotient norm is

$$\log \frac{\|\operatorname{As}(g)\mathcal{Z}_0\|}{|s|} = \eta_1 + \eta_2 + O(1), \tag{848}$$

and the other two components give cyclically $\eta_0 + \eta_2$ and $\eta_0 + \eta_1$.

Now combine (848) with the state twist, whose log norm is $\eta_0 - \eta_1 - \eta_2$. The squared tensor in (846) has state-frame log norm

$$2\eta_0 + O(1) \tag{849}$$

and finite denominator $(v)^2$. Dividing this squared bookkeeping by two, the completed incidence twist accounts for

$$\log|\mathrm{N}\, v| + \eta_0 + O(1). \tag{850}$$

The stopped cusp-transfer quantity is instead

$$3h([u:v]) + \log\Theta = \log|\mathrm{N}\, v| + 2\eta_0.$$

Thus the arithmetic section and its exact finite divisor supply the norm and one central cusp logarithm. Exactly one more η_0, or $2\eta_0$ on the squared line, must come from the logarithmic Hodge/cotangent metric without introducing another finite divisor. This is a more precise target than "compare two tautological lines." The section $\Xi \in H^0(\Omega^1_{\log} \otimes \mathcal{M}^{-2})$ has the correct bundle type; the remaining task is to prove that its adapted metric realizes this second copy on every stopped chamber and that its global degree is charged to $6h(P)$.

Route II.16.10 (Exact one-copy audit)**.** Equation (847) is elementary tensor functoriality, while (846) is the exact finite-place calculation. Together they prove that the recurrent loss is exactly one central logarithm. They do not prove that the Ξ-metric supplies it; that is now the sharply isolated Hodge–Arakelov boundary assertion.

The sign of that assertion can also be audited exactly at the cusp. Let $L = -\log t$, and use the zero-balanced Schwarz expansions

$$\tau_k(t) = i(L + c_k) + O(tL), \qquad c_k = -2\gamma - \psi(a_k) - \psi(b_k).$$

In normalized Hodge frames one obtains

$$\|\mathcal{M}^{-2}\| = C_{\mathcal{M}} \frac{L + c_1}{(L + c_2)(L + c_3)}(1 + O(tL)), \tag{851}$$

$$\|dt/t\|_{\Omega^1_{\log}} = C_\Omega (L + c_1)(1 + O(tL)), \tag{852}$$

$$\|\Xi\| = C_\Xi \frac{(L + c_1)^2}{(L + c_2)(L + c_3)}(1 + O(tL^2)). \tag{853}$$

Indeed a normalized holomorphic form has Hodge norm proportional to $\sqrt{\operatorname{Im}\tau_k}$, while the pullback hyperbolic norm of dt/t is $\operatorname{Im}\tau_1/|t\tau_1'|$. The last line is equivalently the product of the two outer hyperbolic derivative norms. Consequently

$$\log\|\mathcal{M}^{-2}\| = -\log L + O(1), \qquad \log\|dt/t\| = \log L + O(1), \qquad \log\|\Xi\| = O(1).$$

Route II.16.11 (Higgs-ratio sign audit)**.** The section Ξ is a canonical transfer, not a free gain: it moves one $\log L$ between the Hodge and logarithmic-cotangent factors, with opposite signs, and changes no finite divisor. The remaining chamber lemma must therefore prove that the correctly oriented $\log L$ equals the stopped Rosen cusp depth $\eta_0 + O_7(1)$, then bound the resulting global degree by $6h(P)$. The extended program `hodge_cusp_derivative_tests.py` prints all three factors in (851)–(853); the product tends to one, while $\log L + \log\|\mathcal{M}^{-2}\| \to 0$.

There is also an exact dynamical formula for the logarithm being transferred. If the distinguished Rosen state is $x = v/u \in I \setminus \{0\}$, then $|x| < 1$ and

$$\eta_0 = -\log|x|, \qquad \mathcal{R}'(x) = x^{-2}.$$

Therefore

$$\log|\mathcal{R}'(x)| = 2\eta_0. \tag{854}$$

The section $\Xi \in H^0(\Omega^1_{\log} \otimes \mathcal{M}^{-2})$ is equivalently a bundle map $\mathcal{M}^2 \to \Omega^1_{\log}$. At an interior point $z \in \mathbb{H}$, let $J_{\mathcal{M}^2}(\gamma, z)$ and $J_\Omega(\gamma, z)$ denote the two metric transport factors. If $\beta = \log\|\Xi\|$, equivariance gives

$$\log J_{\mathcal{M}^2}(\gamma, z) = \log J_\Omega(\gamma, z) + \beta(z) - \beta(\gamma z). \tag{855}$$

The two β-terms telescope. They are bounded at the cusp by (853); the only algebraic correction is the known triple zero at the order-seven point. Arnoux–Schmidt's explicit cross-section closes the apparent boundary-to-interior gap. Their cross-section matrices $A_{\pm1}(x, y)$ return under right multiplication by $g_t = \mathrm{diag}(e^{t/2}, e^{-t/2})$ at the exact time

$$t = -2\log|x| = \log|\mathcal{R}'(x)|; \tag{856}$$

Lemma 5.2 gives the matrix calculation and Theorem 6.1 states that it applies to Rosen digits after replacing d by $d\lambda$ [5]. Since the uniformizing variation is maximal Higgs, $\mathcal{L}^2_0 \simeq \Omega^1_{\log}$, and the diagonal flow acts on that square with weight e^t. Hence

$$\log J_\Omega(\gamma, z(x, y)) = t = 2\eta_0 \tag{857}$$

with no compactness or large-digit error. Only the endpoint proximity to the order-seven zero of Ξ appears to remain. The explicit domain shows that it does not occur. For $q = 7$, the original natural-extension domain is a finite union of rectangles $J_j \times [0, L_j]$, with $0 \le L_j \le R$, where $R^2 + (2-\lambda)R - 1 = 0$ [25, Section 5.1]. After the coordinate change $Y = y/(1+xy)$, one still has $Y \ge 0$, since $1 - (\lambda/2)R > 0$.

Write $v = \sin(\pi/7)$ and $\rho_7 = -\lambda/2 + iv$. Direct solution of $A_{\pm1}(x, Y)i = \rho_7$ gives

$$Y = \pm\sqrt{\frac{1-v}{v}}, \qquad x = \frac{\lambda}{2} + vY.$$

The negative solution has $Y < 0$; the nonnegative solution has $x > \lambda/2$. Thus the compact cross-section is disjoint from the zero of Ξ. Its endpoint norm $\beta = \log\|\Xi\|$ is uniformly bounded, and no order-seven divisor term enters a Rosen edge.

Route II.16.12 (The explicit stopped edge bundle)**.** Tensor the squared incidence twist in (846) with the $\mathcal{M}^2$-factor on the next Rosen edge. The incidence factor has finite divisor $-2(v)$ and state-frame log norm $2\eta_0 + O(1)$. The edge factor has no finite divisor and contributes the other $2\eta_0$ by (857), up to the telescoping endpoint norm of Ξ. Division by two gives

$$\log|\,\mathrm{N}\,v| + 2\eta_0 + O_7(1) = 3h([u : v]) + \log\Theta + O_7(1).$$

Thus the boundary coefficient, the finite divisor, the boundary-to-interior lift, and the endpoint term are exact. The remaining assertion is the global Arakelov degree bound for this explicit edge section by $6h(P)$, including only its primitive positive-weight leakage. This is narrower than the earlier request to invent an unspecified metric comparison.

Experiment II.16.13 (Exact implementation audit)**.** The program `rosen_incidence_plucker_tests.py` constructs all 28 rational Plücker coordinates, checks the invariant pairing

$$\Omega(\Phi(P), \Phi(w^\sharp)) = \mathrm{N}(ap + bq),$$

and asks PARI/GP for the ideal sum of the nine $\mathcal{Z}_0$-coordinates.

The two named examples give the following decisive records:

orbit	j	$\log\lvert s\rvert$	$h(\mathcal{Z})$	$\mathrm{N}(I_0)/\lvert\mathrm{N}(r)\rvert$
low height	7	14.454175	58.915363	1
denominator 20	38	24.281495	172.319092	1.

The rational coordinate content is one in both rows. A second audit used 500 random pairs (p, N), chosen primitive by $\gcd(\lvert\mathrm{N}(p)\rvert, N) = 1$, and produced 453 positive-step records. Every one satisfied

$$I_0(w, P) = (r).$$

The full rational Plücker vector sometimes had small content—as large as 26 in this sample—but the outer component never had ideal content beyond the visible factor r.

Lemma II.16.9, rather than the sample, rules out a prime-birth or hidden-gcd proof. It does not oppose an Arakelov proof. Rather, it says exactly what such a proof must use: the finite term is the complete divisor $(ap + bq)$, with no bonus, and the missing decrease must be supplied by the archimedean Hodge metrics in the same global degree.

Remark II.16.14 (The component is not yet a global metrized quotient). The preceding proposition constructs an algebraic section with the right finite divisor, but it does not by itself construct the desired Arakelov height. The individual component $\mathcal{Z}_0$ is defined over K; its three real conjugates are $\mathcal{Z}_0, \mathcal{Z}_1, \mathcal{Z}_2$. Therefore the ordinary adelic norm relative to s sums to

$$\begin{aligned}(\lambda_1 + \lambda_2) + (\lambda_0 + \lambda_2) + (\lambda_0 + \lambda_1) + O(1)& \\ = 2(\lambda_0 + \lambda_1 + \lambda_2) + O(1).&\end{aligned}$$

This is global and unit-invariant, but it counts the central proximity twice. The rational 27-dimensional sum has the same problem because an ordinary norm sees all three components.

The remaining construction is consequently very specific. One needs a Hodge-renormalized adelic quotient which selects the distinguished outer component while canceling the two conjugate central contributions. Its finite metric must still be the single $\mathcal{O}_K$-lattice whose content is (r). An arbitrary asymmetric choice of three real norms is not a proof: compatibility with the unit action and automorphy is exactly what earlier scalar metric twists failed to establish.

Route II.16.15 (The tautological correction and its $\mathcal{M}$-boundary model)**.** The arithmetic quantity $\chi(w)$ is the height of a moving prefix row, not literally the Cartan character of a fixed Hodge line. On the Rosen–Hodge correspondence let $\mathcal{T}_i$ be the tautological line spanned by $\sigma_i(w)$, with its Hodge metric, and put

$$\mathcal{T} = \mathcal{T}_1 \otimes \mathcal{T}_2 \otimes \mathcal{T}_0^{-1}.$$

At a fixed Hodge frame, equivalence of Euclidean, Hodge, and sup norms gives

$$\log \|\mathcal{T}(w)\| = \chi(w) + O(1).$$

The exact metrized candidate is consequently

$$\left(\mathscr{E}_0 \otimes \mathscr{S}^{-1}\right) \otimes \mathcal{T}^{-1},$$

whose logarithmic norm on the incidence section is

$$\lambda_1 + \lambda_2 - \chi + O(1).$$

After adding the initial budget $3h(P)$, this is D_{out}.

The fixed Hodge line

$$\mathcal{M} = \mathcal{L}_1 \otimes \mathcal{L}_2 \otimes \mathcal{L}_0^{-1}, \qquad \deg \mathcal{M} = -\frac{1}{28},$$

is the boundary model of $\mathcal{T}$. Its spectral return weight is $\chi_{\mathrm{sp}} = -a_0 + a_1 + a_2$. In a controlled return frame the tautological line approaches this Hodge line, and the ideal pure-line weight becomes

$$2(a_1 + a_2) - \chi_{\mathrm{sp}} = a_0 + a_1 + a_2,$$

which is strictly positive. Thus the formerly troublesome $\mathcal{Z}_0$-component is expanding at the Hodge boundary.

The square of the actual correspondence bundle is

$$\mathscr{A} = \left(\mathscr{E}_0 \otimes \mathscr{S}^{-1}\right)^{\otimes 2} \otimes \mathcal{T}^{-2}.$$

At the boundary, the Higgs-ratio section

$$\Xi \in H^0\left(\Omega^1_{\log} \otimes \mathcal{M}^{-2}\right)$$

controls the squared comparison $\mathcal{T}^{-2} \sim \mathcal{M}^{-2}$. Its divisor is three times the order-seven orbifold point, and it has neither a zero nor a pole at the cusp. One may not simply replace $\mathcal{T}$ by $\mathcal{M}$: proving their bounded comparison along stopped cusp excursions is part of the boundary lemma.

The finite divisor of the state-tautological route is now exact. Equations (844)– (846) give a compatible integral trivialization and reduce its divisor exactly to $-2(r)$, with no bad-prime error. The prefix/Hodge route uses a different tautological line. Three comparison checks remain:

1. combine the state line carrying Θ with the prefix line $\mathcal{T}^{-1}$ in the full stopped incidence tensor, charging the unbounded mismatch to its initial arithmetic degree rather than asserting a bounded line comparison;
2. extend the comparison between $\mathcal{T}^{-2}$ and $\mathcal{M}^{-2}$ defined by Ξ through every stopped cusp chamber with a bounded metric error; and
3. in the adapted nilpotent-orbit metric, prove that every negative-weight boundary specialization of the primitive incidence tensor forces a cusp incidence or the scalar fixed-line equation.

The terminal shear proposition below settles the second item away from the order-seven orbifold point. The first and third together are the exact stopped Hodge–Arakelov incidence lemma. They are a weight-filtration and global-degree calculation on the Rosen–Hodge correspondence, not a claim about prime production.

Experiment II.16.16 (The tautological and spectral characters differ)**.** The program `rosen_asai_wedge_tests.py` now prints both

$$\chi(w) = H_1(w) + H_2(w) - H_0(w) \quad \text{and} \quad \chi_{\mathrm{sp}}(g) = -a_0(g) + a_1(g) + a_2(g).$$

On the exact 169-step orbit, their differences at steps $5, 10, 20, 40, 60$ are

$$-0.679693, \quad -0.921314, \quad -0.597638, \quad -1.525553, \quad -2.182124.$$

Hence replacing the tautological prefix line by $\mathcal{M}$ is not an identity; a uniform comparison has to be proved from the return geometry, not assumed from the definitions. On the balanced inverse one-branch family the difference is $2.79 \cdot 10^{-3}$ at exponent one, $5.12 \cdot 10^{-7}$ at exponent four, and tends to numerical zero. This confirms the correct interpretation: $\mathcal{M}$ is the asymptotic Hodge boundary model on a controlled face, whereas $\mathcal{T}$ is the actual global line. A proof must establish their comparison chamber by chamber after cusp induction.

Proposition II.16.17 (Uniform tautological–Hodge comparison on a charged core)**.** *Let Q be a nonempty word in A_2, A_3, A_4, A_5, and let w be either row of Q. After the fixed sign gauges defining the positive fusion model, there is a field constant C_{ch} such that*

$$|H_i(w) - a_i(Q)| \le C_{\mathrm{ch}} \qquad (i = 0, 1, 2).$$

Consequently

$$|\chi(w) - \chi_{\mathrm{sp}}(Q)| \le 3C_{\mathrm{ch}}.$$

Proof. Each charged generator has strictly positive entries at each of the three gauged real parameters. Its projective action sends $[0, \infty]$ into the compact interval between its two column slopes. Because the alphabet is finite, one compact interval works for all four generators. Applying the same argument to the transposes controls row slopes. It follows that in every nonempty product the entries within either row, and within either column, are comparable by a uniform constant. Lemma II.2.39 makes the comparison simultaneous across the embeddings.

For a positive determinant-one matrix with those entry ratios, its trace, its Perron eigenvalue, its largest entry, and either row norm are all comparable by a constant depending only on the compact slope interval. Taking logarithms proves the first assertion. The second follows by forming $H_1 + H_2 - H_0$. □

Proposition II.16.18 (Terminal parabolic shear comparison)**.** *Let $Q = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ be a nonempty charged product. For $n \ge 0$, put*

$$U_n = \begin{pmatrix} 1 & n\lambda \\ 0 & 1 \end{pmatrix}, \qquad L_n = \begin{pmatrix} 1 & 0 \\ n\lambda & 1 \end{pmatrix}.$$

Every boundary row of QU_n, QL_n, U_nQ, L_nQ, after the relevant cusp automorphy, is uniformly comparable either to a charged row of Q or to the invariant parabolic line. Thus the comparison between the actual tautological square $\mathcal{T}^{-2}$ and its $\mathcal{M}^{-2}$ boundary model has $O_7(1)$ error through an arbitrary terminal parabolic run.

Proof. The four products are

$$QU_n = \begin{pmatrix} a & an\lambda + b \\ c & cn\lambda + d \end{pmatrix}, \qquad QL_n = \begin{pmatrix} a + bn\lambda & b \\ c + dn\lambda & d \end{pmatrix},$$
$$U_nQ = \begin{pmatrix} a + n\lambda c & b + n\lambda d \\ c & d \end{pmatrix}, \quad L_nQ = \begin{pmatrix} a & b \\ n\lambda a + c & n\lambda b + d \end{pmatrix}.$$

The charged projective-cone lemma bounds every relevant entry ratio at all three embeddings. For bounded n, compactness applies. For unbounded n, the expanding row divided by $n\lambda$ approaches the unchanged charged row or the invariant cusp row with $O_7(1/n)$ error; the other row is already charged. Finally $h_{\mathrm{ad}}(\tau + n) = e^{nN} h_{\mathrm{ad}}(\tau)$ absorbs the shear exactly, and Schmid comparison between the Hodge and adapted metrics is bounded. Squaring removes the finite sign gauges. □

Remark II.16.19. Proposition II.16.17 proves the comparison on the thick charged portion, and Proposition II.16.18 extends it through an arbitrary terminal partial run. The unresolved boundary calculation is therefore no longer the tautological metric comparison; it is the specialization of the primitive incidence tensor in the eight record chambers.

Route II.16.20 (A concrete Arakelov quotient)**.** Proposition II.16.8 replaces an unspecified auxiliary representation by a specific one. The neutral scalar is the integer whose vanishing gives fixed-point closing; all expanding-line leakage lies in the rational bundle $\mathcal{R}_{27}$. The Grassmannian separation is the needed zero estimate: a nonzero secant Plücker vector cannot converge to the scalar line.

For the prefix choice $w^\sharp$, the neutral scalar is instead the exact norm record, and the component $\mathcal{Z}_0$ realizes the outer incidence:

$$D_{\mathrm{out}}(w,P) = \log \frac{\|\mathcal{Z}_0(w,P)\|}{|s(w,P)|} - \chi(w) + 3h(P) + O(1).$$

This is the explicit algebraic section sought in Route II.16.3; at every finite place its scalar coordinate has the divisor $(ap+bq)$, while at infinity its three components are the pairwise Hodge proximity sums. Warning II.16.14 explains why an additional Hodge-renormalized quotient, rather than the ordinary component height, is still required.

The remaining inequality is now to bound the stopped $\mathcal{R}_{27}$-coordinate relative to the scalar coordinate with the adapted Hodge metrics at infinity and the primitive metric at every finite place. Experiment II.16.7 shows why this must be stopped at a norm record. Proving that relative Arakelov degree inequality would force $s_g(P) = 0$, hence closing.

Route II.16.21 (Projective tangent height: the moving-frame bill)**.** Put $C = \Phi(w^\sharp)$ for the rational Asai cusp line selected by the future prefix and consider

$$\mathfrak{z}_C(P) = (\Phi(P) \wedge C) \otimes C^{-2}.$$

This is not an ad hoc quotient: it is the tangent vector from C toward $\Phi(P)$ in projective Asai space. If C were the top Anosov line, a tangent coordinate indexed by a nonempty set $F \subset \{0,1,2\}$ would have twisted weight

$$-4\sum_{i\in F} a_i + 2(-a_0 + a_1 + a_2).$$

The premise is false for the cusp chosen by the prefix. If the selected row is $w = (a,b)$, then $w^\sharp = (-b,a)$ is its exact kernel and the prefix sends it to a coordinate vector. Since w is the expanding left covector on a charged return, $w^\sharp$ is the contracting right line. The applicable weight is therefore

$$4\sum_{i\in F} a_i + 2(-a_0 + a_1 + a_2),$$

whose $F = \{0\}$ value is $2(a_0 + a_1 + a_2) > 0$. The proposed tangent normalization does not remove positive leakage.

The program `rosen_incidence_projective_tangent_tests.py` measures the missing bill. In a fixed Euclidean arithmetic frame, the logarithmic tangent norm change is $58.0414, 65.7672, 67.7408$ at the three low-height records, and reaches 180.0538 at step 38 of the denominator-20 orbit. At the last record $1 - |\langle C, v_{\min}\rangle| < 2.6 \cdot 10^{-55}$, confirming that C is the least singular right line. Adding either sign of the exact Rosen cotangent cocycle does not bound the tangent norm. The same program now computes the ordinary all-place flag height $h(\Phi(P) \wedge C) - 2h(C)$. Its change at

step 38 is still 178.6011 . . ., and every displayed cusp and incidence vector has content one. Thus primitive finite metrics do not pay the bill; the needed all-place object must carry the Hodge metric itself.

The projective cotangent $\kappa_C = \Omega(C, -) \otimes C$ has the correct opposite orientation. Its weight after the $\mathcal{M}^2$-twist is exactly $-4a_0$, so the outer weights cancel. But the scalar proxy $2h(C) - 2 \log J_\Omega$ is still too coarse: the new program `rosen_incidence_cotangent_arakelov_tests.py` finds its excess over $6h(P)$ growing to $29.2605\ldots$ on the balanced inverse family, while the actual stopped scalar is far below budget. Passing to an operator norm has discarded the decisive incidence proximity. The remaining candidate must therefore be the evaluated cotangent/incidence section with Hodge metrics at infinity and primitive metrics at finite places, not the tangent flag or the raw cotangent line by itself.

More precisely, let L_P, L_C be the rational lines spanned by $\Phi(P), C$. Restriction gives a map of rational lines

$$\kappa_C : L_P \longrightarrow L_C, \qquad \Phi(P) \longmapsto \pm \mathrm{N}(ap + bq)C.$$

Its finite divisor is the single norm record, while its normalized real norm relative to the full cotangent norm is the product of the three actual incidence sines. This follows immediately by writing the Asai form as the tensor product of the three determinant forms. It is the first candidate here which simultaneously uses one global line, has the correct cotangent Cartan sign, and retains the proximity discarded by a full operator norm. The remaining proof must bound the Arakelov slope of this restricted map after the state-tautological and Hodge twists.

There is a sharper two-factor version. Define

$$K_{12,C} = \mathrm{id}_{V_0} \otimes \big(\omega_1(c_1, -)c_1\big) \otimes \big(\omega_2(c_2, -)c_2\big).$$

Its normalized evaluated norm is exactly the product of the two outer incidence sines, hence its negative logarithm is $\lambda_1 + \lambda_2 + O_7(1)$. Its algebraic evaluation coefficient is $\tau(r)\tau^2(r)$. Dividing by the already proved state character $\vartheta(r) = \tau(r)\tau^2(r)/r$ leaves r, so the finite degree is $\log|\mathrm{N}\, r|$. Formally adding the prefix-tautological metric and the primitive input-line metric gives

$$\lambda_1(w, P) + \lambda_2(w, P) - \chi(w) + 3h(P) = D_{\mathrm{out}}(w, P)$$

up to the fixed comparison of norm models. This proves the required local factorization, but not yet an all-place descent. Galois conjugation cycles $(K_{12,C}, K_{20,C}, K_{01,C})$, so an ordinary descended norm sees all three pairwise incidences. The exact remaining construction is a non-parallel Hodge–Arakelov metric, using $\mathcal{M}$ and the state and prefix lines, which descends the distinguished outer factor without an arbitrary asymmetric rescaling. After that compatibility theorem, the last inequality is the lower slope bound $D_{\mathrm{out}} \geq -O_7(1)$ at a first norm record.

Route II.16.22 (Audit of the synchronized maximum in `curt14a`)**.** The sound algebraic observation in `curt14a.tex` is that the order of maximum and adelic summation can be enforced by a global slope. Give the norm line $\overline{\mathcal{N}}(r)$ the finite lattice $r^{-1}\mathcal{O}_{K,\mathfrak{p}}$ and standard real metrics. For the maximum-norm direct sum one has exactly

$$\widehat{\mu}_{\max}\left(\bigoplus_{k\leq j}^{\infty} \overline{\mathcal{N}}(r_k)\right) = \max_{k\leq j} \log|\mathrm{N}\, r_k|.$$

A mixed K-line cannot choose one coordinate at one place and a different coordinate elsewhere, because the product formula applies to one global coordinate. This is the correct answer to asynchronous prime births.

It is nevertheless a representation of the desired record, not a bound for it: the direct sum grows by one line at every time. More decisively, the "isogeny–shape reservoir" proposed later in `curt14a.tex` is the already rejected inequality

$$\max_{k\le j} \log|\,\mathrm{N}\, r_k| + U_j \le 6h(x_0) + O_7(1).$$

Proposition II.9.46 gives an exact family on which its excess grows linearly. No Bott–Chern correction which is bounded on the state space can repair an unbounded excess. The usable idea is therefore narrower: select the destabilizing record line globally, stop at its first occurrence, and seek a Hodge–Arakelov degree bound for that one line. The projective tangent flag above is a rejected candidate; the evaluated cotangent/incidence section is the finite-rank object which still retains both the correct Cartan sign and the proximity to the initial arithmetic point.

Route II.16.23 (All-place Bost slopes via arithmetic holonomy)**.** Calegari–Dimitrov–Tang, Theorem 2.6 and §2.3, give a particularly close template for the proposed global argument [57]. Their proof starts with an integral lattice of holonomic auxiliary functions, equips it with local metrics, evaluates it at a rational approximant, and invokes Bost's slopes inequality. The local evaluation bounds are then summed over every finite and infinite place. This is the important feature for the present problem: denominator losses and archimedean gains occur in one degree calculation, so loss cannot migrate among a sequence of newly appearing primes.

Write $x = p/q$ for the affine coordinate of the fixed initial point $P = (p, q)$. The natural approximants are the prefix ratios $\beta_j = -b_j/a_j$. The corresponding fixed local targets would be

$$\eta_0 = \sigma_0(x), \qquad \eta_1 = \xi_1(\sigma_0(x)), \qquad \eta_2 = \xi_2(\sigma_0(x)),$$

where ξ_1, ξ_2 are the boundary limit maps of the two outer Hodge representations. This makes the analogy with their simultaneous approximation theorem exact at the level of local targets.

The hypotheses needed to turn the analogy into an application are not present. Their theorem requires a fixed finite-dimensional collection of holonomic power series which realizes the targets and has a controlled denominator type. The outer period maps are hypergeometric, hence holonomic, in the algebraic coordinate t; but the arbitrary endpoint $\sigma_0(x)$ is given in the boundary coordinate of $\mathbb{H}$. We have no arithmetic specialization map taking that endpoint to an algebraic t-value, and hence no uniform coefficient-denominator estimate for power series centered at it. Their theorem would in any case give a constant depending on the local targets, whereas the norm-record theorem needs that dependence bounded explicitly by $h(P)$.

Thus this route neither proves nor contradicts Conjecture II.16.1. It identifies a concrete construction problem: build the auxiliary bundle on the Rosen–Hodge correspondence itself, use the evaluation section $ap + bq$, and prove that its Bost degree is bounded uniformly over the symbolic cylinders after the adapted cusp normalization. If this is achieved, all finite valuations are automatically included in the same slope inequality.

II.17 Reproducibility index

All arithmetic state updates in the programs below are exact. Rosen branch selection uses rational root isolation as proved above; floating point is used only to design weak-approximation targets or for displayed archimedean measurements. The cusp and special certificates are checked as identities in K.

program	role and outcome
`curt10_spectral_tests.cpp`	Exact fusion-coefficient word search; led to the proved uniform outer gap and now audits the conjectural gap per induced block.
`intermediate_unit_word_search.py`	Bidirectional exact search in matrices over $\mathbb{Z}[\lambda]/(\lambda^3-\lambda^2-2\lambda+1)$. Its two depth-seven balls contain (8400) matrices each and recover the displayed length-fourteen word for $\mathrm{diag}(\lambda,\lambda^{-1})$ in $\langle S, T(1), T(\lambda)\rangle$; exact reevaluation verifies the identity.
`levi_residue_tests.py`	Computes the ray-class obstruction to lifting a residue Borel to an integral Borel. It gives an exact three-digit Rosen counterexample of norm (769), and on the (62) nontrivial prefix residues of the special orbit of (671) finds (30) unit-liftable and (32) non-liftable classes. PARI ideal-star discrete logarithms are exact.
`trace_parameter_monotonicity_tests.py`	Builds the exact trace polynomials $F_w(x)$ and gives rational Bernstein certificates for $xF_w' - s(w)F_w \geq 0$ on $[6/5, 2]$; all 279808 charged words of length seven certify, with equality only on the sharp one-turn rays.
`hyperbolic_length_derivative_tests.py`	Certifies exactly $4x^2F_w'(x)^2 - \nu(w)^2(F_w(x)^2-4) \geq 0$ on $[6/5, 2]$ for all 279934 nontrivial words of length seven, with no interval subdivision. This is the infinitesimal sharp fusion candidate.
`trace_induced_derivative_semiring_tests.py`	Decomposes $2xF_w' - \nu(w)F_w$ in $\mathbf{Z}[y] \oplus x\mathbf{Z}[y]$, $y = x^2-1$. All 279808 charged words of length seven are coefficientwise nonnegative; the only identically zero words are the $2n$ sharp one-turn rotations.
`charged_derivative_macro_certificates.py`	Checks all $16+64$ charged macros of lengths two and three. Their four entries contain 1362 positive derivative-gap coefficients, the least one, and no negative coefficient. The product rule turns this finite certificate into Theorem II.2.26 at every length.
`trace_differential_macro_tests.py`	Propagates the cleared logarithmic derivative and two endpoint potentials in $\mathbf{Z}[z]$; all 16256 four-letter-core words of length seven with an A_2 or A_5 satisfy the conjectural entrywise macro.
`trace_open_charged_tests.py`	Certifies all four entries of $xQ' - (s-1)Q$ for every A_2, A_5-word through length twelve on $[6/5, 2]$; all certificates are direct Bernstein-positive, with no subdivision.
`second_outer_alternating_certificates.py`	Finishes by nine reductions in $\mathbf{Z}[L]/(L^3-L^2-2L+1)$ used in Theorem II.2.35. The alternating-block Chebyshev recursion proves the sharp endpoint comparison for every length and hence the four-letter-core second-outer trace theorem.

program	role and outcome
`strong_replacement_certificate.py`	Certifies the cross-embedding projective cone and all four multi-affine corners of Lemma II.2.40. It checks 72 nonzero coefficient polynomials; the least rational Bernstein coefficient is 5339/10000. This upgrades the second-outer theorem from the four-letter core to all six generators.
`two_charge_parabolic_kernel.py`	Prints the sixteen exact matrix-unit derivative kernels in Theorem II.2.31; degree parity leaves the single paired remainder $2(x^4-2)$.
`first_outer_charged_certificate.py`	Checks the 84 charged cyclic base words needed in Theorem II.2.33; the least exact Bernstein coefficient on $[9/5, 181/100]$ is $54/25$.
`fusion_additive_gap_tests.py`	Stable high-precision search on words with up to thirty induced parabolic blocks; tests additivity of the outer surplus.
`parabolic_relative_gap_tests.py`	Exact coefficientwise-Wronskian audit for the proved relative gap on arbitrary alternating A_1, A_6 run words.
`active_macro_gap_tests.cpp`	Exhausts the admissible alphabet $\{A_2, A_3, A_4, A_5, A_1A_6, A_6A_1\}$. It finds short negative-character counterexamples to the overcharged $s+m$ bound, but no failure through length ten for the corrected weight $m + s/2$.
`relative_fusion_block_certificate.py`	Works exactly in $\mathbf{Z}[\sqrt{R}]/((\sqrt{R})^6 - 4(\sqrt{R})^4 + 3(\sqrt{R})^2 + 1)$. It certifies the ten-parabolic peeling block, classifies all 13600 residual second-outer blocks with $p \le 10, \ell \le 4$, finds exactly 22 weak-letter exceptions, and classifies all 242 compatible exception pairs, of which 30 remain open entrywise. Cyclic trace comparison is exact and nonnegative for all $12, 122, 1332$ exceptional cycles of lengths one, two, and three.
`curt10_rosen_tests.py`	Exact Rosen grid enumeration: 36111 terminating/periodic states, none unresolved.
`roth_target_tests.py`	Vojta Gaussian-parameter optimization; shows fixed positive target charge cannot survive the parabolic ray.
`legendre_threshold_tests.py`	Computes the principal and mediant Rosen constants and the first principal failure at $A_2A_1^8$.
`rosen_certificate_tests.py`	Constructs exact cusp matrices and exact special fixers/eigenvalues.
`rosen_polarization_tests.py`	Searches all four cubic polarizations; finds no vector combining central height with tangent proximity.
`height_block_criterion.py`	Evaluates the cubic Panti conorm product on $A_2A_1^N$; raw descent fails asymptotically like N^{-2}.
`rosen_height_barrier_tests.py`	High-precision Weil-height excursions; corrected a false low-precision counterexample.
`rosen_block_excursion_tests.py`	Constructs adversarial exact digit blocks with optimized outer directions; supports a quadratic multiplicative-height barrier.

program	role and outcome
`rosen_unit_perturbation_tests.py`	Tests cusp endpoints perturbed by tiny distinguished units; finds the sharp candidate ratio tending to 2, with exact cusp/special outcomes.
`rosen_sunit_transition_tests.py`	Enumerates exact unit-to-unit Rosen transitions and the unit-core cycles.
`rosen_sunit_degeneracy_tests.py`	Constructs 6851 exact degenerate starts and verifies the proved two-step separation throughout their first 30 iterates.
`rosen_sunit_norm_identity_tests.py`	Checks the exact digit divisibility by neighboring signed norms and audits the unit-perturbation height and sharp sextic norm profiles with adaptive precision.
`rosen_asai_norm_cocycle.py`	Derives the integral 8-dimensional branch matrix (II.9.29), proves determinant one symbolically, and checks it, the symplectic norm pairing, and the twisted-Segre ideal relations on random exact field pairs.
`rosen_norm_barrier_random_tests.py`	Tests the sextic norm barrier on 300 exact random primitive inputs; all close and the largest observed slack is $-14.0783\ldots$.
`rosen_norm_reservoir_tests.py`	Tests the exact decomposition $6h = 2M + U$ and the stronger running-record inequality $\max_{k\le j} M_k + U_j \le 6h(x_0)+B_7$. It finds no growing excess on random, rational-denominator, or sharp unit-ray samples; the balanced inverse program below explains why these samples were not decisive.
`rosen_balanced_unit_inverse_tests.py`	Constructs the exact inverse branch $-1/(\theta^m+n_m\lambda)) \to \theta^m$ and proves computationally that the rejected reservoir excess grows linearly.
`rosen_record_hodge_compensation_tests.py`	Balances the initial representative by one unit and shows on 1146 strict norm records that the Hodge drop and cusp correction cannot be estimated independently.
`rosen_record_prefix_height_tests.py`	Tracks $r_j = a_j x + b_j$ exactly and shows that prefix height can exceed the initial height by a factor 16.39 at a safe norm record.
`rosen_record_residual_prefix_tests.py`	Transports the exact primitive prefix row, checks the global incidence identity at strict norm records, refutes the residual-prefix-only bound, and audits the surviving outer-incidence defect.
`rosen_outer_incidence_dynamics.py`	Audits every state, finds the positive low-height sextic surplus, and disproves separate nonnegativity of the outer-incidence defect.

program	role and outcome
`rosen_cusp_transfer_tests.py`	Audits the proved all-place cusp-transfer inequality, measures its single factor $\Theta = e^{\eta_0-\eta_1-\eta_2}$ immediately before strict norm records, and checks the stopped Θ-envelope on random, low-height, and balanced inverse-unit inputs. Its exact weak-approximation mode constructs first records with unbounded signed state–prefix defect while the stopped all-place excess remains negative.
`hodge_cusp_derivative_tests.py`	Evaluates the three hypergeometric Schwarz maps and checks the cusp asymptotic in which the smooth three-factor derivative bound tends to 1. It also verifies that $\mathcal{M}^{-2}$ has cusp norm $L^{-1+o(1)}$, the logarithmic cotangent has norm $L^{1+o(1)}$, and their Higgs-ratio product remains bounded.
`rosen_asai_wedge_tests.py`	Constructs the exact primitive Plücker vector $\Phi(P)\wedge\Phi(gP)$; it finds strong cancellation of unit height but also shows that the formal negative nearest-wall weight does not control the positive-weight leakage of the actual vector.
`rosen_incidence_plucker_tests.py`	Constructs $\Phi(P) \wedge \Phi(w^\sharp)$, verifies that its scalar is $\mathrm{N}(ap + bq)$, and computes the outer component's exact $\mathcal{O}_K$-content ideal. On 453 random positive-step records that ideal is precisely $(ap+bq)$, so there is no hidden finite-place content to replace the Hodge estimate.
`rosen_incidence_projective_tangent_tests.py`	Tests the projectively normalized incidence $(\Phi(P) \wedge C) \otimes C^{-2}$. It verifies that the prefix cusp is the contracting Cartan line; the correct tangent weights therefore have a positive component, explaining the growth to $180.0538\ldots$ in the denominator-20 example.
`rosen_incidence_cotangent_arakelov_tests.py`	Tests $\kappa_C = \Omega(C, -) \otimes C$, whose $\mathcal{M}^2$-twisted weight is exactly $-4a_0$. Its raw operator-height proxy still grows on the balanced inverse family, while the evaluated stopped scalar stays below budget; incidence proximity cannot be discarded.
`rosen_reservoir_monotonicity_tests.py`	Tests the raw energy $\max_{k\le j} M_k + U_j$ one step at a time on the adversarial unit rays. Positive jumps are frequent and grow to $19.6300\ldots$ in the first twelve rays, ruling out an uncorrected branchwise monotonicity proof.
`rosen_record_recharge_tests.py`	Tests the total reservoir rebuilding after the most recent norm record in seven random coefficient boxes from 10 to 10^7. The largest observed recharge stays between 12.18 and 12.95; the balanced inverse family refutes the apparent additive bound.
`rosen_reservoir_transfer_tests.py`	Solves the scalar-transfer linear programs on 2387 exact unit-ray transitions. A signed Hodge correction works only by becoming unbounded below in a cusp; every globally nonnegative linear toric correction still has a positive minimax step of $10.0141\ldots$.

program	role and outcome
`rosen_norm_radical_tests.py`	Factors the adversarial record norms; the dominant records tested through $m = 16$ are squarefree, so a size-only radical bound cannot suppress the excursion.
`rosen_sunit_prime_dynamics.py`	Uses PARI prime ideals to verify the birth automaton and return-continuant divisibility; the dominant record primes are time-truncated split-prime births.
`rosen_prime_return_delay_tests.py`	Factors every remainder into actual prime ideals and measures first-return delay and returned norm mass. Only 1.65%–7.16% of total birth mass and 0%–21.13% of dominant-record birth mass returns on the twelve tested adversarial orbits; bulk return is not the missing bridge.
`rosen_plucker_omitted_boundary_tests.py`	Computes the exact prime-ideal radical after formally omitting the central Plücker boundary. The apparent surplus is negative at the eight selected norm records and large only immediately before closing; Route II.14.66 proves that the tested H-height inequality is invalid after boundary reduction.
`rosen_norm_polynomial_discriminant_tests.py`	Builds $R_j(N) = \mathrm{N}(r_{j-1} + N\lambda r_j)$ exactly and compares its discriminant with the next remainder norm. Every dominant record tested is coprime to the discriminant, so its newly occurring factors are simple residue roots rather than ramified exceptions.
`rosen_sunit_tropical_tests.py`	Finds rotating unbounded cancellation cones, disproving a dominance-only proof of the moving S-unit principle.
`rosen_preperiod_height_search.py`	Builds exact preperiods backward from special cycles and searches for multiple compounded height excursions; none occurred in the tested beam.
`rosen_preperiod_norm_search.py`	Carries primitive pairs on certified closing preperiods and separately optimizes the sextic surplus and the second running norm-record jump. Across 50,535 exact inverse edges it found no positive sextic surplus and no large compounded record jump.
`rosen_preperiod_reservoir_search.py`	Searches certified closing preperiods for the full reservoir excess and separately for excess after two later norm records. Across 204,257 exact inverse edges it found only a bounded small-height excess and no second sharp conversion; its beam score missed the balanced unit direction.
`parabolic_multiblock_tests.py`	Verifies the full-monomial factorization for 200 exact random parabolic connector families with up to five unbounded blocks.
`parabolic_dominant_block_tests.py`	Tests the moving one-block coefficient criterion. The adversarial unit ray approaches the balanced cubic boundary rather than the small-connector branch.

program	role and outcome
`rosen_pisot_rounding_tests.py`	Audits the exact balanced-ray coefficients and proves computationally the exact coordinate obstruction used in the finite-residue Pisot argument; the apparent limiting exponents are thereby certified.
`rosen_pisot_cone_tests.py`	Tests the full two-dimensional Pisot cone and scans 23948 exact admissible blocks for the rational-prefix obstruction; every observed exception is symmetric or parabolic. Its exhaustive mode also checks bounded unrestricted digit alphabets for nonintegral rational shears.
`hodge_compression_continuant_tests.py`	Verifies exactly that the parabolic Hodge pairing, squared cusp determinant, and squared Rosen return continuant are the same covariant.
`tensor_induction_obstruction_tests.py`	Linearizes exact diagonal norm evaluations on symmetric powers. The quadratic norm form has no nonzero rational diagonal zero but its linearization has a two-dimensional kernel; three random tensor-jet models likewise have large non-diagonal kernels.
`rosen_cumulative_birth_tests.py`	Factors every remainder norm on adversarial closing orbits. The cumulative new-divisor radical exceeds $6h(x)$ by hundreds or thousands, disproving an initial-height budget for all past births.
`rosen_record_birth_budget_tests.py`	Retains only first-occurrence primes born at new norm records. Their charges still exceed the final record by up to 246.94, and one record can have time-truncated charge 55.82 for a record increment of only 0.58.
`rosen_hodge_birth_potential_tests.py`	Tests whether the positive one-step drops of the two pairwise archimedean Hodge energies pay a time-truncated divisor. They leave unpaid excess up to 54.67, so the high-norm alternative must be applied at the current state rather than telescoped from raw local drops.
`rosen_normalized_sunit_abc_tests.py`	Verifies the canonical two-transition glue exactly and uses PARI/GP to factor the fractional ideals of u_j and $1-u_j$. At every high-height diagnostic state the exact abc radical exceeds $h(u_j)$ by 14.94 to 26.70, so scalar normalization alone does not create the missing surplus.
`rosen_chained_log_voyta_tests.py`	Factors the reduced union of all boundary prime ideals on a three-transition window. Its initially positive formal-adjunction surplus is exactly the omitted local different $2E_j$ from (824); it is a negative control showing that ordinary truncated Vojta on the chain remains critical.
`fixed_form_resultant_tests.py`	Verifies symbolically that the resultant of two integral fixed-point forms is $\mathrm{tr}([g,h])-2$, proves the exact consecutive digit-difference identity, and measures the exterior-square carrier on the exact $p=11, x=2$ orbit. The finite divisor fuses correctly, but the wedge-angle slopes are zero rather than negative.

program	role and outcome
`rosen_commutator_trace_audit.py`	Computes exact commutator resultants of consecutive norm-record interval words. Trace dominance collapses the outer sizes to the central one, but the central resultant can grow at the full sum of the two translation lengths (and its global norm at nearly twice that sum), so a bounded-angle shortcut is false.
`rosen_parabolic_ray_audit.py`	Tests the fixed-form exactification threshold obtained by applying trace dominance along every conjugate parabolic ray. Exact interval matrices on a 235-step closing orbit give central approximation exponents tending toward one, not the two required in degree three; conjugation to the moving frame isolates the remaining outer transported-state height.
`rosen_adjacent_parabolic_form_audit.py`	Tests the fixed quadratic obtained by multiplying the primitive parabolics at two adjacent cusp columns. Its norm exceeds the natural $6h(P)$ budget by 18.59 on an exact record state, showing that fixing the quadratic transfers rather than removes the moving-state unit-shape obstruction. It also tests the two linear factors special to $q = 7$; the small sample satisfies the constant-zero budget, while wider boxes show that an additive constant is necessary.
`rosen_adjacent_parabolic_p11_audit.py`	Runs the same fixed-form test in the exact quintic field. The form is irreducible there, and on 30 deterministic inputs its norm exceeds the natural $10h(P)$ budget by 610.96 on 1724 exact record rows.
`adjacent_torus_prime_audit.py`	Computes the discriminant norms in the odd-prime family. For $p = 5, 7, 11, 13, 17, 19, 23, 29, 31$, only $p = 7$ has square norm; the exact formula identifies these norms as the Pell numbers P_p.
`rosen_split_torus_double_coset_audit.py`	Enumerates the exact free-product normal forms through length ten and quotients by both actions of the special cyclic torus. Even with $\lvert \operatorname{Nm} \det(e, ge) \rvert = 1$, the number of observed double cosets grows from 1 to 87. This is a negative control: the unmetrized factor norm does not control the real unit directions.
`rosen_split_factor_all_states_audit.py`	Tests both orientation minima at every exact Rosen state. On 200 random closing orbits the largest excess over $6h(P)$ is -12.14; the adversarial unit rays through exponent 120 also satisfy the bound at every one of up to 3340 states.
`rosen_split_sunit_degeneracy_certificate.py`	Performs the exact 32-case calculation for degenerate consecutive split-coordinate equations over $\mathbf{Q}(n)[\lambda]$. It proves that the factor-ratio norm is nonconstant of degree at most three unless the projective state itself is one of six constant cases.

program	role and outcome
`rosen_split_sunit_return_audit.py`	Computes the conjugated intervening-word height between repeated torus types. On the sharp exponent-80 unit ray its worst ratios to the four-term unit height are $0.9059, 0.9794, 0.9852$, rejecting a little-oh slow-word claim on the critical branch.
`weak_approx_barrier.py`	Constructs exact weak-approximation witnesses proving additive height excursion is unbounded.
`rosen_norm_descent_tests.py`	Tests $\lvert \mathrm{N}\, p\rvert + \lvert \mathrm{N}\, q\rvert$; pointwise descent fails, while logarithmic record growth remains modest in the sample.
`rosen_unit_lattice_tests.py`	Exact 81^2 unit-lattice certificates and the mod–2 residue rule.
`mod2_outcome_tests.py`	The exact $1735/265$ contingency test between residual type and global certificate.
`mod2_power_orbits.py`	Exact orbit enumeration on $\mathbf{P}^1(\mathcal{O}_K/2^n)$ through $n = 6$.
`mod2_lie_certificate.py`	Exact $\mathbf{F}_2$ row-reduction certificate proving the all-n two-orbit theorem.

II.18 Pairwise Arakelov energy: the August 5 advance

The attempt to descend only the outer component $\mathcal{Z}_0$ mixed two separate questions. A metrized line over the totally real field K need not have permutation-invariant metrics at its three real places; that condition would be needed for descent to $\mathbf{Q}$, not for an Arakelov line over K. Nevertheless the three metrics must be specified before the arithmetic point and prefix are chosen. The following construction does so directly and replaces the one-component route by the sum of the two established pairwise Hodge descents.

For a primitive Rosen state $X = (u, v)$, put

$$H_i = \log \max\{|\sigma_i(u)|, |\sigma_i(v)|\}, \qquad \eta_i = H_i - \log |\sigma_i(v)|$$

and define

$$\mathscr{E}_{\mathrm{pair}}(X) = \log |\,\mathrm{N}\, v| + 2\eta_0 + \eta_1 + \eta_2. \tag{858}$$

There are three elementary identities:

$$\mathscr{E}_{\mathrm{pair}}(X) = 3h(X) + \eta_0, \tag{859}$$

$$3h(X) + \log(1 + \Theta(X)) \le \mathscr{E}_{\mathrm{pair}}(X) + \log 2, \tag{860}$$

$$\mathscr{E}_{\mathrm{pair}}(P) \le 6h(P) \quad \text{for the initial primitive state.} \tag{861}$$

For the middle inequality, use the exact log-sum-exp form

$$3h(X) + \log(1 + \Theta) = \log |\,\mathrm{N}\, v| + \log\left(e^{\eta_0+\eta_1+\eta_2} + e^{2\eta_0}\right).$$

Both exponents are at most $2\eta_0 + \eta_1 + \eta_2$. For the initial bound, the product formula for the nonzero algebraic integer v gives $-\log |v_0| \le \log |v_1| + \log |v_2|$, hence $\eta_0 \le H_0 + H_1 + H_2 = 3h(P)$.

The energy is a genuine all-place metric on one tautological K-line. On the chart $v \neq 0$, set

$$\|(u, v)\|_{\mathrm{pair},0} = \frac{\max\{|u_0|, |v_0|\}^2}{|v_0|}, \tag{862}$$

$$\|(u, v)\|_{\mathrm{pair},i} = \max\{|u_i|, |v_i|\}, \qquad i = 1, 2. \tag{863}$$

Every local norm is homogeneous of degree one under scaling by $K^\times$. With primitive metrics at finite places, the sum of their logarithms is exactly (858). The central metric has the expected logarithmic singularity along the cusp divisor $v = 0$.

There is also an exact Hodge and integral realization. In the Asai wedge splitting, let $\mathcal{Z}_i$ be the component in which V_i is alternating, and let $s = \mathrm{N}(r)$ be the scalar component for $r = wP$. If τ cycles the embeddings, then the coefficient ideals are

$$I(\mathcal{Z}_0) = (r), \qquad I(\mathcal{Z}_1) = (\tau r), \qquad I(\mathcal{Z}_2) = (\tau^2 r).$$

Consequently the evaluated tensor

$$\mathscr{P}_{w,P} = K \cdot \frac{\mathcal{Z}_1(w,P) \otimes \mathcal{Z}_2(w,P)}{s(w,P)} \tag{864}$$

has exact finite fractional ideal

$$\frac{(\tau r)(\tau^2 r)}{(r)(\tau r)(\tau^2 r)} = (r)^{-1}.$$

In the transported state frame, the two factors have relative Hodge norms

$$\eta_0 + \eta_2 + O(1), \qquad \eta_0 + \eta_1 + O(1).$$

These have the same pointwise weights as the two quantities in McMullen's pairwise contraction inequalities. Thus the nonstandard central metric in (862) is not chosen to fit the answer: it is the evaluated norm of the corresponding two Hodge components, and its integral lattice is the reciprocal remainder ideal. This is only a pointwise identification. McMullen transports one fixed flat covector, whereas the prefix row w, and therefore the evaluated tensor above, changes at every return. The displayed calculation does not prove that these changing lines contract.

The precise sufficient statement is now

$$\mathscr{E}_{\mathrm{pair}}(X_j) \le 6h(P) + B_{\mathrm{pair}} \tag{865}$$

immediately before every transition which creates a new norm record. Equations (860) and (831) then prove the sextic norm bound.

The first natural strengthening fails: the pairwise energy is not nonincreasing. On the balanced inverse-unit family its increase from the initial state is $57.1719937053\ldots$ at exponent 48. The same row still satisfies

$$\mathscr{E}_{\mathrm{pair}} - 6h(P) = -92.7052826767\ldots.$$

Thus the usable assertion is a bound by the initial sextic budget, not a Lyapunov law. The increase is unit realignment paid for by the initial arithmetic height.

The extended program `rosen_cusp_transfer_tests.py` computes this energy. The exhaustive coefficient-5, denominator-20, 120-step run contained 9862 record-producing transitions and found

$$\max\bigl(\mathscr{E}_{\mathrm{pair}} - 6h(P)\bigr) = 10.7369152971\ldots.$$

The maximum stopped envelope was $10.8197793730\ldots$, and the largest difference between that envelope and the pairwise energy was $0.6407194384\ldots < \log 2$. A smaller coefficient-3, denominator-10 grid had maximum pairwise excess $5.8357960267\ldots$. Five hundred random inputs had maximum energy increase $16.5006049166\ldots$ but negative excess over the sextic budget in that run.

Finally, the stopping condition is globally synchronized. If $\overline{\mathcal{N}}(r)$ denotes the norm line of degree $\log|\mathrm{N}\, r|$, then the maximum-norm direct sum satisfies

$$\widehat{\mu}_{\max}\left(\bigoplus_{k\le j}^{\infty}\overline{\mathcal{N}}(r_k)\right)=\max_{k\le j}\log|\mathrm{N}\, r_k|.$$

At a transition creating the first record r_{j+1}, the maximal-slope coordinate is the line with finite ideal $(r_{j+1})^{-1}$. The tensor in the pre-record state has finite ideal $(r_j)^{-1}$. These are not the same coordinate; the previous paragraph silently shifted the index. One may still seek the pre-record estimate (865) directly, because it feeds the cusp transfer inequality. Alternatively, the correctly aligned post-record target is

$$\mathscr{E}_{\mathrm{pair}}(r_j,r_{j+1})\le 6h(P)+B_{\mathrm{post}}.$$

The order of the arguments is forced by the convention that the finite term in $\mathscr{E}_{\mathrm{pair}}(u,v)$ is $\log|\mathrm{N}\, v|$. Neither target follows from the direct-sum slope identity. That identity synchronizes the finite places, but supplies no morphism from the selected record line to the changing Hodge tensor.

II.19 Theoretical correction: fixed flat line versus moving record line

The preceding construction solves the homogeneity problem but not the dynamical problem. This distinction was obscured by the numerical tests. The following calculation supersedes every sentence above which describes the evaluated prefix tensor itself as contracting.

Let $G_j\in \mathrm{SL}_2(\mathcal{O}_K)$ be a prefix matrix and let $P=(p,q)^t$. The row $\ell_P=(-q,p)$ annihilates P. Hence

$$\ell_j=\ell_P G_j^{-1},\qquad X_j=G_jP,\qquad \ell_jX_j=0.$$

This is the flat covector used in McMullen's argument: it is one integral local-system vector written in changing frames. Unimodularity preserves its coordinate-content ideal, so its finite divisor is constant.

By contrast, if $w_j=(a_j,b_j)=e_1^tG_j$, then the prefix cusp column and the remainder are

$$c_j=w_j^{\sharp}=(-b_j,a_j)^t,\qquad r_j=w_jP=\det(P,c_j).$$

This line has exactly the desired divisor, but it changes with j. Thus the two available properties live on different lines:

	flat under Hodge transport	contains (r_j)
$K\ell_j$	yes	no: its content is constant
Kc_j	no	yes

Comparing them costs the incidence determinant r_j, so a comparison which simply declares their metrics uniformly equivalent is circular.

II.19.1 The exact local loss

Normalize the geodesic to the endpoint 0 as $\gamma(t)=ie^{-t}$, and let

$$\beta_d(z)=\log\frac{|z-d|^2}{\Im z}$$

be the Busemann function centered at the zero d of a covector. Direct differentiation gives

$$\frac{d}{dt}\beta_d(\gamma(t)) = \frac{d^2e^{2t}-1}{d^2e^{2t}+1}, \tag{866}$$

$$\beta_d(\gamma(t)) - \beta_d(\gamma(0)) + t = \log\frac{1+d^2e^{2t}}{1+d^2}. \tag{867}$$

Only $d=0$, the true endpoint covector, has derivative identically -1. For a prefix cusp $d\neq 0$, the derivative becomes positive after $|d|e^t = 1$. The correction in (867) is the archimedean Green function of the incidence divisor. Before Möbius normalization its numerator is $|\det(P,c_j)| = |r_j|$. The failed proof had discarded this correction as a bounded metric comparison; it is in fact the target norm in another form.

II.19.2 The exact global identity and why Roth points the wrong way

For primitive columns $P=(p,q)^t$ and $c=(-b,a)^t$, define

$$\lambda_i(P,c) = \log\frac{\max\{|\sigma_i(p)|,|\sigma_i(q)|\}\max\{|\sigma_i(a)|,|\sigma_i(b)|\}}{|\sigma_i(\det(P,c))|}.$$

Summing this definition over the three embeddings gives the exact identity

$$\log|\,\mathrm{N}\, r_j| = 3h(P) + 3h(c_j) - \sum_{i=0}^{2}\lambda_i(P,c_j). \tag{868}$$

Consequently the desired norm estimate is equivalent to

$$\sum_{i=0}^{2}\lambda_i(P,c_j) \geq 3h(c_j) - 3h(P) - B_7 \tag{869}$$

at prefix cusps which create first norm records.

This reverses the earlier Roth perspective. Roth and the Subspace Theorem give upper bounds for exceptionally large proximity to a fixed algebraic target. Equation (869) asks for a lower bound on the total proximity of a highly correlated moving cusp: its three real incidences must absorb almost all of its height. A better approximation constant C/q^2, by itself, cannot supply this. A generic Vojta upper bound for a diagonal divisor also has the wrong direction. To use a standard Diophantine conjecture one would first have to construct a special orbit variety or Hodge incidence line on which the desired lower bound becomes a positivity statement; naming Vojta before constructing that object does not advance the proof.

II.19.3 Exact energy cocycle

The failure of monotonicity is also visible without computation. Use the state convention

$$X_j = (r_{j-1}, r_j), \qquad x_i = \frac{\sigma_i(r_j)}{\sigma_i(r_{j-1})}, \qquad y_i = \frac{\sigma_i(r_{j+1})}{\sigma_i(r_j)}.$$

With $\log^+ t = \max\{0,\log t\}$, direct use of the product formula gives

$$\begin{aligned}\mathscr{E}_{\mathrm{pair}}(X_{j+1}) - \mathscr{E}_{\mathrm{pair}}(X_j) &= 2\log|x_0| + 2\log^+|y_0| - \log|y_0| - 2\log^+|x_0| \\ &\quad + \sum_{i=1}^{2}\big(\log|x_i| + \log^+|y_i| - \log^+|x_i|\big).\end{aligned} \tag{870}$$

Equivalently,

$$\mathscr{E}_{\mathrm{pair}}(X_j) = \log|\,\mathrm{N}\, r_{j-1}| + |\log|x_0|| + \log^+|x_1| + \log^+|x_2|.$$

After the first Rosen step, $|x_0|, |y_0| < 1$, so the central term in (870) is $2\log|x_0| - \log|y_0|$. The first-record condition controls only $\sum_i \log|y_i| > 0$. It imposes no sign on the full cocycle. This proves theoretically why attempts to manufacture a one-step Lyapunov function kept failing.

II.19.4 What a correct Hodge–Arakelov route must construct

There is an explicit algebraic incidence line before any Hodge theory is used. On $\mathbf{P}^1 \times \mathbf{P}^1$, the determinant is the canonical section of $\mathcal{O}(1,1)$ vanishing on the diagonal. Pulling it back along the pair consisting of the fixed endpoint section P and the parabolic prefix section c_j gives finite divisor (r_j). Its real Green function is exactly the angular correction above. Therefore the next theoretical task is not to search for another scalar but to do the following.

1. Construct the Rosen return correspondence as an algebraic or logarithmic correspondence carrying both the flat endpoint local system and the parabolic vanishing-cycle section.
2. Pull back the incidence line $\mathcal{O}(1,1)$, equip its three archimedean factors with the modular-embedding Hodge metrics, and verify that its integral model has divisor exactly (r_j), with no index shift.
3. Prove an arithmetic degree or Hodge-index inequality whose boundary term is (867) and whose stopped conclusion is (869). The stopping condition must select one global record line; it may not maximize separately at different places.
4. Invoke the four-term S-unit theorem only after this estimate has bounded the norms. The S-unit theorem supplies closing from fixed support, but contains no mechanism for proving the incidence lower bound.

One may call the resulting extension of the flat endpoint line by the parabolic line an incidence or biextension line. The name is secondary. The required content is concrete: its finite divisor must be (r_j), its archimedean Green function must be (867), and its curvature must yield (869) at first records. This is the point at which new Hodge theory could enter.

II.19.5 A relative-cohomology construction of the moving cusp

The parabolic line which changes with the prefix is not an arbitrary moving target. It is generated by one canonical relative cohomology class. Let Γ_∞ be the stabilizer of the standard cusp and choose a primitive column $e \in \mathcal{O}_K^2$ fixed by Γ_∞. Define

$$\kappa(g) = g^{-1}e - e. \tag{871}$$

Then

$$\kappa(gh) = h^{-1}\kappa(g) + \kappa(h), \qquad \kappa|_{\Gamma_\infty} = 0.$$

Thus κ represents the image of $e \in H^0(\Gamma_\infty, K^2)$ under the connecting map

$$H^0(\Gamma_\infty, K^2) \longrightarrow H^1(\Gamma, \Gamma_\infty; K^2).$$

It is nonzero as a relative class precisely because the parabolic invariant vector does not extend to a Γ-invariant vector.

For a prefix G_j, the cusp column is $c_j = G_j^{-1}e = e + \kappa(G_j)$. Pairing with the fixed endpoint covector ℓ_P gives the exact formula

$$r_j = \ell_P(c_j) = \ell_P(e) + \ell_P(\kappa(G_j)). \tag{872}$$

The remainder sequence is therefore a sequence of periods of the cusp extension class, evaluated against one fixed integral covector. This is the first construction which explains both sides of the flat-line/record-line dichotomy: ℓ_P is flat, while all motion of the record line is carried by the single relative class κ.

For a weight-minus-one variation, an extension of the trivial variation by the homology variation is the Hodge-theoretic datum underlying a normal function. Here the absolute extension defined by (871) is split; what is nontrivial is the splitting framed at the cusp. The natural candidate is consequently a *framed relative biextension*: pair the cusp-framed extension $\delta(e)$ with the flat dual line $K\ell_P$. Its expected local height is the Green function in (867), and its integral evaluation is (872).

For general P, the covector line $K\ell_P$ is horizontal on the universal cover but is not a section of the triangle quotient. The correct algebraic base for a universal construction is therefore

$$\mathcal{Y} = \mathbf{P}(\mathcal{V}^\vee) \longrightarrow C,$$

the projective flat bundle of the homology variation, equipped with its tautological covector line. Each endpoint gives a horizontal real leaf of $\mathcal{Y}$. The parabolic invariant e defines a boundary incidence divisor in the compactified bundle, and the pairing of the tautological line with $\delta(e)$ is the universal candidate normal function. The biextension must be constructed on $\mathcal{Y}$ before restricting it to the leaf belonging to P; otherwise one has silently treated a non-descending endpoint covector as an algebraic section of C.

There are genuine positivity theorems for biextension metrics and their boundary height jumps: under the hypotheses of an admissible normal function, the relevant jump divisor is effective [65, 66]. Those theorems do not yet prove the required record inequality. Three identifications must be made first:

1. realize the relative class $\delta(e)$, with its cusp framing, as an admissible normal function for the heptagonal variation or for its weight-three Asai realization;
2. prove that the resulting biextension metric is the pairwise metric, including the extra central logarithmic singularity and the finite evaluation ideal (r_j); and
3. show that effectiveness of its height jump yields the stopped lower bound (869), rather than an all-state monotonicity statement already disproved by the inverse-unit family.

This route is therefore concrete but incomplete. Its advantage over an unspecified "Arakelov height" is that the extension class, the finite period, and the archimedean Green function are now all explicit.

There is, however, a further structural limitation which is easy to miss. It rules out using positivity of the relative biextension by itself.

Proposition II.19.1 (The cusp extension is supported at the boundary)**.** *The image of the relative class represented by κ in ordinary cohomology $H^1(\Gamma, K^2)$ is zero. Equivalently, the associated extension of local systems on the open triangle curve is split. Hence a biextension formed from this class alone has no interior curvature capable of paying the two outer Hodge defects. Its nontrivial information is the cusp framing and the resulting boundary Green function.*

Proof. In ordinary group cohomology,

$$\kappa(g) = g^{-1}e - e$$

is the coboundary of the zero-cochain e. It is nonzero only after the splitting is required to agree with the fixed vector e on Γ_∞. Thus its class can be nonzero in $H^1(\Gamma, \Gamma_\infty; K^2)$, but its image in absolute cohomology is zero. Changing from the cusp-framed splitting to the absolute splitting trivializes the extension on the open curve.

The metric calculation says the same thing. If β_d and β_0 are the two Busemann functions used in (866)–(867), then away from their boundary zeros

$$\beta_d(z) - \beta_0(z) = \log|z-d|^2 - \log|z|^2$$

is the logarithm of the absolute square of a meromorphic function. Its dd^c is zero in the interior. It records the change of boundary framing, not a new negative curvature term. □

This proposition sharpens the construction problem. The relative class supplies the correct finite period r_j and the exact angular correction, but it cannot supply strict descent through the thick part. The latter must come from a genuinely nontrivial Hodge line, such as $\mathcal{M} = \mathcal{L}_1\mathcal{L}_2\mathcal{L}_0^{-1}$, or from the equivalent evaluated log-conormal in the Asai variation. A successful object must therefore be a *Hodge-twisted relative incidence line*: the relative factor fixes the divisor and the angle, while the Hodge factor provides curvature. Citing height-jump effectivity before making this twist cannot prove the stopped inequality.

II.19.6 The exact negative line and the common source of the failed arguments

The required Hodge twist is more precise than the character $\mathcal{M}$ by itself. Define

$$\mathcal{N} = \mathcal{L}_1 \otimes \mathcal{L}_2 \otimes \mathcal{L}_0^{-2} = \mathcal{M} \otimes \mathcal{L}_0^{-1}. \tag{873}$$

It is the line whose weight is the sum of McMullen's two pairwise weights. Consequently

$$\deg \mathcal{N} = \frac{3}{28} + \frac{1}{28} - 2\frac{5}{28} = -\frac{3}{14}. \tag{874}$$

Along a central geodesic the corresponding logarithmic norm has derivative at most

$$-2 + \|DF_1\|_{\mathrm{hyp}} + \|DF_2\|_{\mathrm{hyp}} < 0. \tag{875}$$

This inequality is pointwise in the interior. It becomes asymptotically neutral only during a pure parabolic excursion. It is therefore stronger than the negative total degree and is exactly the doubled-central contraction demanded by the cubic norm.

There is also an exact algebraic check. Since the uniformizing summand is maximal Higgs, $\mathcal{L}_0^2 \simeq \Omega^1_{\log}$. Hence the Higgs-ratio section has bundle

$$\Xi \in H^0(\Omega^1_{\log} \otimes \mathcal{M}^{-2}) = H^0(\mathcal{N}^{-2}). \tag{876}$$

Its divisor is $3P_7$, where P_7 is the order-seven orbifold point. It follows that

$$\mathcal{N}^{-2} \simeq \mathcal{O}(3P_7), \qquad \mathcal{N}^2 \simeq \mathcal{O}(-3P_7). \tag{877}$$

The orbifold degree of $3P_7$ is $3/7 = -2\deg\mathcal{N}$, as it must be.

This identifies the single obstruction against which the preceding computations repeatedly collided. There are two mathematically correct objects:

	finite divisor (r_j)	strict interior curvature
cusp-framed incidence line $\mathcal{N}$	yes no flat integral lattice	no yes.

The theorem requires their tensor product as one integral evaluated section on the Rosen return correspondence. The scalar substitutes failed because they used the incidence divisor at finite places and the $\mathcal{N}$-metric at real places without retaining a common ambient state. The normalized tensor (800) now supplies that state: it has fractional content $(r_j)^{-1}$, pairwise norm (801), and exact transport (802). The unbounded discrepancy has consequently been reduced to the explicit scalar cocycle r_{j+1}/r_j. Unit realignment, positive-weight leakage, and the former index shifts are three manifestations of this one cocycle, not independent phenomena.

The construction now meets the three structural requirements for a line $\widehat{\mathcal{I}}_P$ on the stopped Rosen correspondence:

1. its integral evaluation at the j-th prefix has fractional ideal $(r_j)^{-1}$ (or its square, with indices stated explicitly);
2. its archimedean metric is the incidence Green metric tensored with the metric of $\mathcal{N}$; and
3. its zero-weight specialization retains the full Rosen state, so that repetition gives a hyperbolic fixed line rather than merely a repeated scalar.

The remaining assertion is no longer construction of the object. It is the stopped upper bound for its Arakelov norm.

II.19.7 The formal line and its higher-rank integral realization

One can specify the line which would carry the proof, but specifying its two factors is not yet a construction of an arithmetic transport map. On the universal projective flat bundle

$$\widetilde{\mathcal{Y}} = \mathbb{H} \times \mathbf{P}(V^\vee)$$

let D_c denote the hyperplane $\ell(c) = 0$. Fix the primitive cusp vector e. A prefix g selects $c_g = g^{-1}e$, and

$$s_g(\ell) = \frac{\ell(c_g)}{\ell(e)} \tag{878}$$

is the canonical meromorphic section of $\mathcal{O}(D_{c_g} - D_e)$. Replacing the prefix by a representative differing by the peripheral stabilizer does not change this section, because that stabilizer fixes e. On the horizontal leaf determined by P, equation (872) gives

$$s_g(\ell_P) = \frac{r_g}{r_0}. \tag{879}$$

Thus its finite divisor is $(r_g) - (r_0)$, with no moving-prime qualification and no hidden index shift. The denominator is fixed and is charged to the initial arithmetic height.

Write p_1, p_2 for the source and target maps on a branch of the Rosen return groupoid. The formal candidate is

$$\widehat{\mathcal{I}}_g = \mathcal{O}(D_{c_g} - D_e) \otimes p_2^*\mathcal{N} \otimes p_1^*\mathcal{N}^{-1}. \tag{880}$$

The first factor has the exact finite evaluation (879). With its relative Green metric it has zero curvature off the two incidence divisors, by Proposition II.19.1. The fiber ratio of $\mathcal{N}$ has the desired Hodge curvature only after one chooses a comparison between its source and target fibers. No such comparison is supplied by the displayed tensor product.

Indeed, interpreting the arrow as a deck transformation makes the equivariant Hodge identification an isometry and loses the contraction. Interpreting it as geodesic transport gives a different problem: the Gauss–Manin connection transports the ambient flat bundle, not its Hodge subline. Orthogonal projection between the Hodge lines is analytic and has period-valued coefficients; an integral structure on its determinant has not been constructed. This is the same flat-line/record-line problem in a more invariant form.

Inside the one-dimensional associated graded, one could still ask for an arithmetically normalized transport section

$$\tau_g \in p_2^*\mathcal{N} \otimes p_1^*\mathcal{N}^{-1}, \qquad \mathfrak{t}_g = s_g \otimes \tau_g \in \widehat{\mathcal{I}}_g, \tag{881}$$

whose combined finite evaluation is r_g/r_0 and whose archimedean norm is the pairwise Schwarz–Pick defect. It is unnecessary to project to that line. The higher-rank tensor (800) is integral before projection, has the same extremal Hodge weight, and obeys the transport law (802). Thus it supplies the finite system which plays the role of $\mathfrak{t}_g$ while retaining the full state.

The remaining statement is the following arithmetic maximum principle.

Question II.19.2 (Stopped arithmetic maximum principle)**.** *At every arrow g which creates a first norm record, prove that the Arakelov degree of the evaluation of $\widehat{\mathcal{I}}_g$ at ℓ_P is bounded above by its initial degree plus $O_7(1)$. Prove moreover that a zero-weight infinite family forces repetition of the full reduced Rosen state, not merely repetition of the scalar r_g.*

The first sentence is equivalent, after the established metric identifications, to the pairwise stopping inequality. The second follows from state separation for $\mathbf{U}$ once boundedness has been proved. The scalar section alone would not determine the state; the normalized tensor does.

There is also a weight-three interpretation. The same relative incidence data has a realization in the tensor-induced Asai variation of Hodge type $(1,3,3,1)$. On a rank-two Asai secant the hyperdeterminant is $\prod_i \det(P_i, c_{j,i})^2 = |\,\mathrm{N}\, r_j|^2$. Its logarithmic conormal has coefficient $1/\sigma_i(r_j)$ along variation of the i-th secant factor; the distinguished factor is the reciprocal ideal $(r_j)^{-1}$ occurring in the pairwise tensor. Thus the pairwise line can be viewed canonically as a component of the logarithmic conormal to the hyperdeterminant divisor, not as an independently invented asymmetric metric. What remains is still the same relative transport problem: the secant endpoint c_j changes by the cocycle κ.

II.19.8 Why the expected alternatives are negative and zero Hodge weights

There is a simple integral-lattice argument behind the cusp/special dichotomy. It is useful because it identifies exactly what the missing Hodge twist has to accomplish.

Proposition II.19.3 (Integral nonpositive-weight dichotomy)**.** *Let $X = \mathbb{H}/\Gamma$ have finite hyperbolic area, let $\mathbb{V}_{\mathbb{Z}}$ be an integral local system with an adapted metric, and suppose a boundary point p determines a nonpositive subspace $E^{\le 0}(p)$ in the fiber at a base point. Assume the following two properties.*

1. *If the geodesic to p returns to a fixed compact set at times t_n, then the reduced transports of every $v \in E^{\le 0}(p)$ have uniformly bounded norm; for $v \in E^{<0}(p)$ their norms tend to zero.*

2. *The pair consisting of a nonzero integral vector in the transported nonpositive subspace and the reduced boundary state determines that state: if the same vector occurs at two returns, then the two reduced endpoints are equal.*

   Then a nonzero integral vector in $E^{<0}(p)$ forces p to be a parabolic endpoint. A nonzero integral vector in $E^{\le 0}(p)$, satisfying the state-separation hypothesis, forces either the same conclusion or a hyperbolic element of Γ to fix p. If $p \in \mathbf{P}^1(K)$ and $\Gamma \subset \mathrm{SL}_2(K)$, that hyperbolic element is special.

Proof. If p is not parabolic, its geodesic has infinitely many returns to a fixed compact part of X. On that compact set every nonzero integral vector has norm bounded below. A vector in $E^{<0}(p)$ would have norm tending to zero at these returns, a contradiction.

For a vector in $E^{\le 0}(p)$, all reduced transports lie in the intersection of a fixed lattice with a bounded ball. This set is finite, so two reduced transports are equal. By state separation the corresponding reduced endpoints are also equal. The group element between the two returns fixes p. It is nontrivial; it cannot be elliptic because it fixes a real boundary point, and if it is parabolic then p is a cusp. In the remaining case it is hyperbolic. For $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ fixing $p \in K$, one eigenvalue is $cp + d \in K$, and the other is its inverse. □

Filip's integral-vector theorem is the strict negative-weight part of this argument. The zero-weight extension explains the second alternative in the present problem: bounded integral return data repeats, and repetition is a special eigenline. The heptagonal difficulty is now visible in one symbol. For the untwisted Asai point the dangerous weight is

$$\chi = -a_0 + a_1 + a_2,$$

which can be positive. Neither lattice discreteness nor a standard log-Anosov gap applies to that component. The record-incidence/Hodge twist must move the *evaluated* section, including its ideal $(r_j)^{-1}$, into the nonpositive bundle. Once that is done, the proposition supplies the cusp and special alternatives without a further Diophantine approximation theorem.

II.19.9 Audit of the arithmetic-holonomy route

The arithmetic holonomy bounds of Calegari–Dimitrov–Tang offer a possible way to construct the missing auxiliary sections, but their logical input must be stated accurately. Their theorem bounds the dimension of a space of holonomic power series having a fixed denominator type and admitting analytic continuation through a fixed holomorphic map ϕ. It does not, as stated, give a word-dependent moving-target Roth theorem.

An application here would require three new identifications. For every return word g, one would need a continuation domain and a map ϕ_g for the $(2, 7, \infty)$ hypergeometric system; one would need to prove that

$$\log|\phi_g'(0)| - \tau(\text{denominator type}) = d\big(a_0(g) - \max\{a_1(g), a_2(g)\}\big) + O_7(d);$$

and one would need to identify failure of the resulting auxiliary system to have full rank with the parabolic incidence or fixed-eigenline locus. McMullen's cusp compression is a plausible source of the conformal term, while the hypergeometric G-function coefficients provide the denominator term. None of these three identifications follows formally from the published holonomy bound. This route is therefore worth developing, but it is not presently a proof and no numerical test of word matrices can replace the missing denominator identity.

II.19.10 A relative-dominated representation for the outer gap

The two Hodge contractions also have a concrete representation-theoretic package. Let $\rho_i = \sigma_i \circ \rho$ be the three real rank-two representations and consider the block representation

$$\rho_\star = \operatorname{Sym}^2 \rho_0 \oplus (\rho_1 \otimes \rho_2). \tag{882}$$

If $a_i(g) \geq 0$ is the logarithm of the largest singular value of $\rho_i(g)$, the top exponents of the two blocks are

$$2a_0(g), \qquad a_1(g) + a_2(g).$$

Their difference

$$\Delta_\star(g) = 2a_0(g) - a_1(g) - a_2(g) = (a_0 - a_1) + (a_0 - a_2) \tag{883}$$

is exactly the sum of the two modular-embedding gaps.

This character has the correct cusp behavior. For a parabolic shear T^n, each $a_i(T^n) = \log|n| + O_K(1)$, and therefore $\Delta_\star(T^n) = O_K(1)$. Compact returns, on the other hand, supply positive integrated defects for each of the two pairwise Hodge cocycles by strict Schwarz–Pick contraction. Converting those defects into the singular-value expression (883), with controlled endpoint angles, is part of the proposed theorem. Thus (882) is a candidate *relatively dominated* cocycle: long cusp motion is neutral, while motion through the thick part opens a block gap. This is the precise connection with log-Anosov or cusped-Anosov structures; Filip's results provide the model for such a conclusion, but his theorem is not being asserted here for the present rank-eight Asai variation.

An operator gap alone is insufficient. The desired scalar is the logarithmic conormal evaluated on the endpoint–cusp secant, and the operator-norm experiments above show that evaluation angles can carry an unbounded error. What would suffice is an angle-refined relative domination theorem with the following three features:

1. its attracting line is the central square boundary line and its opposing block is the outer tensor boundary line;
2. its transversality term is the incidence Green function in (868); and
3. its equality trajectories are exactly parabolic boundary axes or unit eigenlines, which give the cusp and special alternatives.

Combined with the relative cocycle κ, such a theorem would supply the archimedean transport law for the same incidence line whose finite period is r_j. Proving only $\Delta_\star(g) \geq c\,|g|_{\mathrm{rel}} - C$ would not suffice, because it still discards the evaluated angle. The useful target is domination of the logarithmic conormal line itself.

II.19.11 What the computations did and did not establish

The numerical data above remain useful for rejecting false monotonicity, state–prefix comparisons, and operator-norm replacements. They also check the algebraic identities and recurrence implementation. They do not test whether a prefix-dependent tensor is flat under the Gauss–Manin connection. Consequently the finite maxima reported for $\mathscr{E}_{\mathrm{pair}} - 6h(P)$ provide no evidence for the missing Hodge transport. The obstruction in this section is exact and survives every finite computation.

II.20 The rank-two correction: what the thin group actually supplies

The preceding scalar construction still compresses too much information. This section records a correction found by returning to the full triangle group representation. It changes the proposed closing object from one determinant line to a rank-two relative object.

II.20.1 Why the determinant cannot control the units

For a unit $\epsilon \in \mathcal{O}_K^\times$, put

$$t_i(\epsilon) = \log|\sigma_i(\epsilon)|.$$

Dirichlet's theorem places the logarithmic unit lattice in the plane $t_0+t_1+t_2 = 0$. The two pairwise characters are

$$(u_1, u_2) = (t_1 - t_0, t_2 - t_0). \tag{884}$$

They recover all three logarithms:

$$t_0 = -\frac{u_1+u_2}{3}, \qquad t_1 = \frac{2u_1-u_2}{3}, \qquad t_2 = \frac{-u_1+2u_2}{3}.$$

Thus (884) is proper on the rank-two unit lattice.

The determinant character used above is only the sum

$$u_1 + u_2 = -3t_0. \tag{885}$$

It is not proper. Indeed, for independent units ϵ, η, the ratio $\log|\sigma_0\epsilon|/\log|\sigma_0\eta|$ is irrational. Otherwise a nontrivial product of the two units would have distinguished conjugate equal to ± 1, and hence would itself be ± 1. Therefore integer linear combinations of the two distinguished logarithms approach zero along unbounded exponent vectors. Along this sequence (885) stays bounded while the pair (884) diverges.

This is an exact explanation for the repeated "unit realignment" in the failed scalar experiments. A bound for the Arakelov degree of $\mathcal{N}$ controls only the sum of two slopes. It permits one slope to tend to plus infinity and the other to minus infinity. A scalar proof would therefore need an additional theorem saying that escape in this kernel is periodic unit holonomy; sharpening the scalar bound alone does not supply that theorem.

There is no hidden flat rational line which retains either slope. Put $H = \mathrm{Res}_{K/\mathbf{Q}}\,\mathrm{SL}_2$. Over a splitting field its diagonal torus has weight lattice $\mathbb{Z}^3$, with the cyclic Galois group permuting the coordinates. The weights

$$\chi_1 = (-1,1,0), \qquad \chi_2 = (-1,0,1), \qquad \chi_1+\chi_2 = (-2,1,1)$$

are asymmetric and have three-element Galois orbits.

Proposition II.20.1 (No flat rational pairwise line)**.** *No nontrivial one-dimensional* $\mathbf{Q}$*-subquotient of an algebraic representation of* H *carries any of the three displayed weights. The smallest rational flat package retaining one of them contains its full cyclic Galois orbit.*

Proof. The connected group H is semisimple and therefore has no nontrivial algebraic character. Thus every one-dimensional subquotient of an algebraic H-representation has trivial weight. Already on the split torus, a one-dimensional weight defined over $\mathbf{Q}$ would have to be fixed by the cyclic Galois action, which none of the three weights is. □

This proposition separates two roles which had repeatedly been conflated. The pairwise line exists after the Hodge filtration reduces the flat H-bundle to a Borel; its nonzero degree is its curvature. The integral lattice exists before that reduction and is necessarily Galois stable. The three Asai components below are not an ornamental enlargement: they are the smallest concrete package in the present construction which can carry integrality and the two asymmetric Hodge slopes simultaneously.

II.20.2 Why ordinary principal-bundle stability is circular here

Let $C = G_7\backslash\mathbb{H}$, and let $\mathscr{F}$ be the flat bundle over C with fiber the product of the three real flag varieties. The Hodge flag

$$z \longmapsto (z, F_1(z), F_2(z))$$

is equivariant, so it is a genuine Borel reduction over C. For $x \in \mathbf{P}^1(K)$, the constant rational flag $\Delta_K(x) = (\sigma_0 x, \sigma_1 x, \sigma_2 x)$ descends only to

$$C_x = \mathrm{Stab}_{G_7}(x)\backslash\mathbb{H}.$$

This follows directly from the descent condition: a deck transformation preserves the constant section precisely when it fixes x.

The stabilizer of a real boundary point in a Fuchsian group is trivial, parabolic cyclic, or hyperbolic cyclic. Hence the desired classification is precisely the assertion that C_x is a cusp or hyperbolic cylinder, and never the universal cover $\mathbb{H}$. Standard Ramanathan or Harder–Narasimhan stability on the finite-area base cannot compare the two reductions, because the rational one is not present on that base. In the case to be excluded it is present only on the universal cover.

This diagnoses the exact additional theorem required by the principal bundle idea. It must be a boundary or stopped Hilbert–Mumford theorem on C_x, not an ordinary degree inequality on C. Its generalized minors are the incidences $r_g = \det(P, g^{-1}e)$; its two numerical characters are χ_1, χ_2; and its equality cases must say that the deck group of C_x is parabolic or hyperbolic cyclic. Any argument which begins by assigning a global degree on C to the rational Borel has already assumed the descent, and therefore the conclusion.

II.20.3 The order-seven monodromy splits the defect

Put $R = ST$. Projectively the full group has the presentation

$$G_7 = \langle S, R \mid S^2 = R^7 = 1\rangle \simeq C_2 * C_7. \tag{886}$$

This presentation, together with the Galois conjugate representations, is the first group-wide structure used here which cannot be reduced to a list of generator inequalities.

There is also an indexing correction. With

$$\sigma_1(\lambda) = 2\cos(3\pi/7), \qquad \sigma_2(\lambda) = -2\cos(2\pi/7),$$

the first outer representation has Lyapunov exponent $1/5$, while the second has exponent $3/5$. Consequently the Hodge degrees, with the embeddings carrying the same indices, are

$$\deg \mathcal{L}_0 = \frac{5}{28}, \qquad \deg \mathcal{L}_1 = \frac{1}{28}, \qquad \deg \mathcal{L}_2 = \frac{3}{28}. \tag{887}$$

The earlier ordering of the last two entries did not affect their sum, but it obscured the individual geometry.

There is a purely group-theoretic reading of the same three numbers. With the positive cusp lift fixed, the rotation numbers of S are $1/2$, while the rotation numbers of R in the three signed representations are $1/7, 3/7, 2/7$. The relative Euler-number formula for a $(2,7,\infty)$ representation gives

$$e(\rho_0) = \frac{5}{14}, \qquad e(\rho_1) = \frac{1}{14}, \qquad e(\rho_2) = \frac{3}{14}. \tag{888}$$

The Hodge-line degree is half the corresponding Euler number, recovering (887). In particular

$$2e(\rho_0) - e(\rho_1) - e(\rho_2) = \frac{3}{7}. \tag{889}$$

This equals the degree of the effective divisor $3P_7$ below. Thus the complete scalar charge is a relative bounded-Euler-class defect of the three full representations. A charged letter only records where a chosen tessellation crosses a representative of this global class.

This observation also sets a limit. An Euler class controls periods on closed surfaces or closed orbits. The rational flag does not descend to the compact triangle quotient unless it is stabilized, as Proposition II.7.16 makes explicit. Consequently the Euler defect explains the Hodge surplus and the order-seven equality divisor, but cannot by itself supply the stopped fixed-flag transport.

Define the two negative lines

$$\mathcal{N}_i = \mathcal{L}_i \otimes \mathcal{L}_0^{-1}, \qquad i = 1, 2.$$

The ratios of the Higgs fields give

$$\Xi_i = \frac{\theta_i}{\theta_0} \in H^0(C, \mathcal{N}_i^{-2}).$$

At the order-seven point, the σ_1-Schwarz map has local degree three and the signed σ_2-Schwarz map has local degree two. Hence

$$\operatorname{div}(\Xi_1) = 2P_7, \qquad \operatorname{div}(\Xi_2) = P_7. \tag{890}$$

The degree check is exact:

$$\deg \mathcal{N}_1^{-2} = \frac{2}{7}, \qquad \deg \mathcal{N}_2^{-2} = \frac{1}{7}.$$

Thus the previously found divisor $3P_7$ is the determinant of two unequal global monodromy defects. Keeping only their product is exactly the scalar collapse in (885).

There is a useful cyclic completion of this calculation. Read indices modulo three and define

$$\mathcal{Q}_j = \mathcal{L}_{j+1} \otimes \mathcal{L}_{j+2} \otimes \mathcal{L}_j^{-2}, \qquad \Upsilon_j = \frac{\theta_{j+1}\theta_{j+2}}{\theta_j^2}.$$

Then Υ_j is a meromorphic section of $\mathcal{Q}_j^{-2}$. Since the Higgs divisors are $0, 2P_7, P_7$, direct subtraction gives

$$\operatorname{div}(\Upsilon_0) = 3P_7, \qquad \operatorname{div}(\Upsilon_1) = -3P_7, \qquad \operatorname{div}(\Upsilon_2) = 0, \qquad \prod_j \Upsilon_j = 1. \tag{891}$$

Equivalently,

$$\mathcal{Q}_0^2 \simeq \mathcal{O}(-3P_7), \qquad \mathcal{Q}_1^2 \simeq \mathcal{O}(3P_7), \qquad \mathcal{Q}_2^2 \simeq \mathcal{O}. \tag{892}$$

Their degrees are $-3/14, 3/14, 0$, exactly the three cyclic degrees of the doubled-central pairwise state.

The zero-degree assertion can be sharpened. The orbifold characters of $\mathcal{Q}_2 = \mathcal{L}_0\mathcal{L}_1\mathcal{L}_2^{-2}$ vanish. At the order-two point all three systems have the same character; at the order-seven point the rotation labels give $1+3-2\cdot 2 = 0 \pmod 7$; and the unipotent cusp residues cancel. Together with

$$\deg \mathcal{Q}_2 = \frac{5}{28} + \frac{1}{28} - 2\frac{3}{28} = 0$$

this proves the unsquared identity

$$\mathcal{Q}_2 \simeq \mathcal{O}_C. \tag{893}$$

The Hodge metric on this trivial line is bounded relative to the constant metric. Indeed, near the cusp its squared norm is comparable to

$$\frac{\operatorname{Im} z \, \operatorname{Im} F_1(z)}{(\operatorname{Im} F_2(z))^2},$$

which tends to a positive constant by the signed cusp expansions. It is smooth on the rest of the compact orbifold. Therefore $\Theta_1^{\mathrm{Hod}} - 2\Theta_2^{\mathrm{Hod}}$ is a bounded coboundary on reduced returns.

The fixed elliptic divisor also has a completely algebraic orbit coordinate. Put $\eta = e^{\pi i/7}$, so that $\eta + \eta^{-1} = \lambda$, and define

$$Q_7(u,v) = u^2 - \lambda uv + v^2 = \mathrm{Nm}_{K(\eta)/K}(u - v\eta).$$

For $R = ST$, right multiplication by R multiplies $u - v\eta$ by η; hence $Q_7(XR) = Q_7(X)$. If $\zeta = -\eta^{-1}$ is the order-seven fixed point and $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, set

$$D_g = c^2 - \lambda cd + d^2, \qquad E_g = 2ac - \lambda(ad + bc) + 2bd.$$

Then

$$\operatorname{Re}(g\zeta) = \frac{E_g}{2D_g}, \qquad \operatorname{Im}(\sigma_i(g)\zeta_i) = \frac{\sqrt{4 - \lambda_i^2}}{2\sigma_i(D_g)}.$$

For $x = p/q$ the horizontal incidence integer

$$A_g(P) = 2D_g p - E_g q$$

satisfies the CM factorization

$$A_g(P) - q(\eta - \eta^{-1}) = 2(c - d\eta^{-1})\big((pc - qa) - (pd - qb)\eta\big). \tag{894}$$

Indeed, if $X_g = a - b\eta$, $Y_g = c - d\eta$, and $Z_g = pY_g - qX_g$, then $X_g\overline{Y_g} - \overline{X_g}Y_g = \eta - \eta^{-1}$, because $ad - bc = 1$, and the displayed factorization follows by expanding $2\overline{Y_g}Z_g$.

There is a useful way to see exactly where the thin group enters this factorization. Since $\eta = -\zeta_7^4$, one has $\mathcal{O}_L = \mathcal{O}_K[\eta]$. For every $g \in \mathrm{SL}_2(\mathcal{O}_K)$, the ordered pair

$$\mathcal{B}_g = (X_g, Y_g) = (a - b\eta, c - d\eta) \tag{895}$$

is an $\mathcal{O}_K$-basis of $\mathcal{O}_L$, and its orientation is the exact identity

$$X_g\overline{Y_g} - \overline{X_g}Y_g = \delta, \qquad \delta = \eta - \eta^{-1}. \tag{896}$$

Conversely, every ordered basis with this orientation comes from a unique element of $\mathrm{SL}_2(\mathcal{O}_K)$. This is simply the determinant calculation in the basis $(1, \eta)$. Moreover the linear identity

$$(u,v)R\begin{pmatrix} 1 \\ -\eta \end{pmatrix} = \eta(u - v\eta)$$

gives

$$\mathcal{B}_{gR} = (\eta X_g, \eta Y_g). \tag{897}$$

Thus the full arithmetic group parametrizes all oriented CM bases, whereas G_7 selects one thin suborbit, considered modulo simultaneous multiplication of both basis vectors by η. The point $g\zeta = X_g/Y_g$ is the projectivization of this basis, and $Z_g = pY_g - qX_g$ has the primitive coordinate vector $(-q, p)$ in it.

This distinction is important. The coprimality and norm identities below are properties of every oriented CM basis and therefore belong to the ambient algebraic-number-theoretic space. The classification problem is the much sharper assertion for the particular thin suborbit cut out by the Teichmüller curve. A calculation with the six Rosen moves sees only a choice of coordinates on that suborbit; it does not use its global distribution among all CM bases.

Taking the relative CM norm gives the two exact identities

$$\prod_i \frac{|\sigma_i(x) - \mathrm{Re}(\sigma_i(g)\zeta_i)|}{\mathrm{Im}(\sigma_i(g)\zeta_i)} = \frac{|\,\mathrm{N}\,A_g(P)|}{\sqrt{7}\,|\,\mathrm{N}\,q|}, \tag{898}$$

$$A_g(P)^2 + (4-\lambda^2)q^2 = 4D_g Q_7(pc - qa, pd - qb). \tag{899}$$

Here $\mathrm{N}(4-\lambda^2) = 7$, and $(pc - qa, pd - qb) = (-q, p)g$ is the full remainder row. The second identity is the determinant-one Gram identity for the polar form of Q_7.

This is a useful whole-group reformulation: the order-seven charge is the orbit $G_7\zeta$, its local equation along a rational vertical leaf is the single algebraic integer $A_g(P)$, and it is joined to the complete remainder state by (899). It does not label primes or letters.

There is a stronger ideal-theoretic consequence which makes the full thin group visible in every good finite quotient. Put $L = K(\eta)$, let a bar denote CM conjugation, and set $\delta = \eta - \eta^{-1}$. If (p, q) is primitive, then

$$(Y_g, \overline{Y_g}) \mid (\delta), \qquad (Z_g, \overline{Z_g}) \mid (\delta), \qquad (Y_g, Z_g) \mid (q\delta). \tag{900}$$

The first and third assertions follow respectively from

$$X_g\overline{Y_g} - \overline{X_g}Y_g = \delta, \qquad \overline{Y_g}Z_g - Y_g\overline{Z_g} = -q\delta.$$

For the second, write $Z_g = r_g - s_g\eta$, where $(r_g, s_g) = (-q, p)g$. This row is primitive. Away from (δ), a common divisor of Z_g and $\overline{Z_g}$ would divide both s_g and r_g.

It follows that every good prime occurring in either Y_g or Z_g splits in L/K: an inert prime is fixed by CM conjugation and would contradict (900), while the ramified primes divide δ. At a split prime $\mathfrak{p} = \mathfrak{P}\overline{\mathfrak{P}}$ not dividing $q\delta$, the two divisibilities are exactly

$$\mathfrak{P} \mid Y_g \iff g\zeta = \infty \pmod{\mathfrak{P}}, \qquad \mathfrak{P} \mid Z_g \iff g\zeta = x \pmod{\mathfrak{P}}. \tag{901}$$

Indeed,

$$g\zeta = X_g/Y_g, \qquad Z_g = Y_g(p - qg\zeta).$$

For almost every $\mathfrak{p}$, strong approximation identifies the reduction of G_7 with $\mathrm{SL}_2(\mathcal{O}_K/\mathfrak{p})$. Hence (901) consists of two complete and distinct Borel cosets in the full reduced group. This is not a generator calculation: it uses the entire congruence image of the thin group and the full order-seven orbit.

The finite representation selected by these two incidences is also exact. If $k_\mathfrak{p} = \mathcal{O}_K/\mathfrak{p}$, $Q_\mathfrak{p} = |k_\mathfrak{p}|$, and

$$e_y = \delta_y - \frac{1}{Q_\mathfrak{p} + 1}\mathbf{1} \in \ell^2(\mathbb{P}^1(k_\mathfrak{p})),$$

then the e_y span the Steinberg representation and

$$\mathbf{1}_{\{g\zeta=y\}} = \frac{1}{Q_{\mathfrak{p}}+1} + \langle \mathrm{St}_{\mathfrak{p}}(g)e_\zeta, e_y\rangle. \tag{902}$$

Here all entries on the right are reduced modulo the chosen split prime. The two factors Y_g, Z_g are integral lifts of (902) for $y=\infty, x$. Thus the Hodge divisor $3P_7$ chooses the source vector e_ζ, the rational Borel chooses e_x, and there is only one nonconstant finite representation in the remaining estimate. Avoidance forces the extremal coefficient $-1/(Q_{\mathfrak{p}}+1)$; persistence of that extremum through coherent growing levels is what must imply a characteristic-zero stabilizer.

The same factorization also packages all finite primes into one canonical height. Put $\delta = \eta - \eta^{-1}$ and

$$t_g(P) = \frac{A_g(P)+q\delta}{A_g(P)-q\delta} = \frac{Y_g\overline{Z_g}}{\overline{Y_g}Z_g}. \tag{903}$$

Then $t_g\overline{t_g}=1$, and every complex absolute value of t_g is one. If

$$h_L^{\mathrm{fin}}(t) = \sum_{\mathfrak{P}} \max\{0, v_{\mathfrak{P}}(t)\}\log \mathrm{N}\,\mathfrak{P},$$

the coprimalities (900) give

$$\begin{aligned} h_L^{\mathrm{fin}}(t_g) &= \log|\,\mathrm{N}_{L/\mathbb{Q}}(Y_gZ_g)| + O_P(1)\\ &= \log|\,\mathrm{N}_{K/\mathbb{Q}}(D_gQ_7(pc-qa, pd-qb))| + O_P(1)\\ &= \log\left|\mathrm{N}_{K/\mathbb{Q}}\big(A_g(P)^2 + (4-\lambda^2)q^2\big)\right| + O_P(1). \end{aligned} \tag{904}$$

Away from $q\delta$, this is equality prime by prime: a split prime appears with one orientation in Y_g or Z_g, and the valuation of t_g records that orientation. At the fixed bad primes the possible cancellation is bounded by $v(q\delta)$. Since the archimedean contribution is zero, (904) is the full Weil height of a point on the anisotropic norm-one torus, up to normalization. Height zero makes t_g a root of unity by Kronecker; bounded height gives a finite set by Northcott.

The Cayley transform gives a second exact interpretation. For unnormalized projective heights,

$$h_L(t_g) = 2h_K([A_g(P):q]) + O_P(1) = \sum_{i=0}^{2}\log\left(A_{g,i}^2 + |\widetilde{\sigma}_i(\delta)|^2q_i^2\right) + O_P(1). \tag{905}$$

Indeed, $[A:q]\mapsto[A+q\delta : A-q\delta]$ is a fixed projective linear map over L, and extension from K to L doubles unnormalized height. Therefore the proximity product and the required height are genuinely different quantities:

$$\log|\,\mathrm{N}\,A_g| = \sum_i \log|A_{g,i}|, \qquad h_L(t_g) = \sum_i \log\max\big(|A_{g,i}|, |\widetilde{\sigma}_i(\delta)q_i|\big) + O_P(1).$$

Their difference is exactly Galois redistribution into the neutral Asai coordinate. This proves, without a numerical example, why improving the constant in a scalar C/q^2 approximation or invoking fixed-target Roth cannot control the required Arakelov height.

There is also an exact angular interpretation. If $w_{g,i}=\sigma_i(g)\zeta_i$ and $x_i=\sigma_i(x)$, then

$$\widetilde{\sigma}_i(t_g) = \frac{x_i-\overline{w_{g,i}}}{x_i-w_{g,i}}, \qquad \frac{|\widetilde{\sigma}_i(t_g)+1|}{|\widetilde{\sigma}_i(t_g)-1|} = \frac{|x_i-\mathrm{Re}\,w_{g,i}|}{\mathrm{Im}\,w_{g,i}}. \tag{906}$$

This follows at once from $Z_g = qY_g(x - g\zeta)$. Hence the elliptic approximation problem is approximation to -1 on the anisotropic norm-one torus. If the support of t_g is fixed, the equation involving $t_g, t_g + 1$ is an ordinary S-unit equation. When support is allowed to vary, the identity

$$(A_g + q\delta) - (A_g - q\delta) = 2q\delta$$

lies exactly on the coefficient-one *abc* boundary: its radical is already the torus height in (904). Thus ordinary *abc* is compatible with the escape and supplies no strict surplus. Any conjectural number-theoretic completion must couple this torus height to the $3P_7$ Hodge degree or impose a genuine thin-orbit inverse statement.

This is the global Arakelov replacement for prime-by-prime divisor control. The possible support growth is not to be estimated one prime at a time: it is exactly the height of t_g. The remaining theorem must compare this finite torus height with the positive $3P_7$ Hodge degree at a first record, with unit holonomy as the equality case. The norm-one torus is compact at infinity, so no hidden archimedean CM-unit direction remains; all noncompact Hodge shape lies in the two real K-unit coordinates already retained by the projectivized Asai components.

There is a still sharper, and completely presentation-free, state decomposition. Define the Cayley parametrization on K-points

$$c_\eta(y) = \frac{1 + y\eta}{1 + y\eta^{-1}} \in \big(\mathrm{Res}^1_{L/K}\,\mathbb{G}_m\big)(K),$$

and put

$$u_g = \frac{Z_g}{\overline{Z_g}}, \qquad v_g = \frac{Y_g}{\overline{Y_g}}.$$

Writing $(r_g, s_g) = (-q, p)g$, one has

$$g^{-1}x = -\frac{s_g}{r_g}, \qquad g^{-1}\infty = -\frac{d}{c},$$

and therefore

$$u_g = c_\eta(g^{-1}x), \qquad v_g = c_\eta(g^{-1}\infty), \qquad t_g = \frac{v_g}{u_g}. \tag{907}$$

The same coprimalities used in (900) give the exact height ledger

$$\begin{aligned} h_L(u_g) &= \log|\,\mathrm{N}_{K/\mathbb{Q}}\,Q_7(pc - qa, pd - qb)| + O_P(1) = 2h_K(g^{-1}x) + O_P(1),\\ h_L(v_g) &= \log|\,\mathrm{N}_{K/\mathbb{Q}}\,D_g| + O(1) = 2h_K(g^{-1}\infty) + O(1),\\ h_L(t_g) &= h_L(u_g) + h_L(v_g) + O_P(1). \end{aligned} \tag{908}$$

The last equality is genuinely additive, rather than merely the usual triangle inequality for height: away from the fixed primes dividing $q\delta$, the ideals supporting Y_g and Z_g cannot cancel in the quotient v_g/u_g.

This pair is literally one orbit of the full thin group. Put

$$C_\eta = \begin{pmatrix} \eta & 1 \\ \eta^{-1} & 1 \end{pmatrix}, \qquad \widehat{g} = C_\eta g^{-1} C_\eta^{-1}.$$

The matrix C_η represents c_η, so, with $u_* = c_\eta(x)$ and $v_* = c_\eta(\infty) = \eta^2$,

$$(u_g, v_g) = (\widehat{g}u_*, \widehat{g}v_*). \tag{909}$$

At each complex embedding C_η carries the real projective line to the unit circle; hence $\widehat{G}_7$ acts by unit-circle Möbius transformations. Moreover

$$\mathrm{Stab}_{G_7}(x) = \{g : \widehat{g}u_* = u_*\}.$$

Thus neither coordinate depends on a Rosen presentation: they are the two images of a fixed ordered boundary pair under an arbitrary element of G_7. The quotient $t_g = v_g/u_g$ is not a group character; it is a relative incidence coordinate of this diagonal pair orbit.

The two heights are also exact counts on the full pointed congruence tower. Let S_0 be the primes of L over (δ), and put

$$\iota_{\mathfrak{P},m}(g; y) = \mathbf{1}_{\{g\zeta \equiv y \pmod{\mathfrak{P}^m}\}} \qquad (y = x, \infty).$$

Primitivity of $[X_g : Y_g]$ and $[p : q]$ away from S_0 gives

$$v_{\mathfrak{P}}(Z_g) = \sum_{m\geq 1} \iota_{\mathfrak{P},m}(g; x), \qquad v_{\mathfrak{P}}(Y_g) = \sum_{m\geq 1} \iota_{\mathfrak{P},m}(g; \infty).$$

The CM coprimalities then give

$$\begin{aligned} h_L(u_g) &= \sum_{\mathfrak{P}\notin S_0} \sum_{m\geq 1} \iota_{\mathfrak{P},m}(g; x) \log \mathrm{N}\,\mathfrak{P} + O(1), \\ h_L(v_g) &= \sum_{\mathfrak{P}\notin S_0} \sum_{m\geq 1} \iota_{\mathfrak{P},m}(g; \infty) \log \mathrm{N}\,\mathfrak{P} + O(1). \end{aligned} \tag{910}$$

All nonzero good summands lie at split primes. Thus (911) below is not a metaphor about prime recycling: it compares the total depths of the two coherent sections $g\zeta = x$ and $g\zeta = \infty$ through every split prime-power quotient of the complete group orbit.

This produces an elementary but decisive closing criterion. If (g_n) is an infinite sequence of projectively distinct full-group returns and

$$\sup_n \big(h_L(t_{g_n}) - h_L(v_{g_n})\big) < \infty, \tag{911}$$

then (908) bounds $h_L(u_{g_n})$. Northcott gives $u_{g_m} = u_{g_n}$ for some $m \neq n$. The Cayley map is injective, so $g_m^{-1}x = g_n^{-1}x$, and $g_n g_m^{-1}$ is a nonidentity stabilizer of x. Likewise a collision among the v_g is a cusp stabilizer.

The proof of this criterion is formal; *it is not being presented as the missing thin-group estimate.* Its importance is diagnostic. It separates the complete state into two full-orbit coordinates: v_g is the cusp or parabolic payment and u_g is a complete projective state for $g^{-1}x$. The remaining estimate is exactly

$$h_L(t_g) - h_L(v_g) = h_L\big(c_\eta(g^{-1}x)\big) + O_P(1),$$

along the Hodge-stopped trajectory. Thus the relevant use of thinness cannot be a bound for six generators. It must be a relative domination or inverse theorem for the orbit of the *pair* (x, ∞): parabolic excursions are allowed arbitrary size but are entirely charged to v_g, while repeated thick returns must keep the residual state u_g in a bounded-height set unless the trajectory lies in one cyclic stabilizer. Strong approximation describes every finite orbit and expansion describes typical group elements; neither yet supplies this pointed, deterministic conclusion.

The factorization gives the precise S-unit bridge. The two conjugate factors $A_g(P) \pm q(\eta - \eta^{-1})$ have gcd supported on the fixed primes dividing $2q(\eta - \eta^{-1})$. Away from that set, their prime ideals

are distributed between the cusp denominator Y_g and the complete state Z_g. With bounded support their quotient is an ordinary S-unit and the established four-term S-unit theorem applies. Without bounded support, the missing assertion is a non-recycling theorem for this one CM factorization at first norm records. This is substantially more precise than tracking newly appearing primes one at a time.

The word "non-recycling" can now be stated without prime-by-prime imagery. Along the stopped orbit $g_n\zeta$, it asks for uniform control of repeated visits to the two projective sections ∞ and x in all split congruence quotients at once. Surjectivity of each reduction does not prove this, because it says nothing about the single path selected by x. Expansion gives equidistribution for most words in a large shell, but a proof for every $x \in \mathbb{P}^1(K)$ must also show that an exceptional pointed path which shadows these Borel cosets indefinitely is eventually a coset of one cyclic stabilizer. That final implication is the exact thin-orbit rigidity statement, rather than a local analysis of new primes.

It is tempting to use $\mathrm{N}\, Q_7$ itself as a greedy height on the Bass–Serre tree. That fails for an exact structural reason. The seven neighbors of an order-seven vertex retain, besides Q_7^7, the independent angular invariant $(u - v\eta)^{14} + (u - v\eta^{-1})^{14}$. Thus their product is not a scalar power of Q_7. There are strict local minima away from norm one; for example

$$|\,\mathrm{N}\, Q_7(3,4)| = 2653,$$

whereas the seven neighboring norms are

$$14197, 117349, 225191, 23521, 5867, 5041, 22807.$$

This exact example rules out a one-vertex greedy descent. A valid Bass–Serre argument would have to use a nonlocal convexity or return theorem on the whole tree; checking the seven moves cannot prove it.

This is a new simplification of the maximal-slope problem. The positive Galois component is not an arbitrary compensating direction: after squaring it is the line of the one fixed effective divisor $3P_7$, and Υ_1^{-1} is its canonical holomorphic section. The neutral component is actually trivial, but this does not make its rational Asai coordinate bounded. It shows exactly where that possible escape lives:

$$u_{\mathrm{neut}}(r) = \log \left| \frac{\sigma_0(r)\sigma_1(r)}{\sigma_2(r)^2} \right|$$

has no unbounded Hodge contribution and is a purely arithmetic unit/ideal coordinate. The complementary direction is charged by the full G_7-orbit of the order-seven point. Thus the rank-two problem splits into one curved Hodge direction and one flat S-unit direction. A proof still requires a stopped arithmetic Jensen or height-jump inequality for the former and support or holonomy control for the latter; no estimate for individual Rosen generators supplies either joining.

This was the correct analytic splitting but not yet the optimal arithmetic formulation. The later identity $\nu(m) = \alpha_2(m)^{-2}\tau(\alpha_2(m))^{-1}$, with $\alpha_2(m) = \tau^2(m)/m$, shows that after Galois completion the flat character is a monomial in the one fundamental K-valued coordinate. Thus a separate support theorem for the flat direction can be replaced by an all-place stopped height bound for α_2; obtaining that bound remains the same equality-case obstacle.

The analytic splitting is equally explicit. For a rational flag x, put

$$G_{i,x}(z) = \Phi(F_i(z), x_i^\dagger) - \Phi(z, x_0^\dagger), \tag{912}$$

$$\Psi_{i,x}(z) = \frac{F_i'(z)(z - x_0^\dagger)}{F_i(z) - x_i^\dagger}. \tag{913}$$

Then

$$dd^c G_{i,x} = 2(\omega_{\mathrm{hyp}} - F_i^* \omega_{\mathrm{hyp}}) > 0, \tag{914}$$

$$\log |\Psi_{i,x}|^2 = G_{i,x} + \log \|\Xi_i\|^2. \tag{915}$$

Poincaré–Lelong gives the two charges $2P_7$ and P_7 separately. These formulas use the equivariance for every element of G_7; no Rosen digit or chosen word occurs in them.

II.20.4 The full Asai orbit retains the same two coordinates

The rank-two correction also appears at finite places. Recall the three Galois-permuted components $\mathcal{Z}_0, \mathcal{Z}_1, \mathcal{Z}_2$ of the exterior square of the Asai tensor. If r is the incidence remainder and τ_i is the automorphism satisfying $\sigma_i = \sigma_0 \circ \tau_i$, their coordinate ideals are

$$I(\mathcal{Z}_0) = (r), \qquad I(\mathcal{Z}_i) = (\tau_i r) \quad (i = 1, 2). \tag{916}$$

Therefore the two projective ideal ratios are

$$\frac{I(\mathcal{Z}_i)}{I(\mathcal{Z}_0)} = \left(\frac{\tau_i r}{r}\right), \qquad i = 1, 2. \tag{917}$$

At the distinguished real place their logarithms are exactly $\log|\sigma_i r| - \log|\sigma_0 r|$, up to the already explicit Hodge angle. Thus the direct sum of all three Asai components, modulo its common radial line, realizes the two-dimensional norm-one torus on which (884) lives.

By contrast, the earlier scalar tensor $\mathcal{Z}_1 \otimes \mathcal{Z}_2 / \mathrm{N}(r)$ has finite ideal $(r)^{-1}$ but replaces the pair (917) by its determinant. It records the radial norm and loses one unit direction. The correct integral object should therefore retain

$$\mathcal{Z}_0 \oplus \mathcal{Z}_1 \oplus \mathcal{Z}_2 \quad \text{modulo common scale},$$

or an equivalent rank-two torus bundle, and use the scalar incidence line only for the radial norm. This object is canonical under the full Galois and thin-group actions; an asymmetric direct sum invented state by state would not be.

There is a useful elementary finiteness consequence. If algebraic integers r_n satisfy

$$|\mathrm{N}(r_n)| \le B, \qquad |\log|\sigma_i r_n| - \log|\sigma_0 r_n|| \le C \quad (i = 1, 2), \tag{918}$$

then the three values $|\sigma_i r_n|$ are bounded. This follows by solving the two difference equations together with $\sum_i \log|\sigma_i r_n| = \log|\mathrm{N}\, r_n|$. The Minkowski embedding of $\mathcal{O}_K$ is discrete, so only finitely many r_n occur. Applied to both entries of a Rosen state, this gives actual state repetition and hence the special fixed-line alternative.

Accordingly, a canonical sufficient theorem should no longer be phrased as one determinant-degree bound. It bounds both maximal slopes of the projectivized three-component Asai lattice at norm-record returns. The radial bound controls $|\mathrm{N}\, r|$; the two split Hodge defects control the torus coordinates. Equality in a rational torus direction is unit holonomy and gives the special alternative. This is a precise use of the thin group as one integral Galois representation. What is not yet proved is either the stopped maximal-slope inequality connecting (914) to (917), or a scalar inequality together with rigidity of its determinant-kernel trajectories.

II.20.5 A record-free route through the integral Asai exponent

There is a second consequence of retaining the full Asai lattice. It does not use a norm record. Let $P \in \mathcal{O}_K^2$ represent $x \in \mathbf{P}^1(K)$, and let

$$Y(P) = \sigma_0(P) \otimes \sigma_1(P) \otimes \sigma_2(P)$$

be the corresponding nonzero vector in the integral rank-eight tensor induction. Along the geodesic ray to $\sigma_0(x)$, the distinguished factor is the contracting endpoint line. In the normalization with uniformizing exponent one,

$$\log \|\sigma_0(P)\|_{0,t} = -t + O_P(1). \tag{919}$$

Proposition II.20.2 (Integral Asai obstruction)**.** *If x is not a cusp and t_n are compact return times of its geodesic ray, then*

$$\limsup_n \frac{1}{t_n} \left(\log \|\sigma_1(P)\|_{1,t_n} + \log \|\sigma_2(P)\|_{2,t_n}\right) \geq 1. \tag{920}$$

Proof. Reduce the n-th return to one fixed compact set using $g_n \in G_7$. Hodge equivariance carries $Y(P)$ to the nonzero integral vector $Y(g_nP)$. On a compact set, the shortest nonzero vector in the integral local-system lattice has a uniform positive norm. Therefore

$$\sum_{i=0}^{2} \log \|\sigma_i(P)\|_{i,t_n} \geq -O(1).$$

Equation (919) gives the result. □

The equality case of this obstruction is exact. If $hP = \varepsilon P$ for a nontrivial stabilizer $h \in G_7$, then ε and ε^{-1} are integral, so $\varepsilon \in \mathcal{O}_K^\times$, and

$$\mathrm{As}(h)Y(P) = \mathrm{N}_{K/\mathbf{Q}}(\varepsilon)Y(P) = \pm Y(P). \tag{921}$$

Conversely, if two distinct return elements g_m, g_n give the same projective decomposable tensor $[Y(g_mP)] = [Y(g_nP)]$, equality of the distinguished factor lines implies that $g_n^{-1}g_m$ fixes $[P]$. Thus projective repetition in the integral Asai lattice is exactly the cusp/special closing event.

For $P = (\lambda^2 - 1, 1)^t$, the elementary special matrix has eigenvalue

$$u_7 = \lambda(\lambda^2 - 1) + 1 = \lambda^2 + \lambda, \qquad \mathrm{N}(u_7) = 1.$$

Both outer conjugates of u_7 lie in $(0, 1)$, so their inverse expansions add exactly to the distinguished contraction. This gives a literal zero Asai exponent on the basic special orbit and shows why the periodic alternative must be retained in every strict-descent proposal.

The generic outer top exponents are $1/5$ and $3/5$, with the corrected embedding labels. Hence at an Oseledets-regular ray every fixed outer vector has upper total growth at most $4t/5 + o(t)$, and

$$\log \|Y(P)\|_t \leq -\frac{1}{5}t + o(t).$$

At compact returns this contradicts the integral lattice lower bound. Thus every noncuspidal trace-field endpoint is forced into an exceptional outer-Lyapunov level set.

This gives a concise conditional route.

Conjecture II.20.3 (Trace-field Oseledets dichotomy)**.** *Every $x \in \mathbf{P}^1(K)$ is either a parabolic endpoint, a hyperbolic periodic endpoint, or is Oseledets regular for the two outer Hodge variations with upper exponents* $1/5$ *and* $3/5$.

The proposition rules out the third alternative, so this conjecture implies the complete cusp/special classification. In the periodic case, $x \in K$ makes the eigenvalue $cx + d$ a unit of K, hence the period is special.

Route II.20.4 (Outcome of the integral-exponent route)**.** This route succeeds unconditionally in proving the exceptional-growth bound (920). It replaces the moving record line by one fixed integral Asai vector and therefore avoids the index-shift and unit-properness defects of the scalar stopping argument. Its remaining assertion is an arithmetic rigidity theorem for exceptional trace-field rays: exceptional outer growth must force periodicity.

Ordinary Roth does not prove that assertion. Roth bounds the depth of an individual exceptionally good cusp excursion relative to the height already accumulated. Oseledets regularity is a statement about the sum of the Hodge cocycle over all excursions. For example, an abstract digit sequence with logarithmic cusp depths of order $\sqrt{n}$ has each new depth negligible relative to the cumulative time, as Roth permits, while the proportion of time in the thick part tends to zero. A fixed-ε Roth theorem therefore supplies no time-average bound. What would suffice is a trace-field equidistribution theorem, or the weaker assertion that every nonperiodic trace-field ray has outer Hodge upper-exponent sum strictly below one.

The last assertion belongs to a standard hierarchy. Let μ_{BM} be the Bowen–Margulis measure on the natural extension of the accelerated Rosen map. Consider the algebraic-genericity conjecture that every $x \in \mathbf{P}^1(K)$ whose Rosen orbit neither terminates nor becomes periodic is generic for μ_{BM}. This is the Rosen analogue of the conjectural Gauss normality of nonquadratic algebraic real numbers. Birkhoff–Oseledets applied to the two outer Hodge cocycles, normalized by the accumulated hyperbolic roof time, then gives exponents $1/5$ and $3/5$, so the fixed integral Asai tensor has exponent at most $-1/5$. Proposition II.20.2 rules this out at compact returns. Thus algebraic Rosen genericity implies the full classification.

This is stronger than necessary. The exact conditional input is only

$$\limsup_{t\to\infty} \frac{1}{t} \left(\log \|\sigma_1(P)\|_{1,t} + \log \|\sigma_2(P)\|_{2,t}\right) < 1 \tag{922}$$

for every nonperiodic, noncuspidal $[P] \in \mathbf{P}^1(K)$. Full digit normality, separate frequency laws, and a chosen generating set are not part of (922). This makes the division of labor precise: algebraic number theory must exclude rational points from one exceptional Lyapunov level set, while Hodge theory computes the generic level and the integral Asai lattice turns the strict gap into closing.

II.20.6 Audit of the integral-vector theorem for log-Anosov systems

Filip's integral-vector theorem for log-Anosov representations has almost the form needed above. If a nonzero integral vector belongs to the strictly decreasing boundary subbundle, it decays at compact returns; lattice discreteness then forces the boundary point to be parabolic. The proof is exactly the compact-return argument in Proposition II.20.2.

There is an important root-theoretic refinement. The product representation

$$\rho_{\mathrm{Gal}} : G_7 \longrightarrow \mathrm{SL}_2(\mathbf{R})^3$$

is log-Anosov for the distinguished simple root: its first projection is the Fuchsian uniformizing representation. This still does not put the Asai endpoint vector in Filip's decreasing bundle. The weights of $V_0 \otimes V_1 \otimes V_2$ above the central contracting line are

$$(-1,+1,+1), \quad (-1,+1,-1), \quad (-1,-1,+1), \quad (-1,-1,-1). \tag{923}$$

The partial order for the distinguished root compares weights only when their two outer coordinates agree. None of the weights in (923) is therefore below the zero weight. Replacing this partial order by the numerical sum of the three coordinates is not permitted. This explains why the theorem does not incorrectly declare a known special endpoint to be cuspidal.

The hypotheses fail here in a measurable and important way. For T^n, let $a_i(n)$ be the logarithm of the largest singular value in the i-th rank-two factor. Directly from the unipotent matrix,

$$a_i(n) = \log n + \log|\lambda_i| + o(1).$$

The numerical weight of the most dangerous ordinary integral Asai endpoint component is therefore

$$-a_0(n) + a_1(n) + a_2(n) = \log n + O_7(1) > 0. \tag{924}$$

Its positivity is the cusp manifestation of the incomparability above. Adding the second central copy gives

$$-2a_0(n) + a_1(n) + a_2(n) = -3\log\lambda_0 + o(1) < 0, \tag{925}$$

where $|\lambda_0\lambda_1\lambda_2| = 1$ was used. This is the desired pairwise contraction, but it is only a constant as the hyperbolic cusp distance grows like $2\log n$. Hence it is not the uniform linear gap required by log-Anosov domination. More importantly, the corresponding doubled-central Hodge line is not itself a flat integral subrepresentation.

Route II.20.5 (Outcome of the log-Anosov audit)**.** The published integral-vector theorem cannot be applied directly. The ordinary Asai representation passes the integrality test, but its endpoint weights are not in the decreasing root-poset bundle; the dangerous scalar combination is positive in (924). The pairwise Hodge line passes the negative-weight test in (925) but fails the flat integral-subrepresentation test. Constructing the projectivized three-component Asai filtration, or proving determinant-kernel rigidity, is precisely what would join the two hypotheses. This is not a missing citation to an existing log-Anosov theorem.

There is a refinement of this audit which uses the full strength of Filip's poset formalism. His Example 4.3.17 is designed for a uniformizing rank-two system together with a Galois-conjugate Hodge system. Strict Schwarz–Pick contraction refines the bare root order and places the three real bundles in the order

$$E_0^+ \succ E_i \succ E_0^-.$$

Applied to the pairwise tensor $\mathbb{W}_i = V \otimes {}^{\tau^i}V$, this proves exactly the real inequality obtained later for the full-group pairwise state: the distinguished component of $w_i(P)$ has weight at most $-a_0 + a_i$, strictly negative whenever positive compact-part length is accumulated. Thus the issue is not that Filip's formalism sees only the six generators, nor that it misses the Schwarz–Pick refinement.

The obstruction occurs one step later, at rational descent. As a $\mathbb{Q}$-local system, $\mathrm{Res}_{K/\mathbb{Q}}\,\mathbb{W}_i$ has three cyclically conjugate real components. A nonzero rational vector has all three components, whereas the decreasing subspace singled out above contains only the distinguished one. Their degrees are the rows of (806); each row has total degree zero and contains a positive conjugate component. Equivalently, taking the Galois norm of the distinguished decreasing line changes its

negative degree into degree zero. Consequently the hypothesis $v \in V_{\mathbb{Z}} \cap V^{-}$ in Filip's Theorem 4.5.3 is not satisfied.

This also identifies the archimedean half of the closed-Bruhat channel in the ambient restriction formula. At a finite prime its vector is the Borel-fixed line $\mathbb{C}\phi_0$. At infinity its distinguished component is the Schwarz–Pick decreasing line. Galois descent restores the compensating components, and the moving reciprocal incidence ideal in (800) is exactly the datum which decides whether that compensation is genuine unit holonomy or unrelated escape. The required extension of the integral-vector theorem can now be stated without symbolic coding:

> At first norm-record returns, a rational vector in the nonpositive pairwise Hodge filtration either enters a strictly decreasing rational subspace, hence is parabolic, or its Galois compensation is one recurrent zero-weight projective line, hence has a nontrivial stabilizer.

Proposition II.19.3 proves the last implication once the stopped maximal-slope bound is supplied. This is an adelic semistable extension of Filip's theorem; applying the strict theorem alone would discard the special hyperbolic alternative.

II.20.7 Full-group parabolic coding versus stammering theorems

The presentation $G_7 \simeq C_2 * C_7$ gives one more number-theoretic test which is invisible in the scalar rational-approximation exponent. A Rosen digit a_j is a run of $|a_j|$ copies of a parabolic letter in a finite-alphabet coding of the full group. Its geometric contribution to cusp time, however, is only

$$d_j = \log |a_j| + O(1).$$

Consequently a theorem about long repeated finite words and a theorem about long cusp time are not the same theorem.

The distinction is witnessed without a numerical experiment. Consider the abstract digit depths

$$|a_j| = \left\lceil e^{\sqrt{j}} \right\rceil .$$

Then

$$\sum_{j \le N} d_j = \frac{2}{3} N^{3/2} + O(N), \qquad \frac{d_N}{\sum_{j<N} d_j} = O(N^{-1}), \tag{926}$$

$$\sum_{j<N} |a_j| \sim 2\sqrt{N} e^{\sqrt{N}}, \qquad \frac{|a_N|}{\sum_{j<N} |a_j|} \sim \frac{1}{2\sqrt{N}}. \tag{927}$$

The first line says that every new excursion is sublinear relative to the height already accumulated, exactly as Roth requires, while the total thick-part time $O(N)$ has density zero. The second line says that the new parabolic run occupies a vanishing, rather than a positive, proportion of the preceding finite word.

Adamczewski–Bugeaud's shifted stammering criterion requires a repeated block whose size dominates its prefix by a fixed amount; the ratio in (927) goes in the opposite direction [4]. In the ordinary positive continued-fraction model one also has

$$\log q_N \asymp \sum_{j \le N} \log |a_j| \asymp N^{3/2}.$$

Hence $\log\log q_N = O(\log N)$, far below the Davenport–Roth transcendence threshold of order $N/\sqrt{\log N}$. The strongest quantitative denominator-growth consequence of Roth does not exclude the model either.

Route II.20.6 (Outcome of the stammering audit)**.** Passing from Rosen digits to the finite alphabet of the actual thin group is conceptually correct, but existing repetition theorems detect a single parabolic block comparable with the full preceding word. The Asai obstruction detects the strictly broader regime in which the *sum of the logarithmic* parabolic depths has density one. A sufficient new number-theory statement is therefore:

A trace-field boundary word whose logarithmic parabolic time has density one is parabolic or eventually periodic.

This is the finite-alphabet version of trace-field Oseledets rigidity. It uses the whole relatively hyperbolic group, rather than one rational approximation, and the calculation above proves that it is not a disguised application of Roth, Davenport–Roth, or the known stammering criteria.

Route II.20.7 (Successful partial theorem: thin-group stammering exactification)**.** The preceding audit concerns the published stammering theorem, whose conclusion does not exclude an endpoint already in the trace field. The fixed-form exactification mechanism developed later gives a sharper statement tailored to precisely that endpoint.

Let p_n/q_n be the Rosen convergents of $\xi \in K \cap [-\lambda/2, \lambda/2)$, and suppose $U_\nu V_\nu^{w_\nu}$ is a prefix of length N_ν, with $u_\nu = |U_\nu|$, $v_\nu = |V_\nu| > 0$, and admissible periodic continuation $U_\nu V_\nu^\infty$. The argument written out in Theorem II.20.8 proves that

$$\frac{(q_{u_\nu} q_{u_\nu+v_\nu})^3}{q_{N_\nu}^2} \longrightarrow 0 \quad \Longrightarrow \quad \xi \text{ is a cusp or a special fixed point.} \tag{928}$$

Theorem II.20.8 (Exactification of a stammering trace-field endpoint)**.** *Implication* (928) *holds. If* $1 < b = \liminf q_n^{1/n} \le \limsup q_n^{1/n} = B < \infty$*, it is enough that*

$$\limsup_\nu \frac{N_\nu}{2u_\nu + v_\nu} > \frac{3}{2}\frac{\log B}{\log b}. \tag{929}$$

Proof. Let $\alpha_\nu = [U_\nu, V_\nu^\infty]$, and use the convergent matrices

$$C_n = \begin{pmatrix} p_{n-1} & p_n \\ q_{n-1} & q_n \end{pmatrix}.$$

The matrix $M_\nu = C_{u_\nu+v_\nu} C_{u_\nu}^{-1}$ fixes α_ν and belongs to the full Hecke reflection group. The central entries of C_n are $O(q_n)$, so

$$H_0(Q_{M_\nu}) \ll q_{u_\nu} q_{u_\nu+v_\nu}. \tag{930}$$

The common-prefix cylinder estimate is

$$|\xi - \alpha_\nu| \ll q_{N_\nu}^{-2}.$$

If β_ν is the other fixed point, then $Q_{M_\nu}(z,1) = c(z-\alpha_\nu)(z-\beta_\nu)$. The root-sum identity gives

$$|c(\xi - \beta_\nu)| \ll_\xi H_0(Q_{M_\nu}),$$

with the linear case $c = 0$ even simpler. For a fixed primitive column P representing ξ, therefore,

$$|\sigma_0 Q_{M_\nu}(P)| \ll_P H_0(Q_{M_\nu}) q_{N_\nu}^{-2}.$$

Parabolic-ray trace dominance applies in the full reflection group and bounds both outer coefficient vectors by H_0. If the integral value $Q_{M_\nu}(P)$ were nonzero, the product formula would give

$$1 \ll_P |\sigma_0 Q_{M_\nu}(P)| H_0(Q_{M_\nu})^2 \ll_P \frac{(q_{u_\nu} q_{u_\nu+v_\nu})^3}{q_{N_\nu}^2},$$

contradicting (928). Hence a return matrix fixes ξ. Squaring it if it reverses orientation gives the required parabolic or hyperbolic element of G_7.

Finally $q_n \geq (b-\eta)^n$ and $q_n \leq (B+\eta)^n$ eventually. Substitution shows that the logarithm of the last quotient is at most

$$3(2u_\nu + v_\nu)\log(B+\eta) - 2N_\nu \log(b-\eta),$$

which tends to $-\infty$ under (929) after choosing η small. □

Why this is a real improvement. The published Rosen stammering argument first takes the integer height of a quadratic-over-K surrogate and then pays the exponent 3 in algebraic approximation. With $[K : \mathbf{Q}] = 3$, its threshold is $9/2$, and its quadratic-or-transcendental alternative cannot exclude $\xi \in K$. Keeping the fixed-point form over K and using the thin-group trace theorem removes that extra factor 3, leaving the $3/2$ in (929). This proves the desired classification in the dominant-repeated-block sector.

Why it is not yet the global theorem. For the test depths $|a_j| = \lceil e^{\sqrt{j}} \rceil$, every new run is a vanishing fraction of the preceding word. Thus even the new $3/2$ criterion fails while total logarithmic cusp time still has density one. That is the balanced many-block obstruction which the stopped Hodge–Arakelov inequality must address.

At $q = 11$, the same proof uses four outer embeddings and replaces the third power in (928) by a fifth power, giving threshold $5/2$. This works without a favorable Hodge sign and therefore clarifies the division of labor: trace dominance is a group-theoretic tool for dominant literal repetition, whereas the Hodge defect is the proposed control on diffuse, nonrepetitive cusp time.

II.20.8 Advance: the shrinking real target is unnecessary

The first version of the congruence route below asked one group element to satisfy three simultaneous conditions: a congruence Borel return, an exponentially accurate real return, and a typical outer Hodge bound. A better choice of defect removes the real return completely.

Fix a primitive column $P \in \mathcal{O}_K^2$, set $x = [P]$, and let $e = (1,0)^t$. For $h \in G_7$, define the cusp-incidence defect

$$r_h = \det(P, he).$$

If the cusp $h\infty$ reduces to x modulo an ideal $\mathfrak{a}$, then

$$r_h \in \mathfrak{a}. \tag{931}$$

Let

$$R(h) = d_{\mathbb{H}}(i, \sigma_0(h)i), \qquad a_i(h) = \log \|\sigma_i(h)\|_{\mathrm{op}}.$$

The three archimedean estimates for this *one scalar* are

$$\log|\sigma_0 r_h| \leq \frac{1}{2}R(h) + O_P(1), \qquad \log|\sigma_i r_h| \leq a_i(h) + O_P(1) \quad (i = 1, 2).$$

They immediately give the following criterion.

Proposition II.20.9 (Optimal cusp-lift closing). *Suppose that* $\mathrm{N}\,\mathfrak{a}_n \to \infty$ *and that* $h_n \in G_7$ *satisfy*

$$\begin{aligned} h_n\infty &\equiv x \pmod{\mathfrak{a}_n}, \\ R(h_n) &\le A \log \mathrm{N}\,\mathfrak{a}_n + O(1), \\ a_1(h_n) + a_2(h_n) &\le \theta R(h_n) + O(1). \end{aligned} \tag{932}$$

If $A(1/2+\theta) < 1$, *then* x *is a cusp.*

Proof. The displayed archimedean bounds give

$$\log |\,\mathrm{N}\, r_{h_n}| \le A(1/2+\theta) \log \mathrm{N}\,\mathfrak{a}_n + O_P(1).$$

If $r_{h_n} \neq 0$, divisibility (931) gives $\log |\,\mathrm{N}\, r_{h_n}| \ge \log \mathrm{N}\,\mathfrak{a}_n$, a contradiction for large n. Thus $r_{h_n} = 0$, so P and $h_n e$ represent the same projective point and $x = h_n\infty$. □

This is the cleanest place found so far where the thin group does real mathematical work. Strong approximation for the Zariski-dense diagonal representation gives, away from finitely many ideals, a projective congruence orbit of size

$$\#X_{\mathfrak{a}} = (\mathrm{N}\,\mathfrak{a})^{1+o(1)}.$$

Central orbital growth is $e^{R+o(R)}$. Reaching one prescribed residue flag should therefore occur at the information-theoretically optimal scale

$$R = (1+o(1)) \log \mathrm{N}\,\mathfrak{a}.$$

The mean sum of the two outer Hodge exponents is $2/5$. Substitution in the closing criterion gives

$$1\left(\frac{1}{2} + \frac{2}{5}\right) = \frac{9}{10} < 1.$$

Thus an entropy-optimal, Hodge-typical lift leaves a genuine margin of $1/10$ and proves that x is a cusp without any real shrinking condition.

Conjecture II.20.10 (Hodge-balanced optimal lifting dichotomy). *For every* $x \in \mathbf{P}^1(K)$, *either* $\mathrm{Stab}_{G_7}(x) \neq \{\pm 1\}$, *or, for every* $\epsilon > 0$, *there are ideals of unbounded norm and elements* h_n *such that*

$$h_n\infty \equiv x \pmod{\mathfrak{a}_n}, \qquad R(h_n) \le (1+\epsilon) \log \mathrm{N}\,\mathfrak{a}_n, \qquad a_1(h_n) + a_2(h_n) \le (2/5+\epsilon) R(h_n).$$

Proposition II.20.9, with ϵ small enough that $(1+\epsilon)(9/10+\epsilon) < 1$, proves the full classification from this conjecture. The first alternative gives a cusp or a special hyperbolic fixed point by the rational-Borel calculation; the second gives a cusp.

This terminology matches Sarnak's optimal strong approximation and the optimal-lifting framework of Golubev–Kamber [37]. Those principles place the lifting radius at the entropy threshold, but their basic form controls almost all pairs of classes in congruence quotients of arithmetic lattices. The current target is stronger in two specific ways: it is the reduction of one fixed rational flag in a thin orbit, and the lift must satisfy a nonunitary weight bound in the two Galois representations. The periodic alternative is exactly the possible exception hidden by an almost-all statement.

There is a sharp warning from the full arithmetic group. Kamber–Varjú study the projective action of $\mathrm{SL}_2(\mathbb{Z})$ itself [38]. With distance equal to logarithmic operator norm, the information-theoretic projective scale for $n = 2$ is $(1/2 + o(1)) \log q$, whereas the limsup of the worst projective distance is $(1 + o(1)) \log q$. The displayed worst targets in their proof use special composite moduli. Doubling to the hyperbolic displacement normalization gives an almost-all scale $\log q$ and a possible all-target scale $2 \log q$. The heptagonal product-formula argument needs less than $(10/9 - o(1)) \log q$.

Consequently the prescribed-target clause is doing real work. Neither Sarnak's density conjecture nor generalized Ramanujan, when used only through an almost-all optimal-lifting theorem, implies the heptagonal statement. Passing to split primes avoids the known composite lower-bound example, but a fixed algebraic section may still lie in the exceptional set at every level. What must be classified is coherence of those exceptions as the level varies: it must force the rational Borel of $[P]$, not merely have small cardinality.

Only a horizontal form is needed. Fix a small $\epsilon > 0$. For arbitrarily large Q, it suffices to find one split prime ideal $\mathfrak{p}$, with $Q \le \mathrm{N}\,\mathfrak{p} \le 2Q$, and one $h \in G_7$ for which

$$h\infty \equiv x \pmod{\mathfrak{p}}, \qquad R(h) \le (1+\epsilon)\log Q, \qquad a_1(h) + a_2(h) \le (2/5+\epsilon)R(h).$$

Indeed Proposition II.20.9 applies with $\mathfrak{a} = \mathfrak{p}$. Equivalently, if x is not a cusp, then every Hodge-balanced element in the entropy ball omits the reduction of x for every prime of size Q. This is violently at odds with the expected equidistribution of the thin cusp orbit.

The corresponding pair count makes the remaining number theory explicit:

$$\sum_{\substack{Q \le \mathrm{N}\,\mathfrak{p} \le 2Q \\ \mathfrak{p} \text{ split}}} \#\left\{ h : \begin{array}{l} R(h) \le (1+\epsilon)\log Q, \quad a_1(h) + a_2(h) \le (2/5+\epsilon)R(h), \\ h\infty \equiv x \pmod{\mathfrak{p}} \end{array} \right\}.$$

Local density predicts $1/\mathrm{N}\,\mathfrak{p}$ for the last condition. Since the ball has $Q^{1+\epsilon+o(1)}$ elements and there are $Q^{1+o(1)}$ split primes in the dyadic interval, the predicted total is $Q^{1+\epsilon+o(1)}/\log Q$, unless the rational section is periodic. Proving a positive lower bound at moduli of the same exponential size as the orbit ball is a level-one distribution problem. This is much stronger than a fixed spectral gap, but weaker and more natural for the application than demanding optimal lifting for every target in every quotient.

Route II.20.11 (What succeeds and what remains in the optimal-lifting route)**.** The successful part is unconditional and global: one prescribed cusp lift modulo an externally chosen ideal, at central exponent 1 and outer exponent 2/5, is enough. This replaces the former three-way shrinking-target count by a two-way optimal-lifting problem.

Strong approximation and expansion alone do not supply the stated conjecture. They give surjectivity and a bound $R \le C \log \mathrm{N}\,\mathfrak{a}$ with an unspecified covering constant. Here one needs $C < 10/9$, essentially the entropy-optimal constant 1, and one needs the lift in a prescribed congruence fibre to obey the outer Hodge large-deviation bound. The residual theorem is therefore precise: entropy-optimal projective lifting with Hodge balance, with periodic concentration as the only exception. This is a property of the whole family of thin congruence Schreier graphs and the full Galois representation; it is not a calculation with the generators.

II.20.9 A two-prime amplification of the fixed-target problem

The pointwise modulus in optimal lifting has size nearly e^R. The same closing can be arranged with a product of two primes, each of size strictly below $e^{R/2}$. This does not prove the needed

distribution theorem, but it changes it into a factorial-moment problem to which bilinear dispersion and a large sieve are naturally adapted.

For $\delta > 0$, let $\Omega_R(\delta)$ be the Hodge-balanced part of the central ball:

$$\Omega_R(\delta) = \{g \in G_7 : R(g) \le R, \quad a_1(g) + a_2(g) \le (2/5 + \delta)R\}.$$

Orbital equidistribution and the Hodge Lyapunov exponents predict, and the usual large-deviation formulation gives, $\#\Omega_R(\delta) = e^{R-o(R)}$. Fix numbers

$$9/10 + \delta < s < 1, \qquad u = s/2 < 1/2, \tag{933}$$

and let $\mathcal{P}_R(u)$ be the good split prime ideals with $e^{uR} \le \mathrm{N}\,\mathfrak{p} \le 2e^{uR}$.

For a primitive $P \in \mathcal{O}_K^2$, define the second prime moment

$$\mathscr{M}_R(P) = \sum_{\substack{\mathfrak{p},\mathfrak{q} \in \mathcal{P}_R(u) \\ \mathfrak{p} \ne \mathfrak{q}}} \#\{g \in \Omega_R(\delta) : g[P] = [P] \pmod{\mathfrak{p}\mathfrak{q}}\}. \tag{934}$$

Strong approximation and the Chinese remainder theorem give the local density

$$\frac{1}{(\mathrm{N}\,\mathfrak{p} + 1)(\mathrm{N}\,\mathfrak{q} + 1)} \tag{935}$$

up to the harmless replacement of the norms by the exact projective orbit sizes.

Conjecture II.20.12 (Hodge-weighted two-prime dispersion)**.** *If $[P]$ has projectively trivial stabilizer, then, for some $\delta > 0$ and parameters as in (933),*

$$\mathscr{M}_R(P) = (1 + o(1))\#\Omega_R(\delta) \sum_{\substack{\mathfrak{p},\mathfrak{q} \in \mathcal{P}_R(u) \\ \mathfrak{p} \ne \mathfrak{q}}} \frac{1}{(\mathrm{N}\,\mathfrak{p} + 1)(\mathrm{N}\,\mathfrak{q} + 1)}. \tag{936}$$

The same conclusion with a positive constant lower bound in place of the asymptotic is sufficient.

The right side of (936) has order $e^{R-o(R)}/R^2$. This is not a pointwise assertion for a modulus of size e^{sR}: it is an average over pairs of prime levels, each below the half-entropy scale. The expected local count for each pair is still exponential, of size $e^{(1-s-o(1))R}$.

There is a standard-shaped sufficient estimate for the displayed asymptotic. Put

$$E_{\mathfrak{p},\mathfrak{q}}(P; R) = \#\{g \in \Omega_R(\delta) : g[P] = [P] \pmod{\mathfrak{p}\mathfrak{q}}\} - \frac{\#\Omega_R(\delta)}{(\mathrm{N}\,\mathfrak{p} + 1)(\mathrm{N}\,\mathfrak{q} + 1)}.$$

Conjecture II.20.13 (Hodge-weighted Barban–Davenport–Halberstam bound)**.** *For the same fixed rational flag and prime range,*

$$\sum_{\substack{\mathfrak{p},\mathfrak{q} \in \mathcal{P}_R(u) \\ \mathfrak{p} \ne \mathfrak{q}}} |E_{\mathfrak{p},\mathfrak{q}}(P; R)|^2 \le e^{R+o(R)}. \tag{937}$$

This is the fixed-residue, Hodge-weighted analogue of a Barban–Davenport–Halberstam variance estimate. It is not a consequence of the classical theorem: the sequence being distributed is a thin orbital shell, the residue target varies coherently with P, and the level is exponential in the

hyperbolic radius. Indeed the exact variance gap is visible before any estimate is attempted. If $[P]$ has no nonidentity stabilizer, the norm bound (939) below makes every pair count equal to the identity contribution. The predicted count for one pair is $e^{(1-s-o(1))R}$, so the left side of (937) then has size

$$e^{(2-s-o(1))R}, \tag{938}$$

rather than $e^{R+o(R)}$. Thus the needed inverse theorem must save the exact factor $e^{(1-s)R}$, or classify its failure as concentration in B_P. This is a numerical measure of the fixed coherent hole, not another vague appeal to mixing.

Proposition II.20.14 (The variance bound implies two-prime dispersion)**.** *Conjecture II.20.13 implies* (936) *with a relative* $o(1)$ *error.*

Proof. The number of ordered prime pairs is $e^{sR+o(R)}/R^2$. By Cauchy–Schwarz and (937), the sum of all errors is at most

$$\exp\left(\frac{1+s}{2}R + o(R)\right).$$

The sum of the local main terms is $e^{R-o(R)}/R^2$. Since $s<1$, the former is exponentially smaller than the latter. □

Proposition II.20.15 (Two-prime dispersion completes the classification)**.** *Conjecture II.20.12 implies that every point of* $\mathbf{P}^1(K)$ *has a projectively nontrivial* G_7*-stabilizer, and hence is a cusp or a special hyperbolic fixed point.*

Proof. For fixed P, put $D_g(P) = \det(P, gP)$. At the three real places,

$$|\sigma_i D_g(P)| \ll_P \|\sigma_i(g)\|.$$

The central displacement bound and the definition of $\Omega_R(\delta)$ therefore give

$$\log|\mathrm{N}\, D_g(P)| \le (9/10+\delta)R + O_P(1). \tag{939}$$

If a term in (934) is counted, then $D_g(P) \in \mathfrak{pq}$. Were it nonzero, ideal inclusion would give

$$\log|\mathrm{N}\, D_g(P)| \ge \log \mathrm{N}(\mathfrak{pq}) \ge sR,$$

contradicting (939) for large R. Thus every counted element fixes $[P]$ exactly.

The identity contributes only $O(e^{sR+o(R)}/R^2)$ to the left side, the number of prime pairs. The conjectural lower bound is $e^{R-o(R)}/R^2$, and $s<1$, so a projectively nonidentity element is counted for large R. This contradicts triviality of the stabilizer. The rational-Borel classification gives the two alternatives. □

This formulation exposes a genuinely whole-group task. Reduction of G_7 at the two primes supplies two independent projective actions; the central lattice ball supplies nearly e^R long elements; and the two Galois Cartan cocycles cut out the norm exponent 9/10. A proof would estimate the second factorial moment of prime divisibility of the single coefficient $D_g(P)$, using a macroscopic factorization of arbitrary group elements. Separate expander estimates still do not give (936); the prospective gain is that each prime variable is below half entropy and may be averaged by a bilinear large sieve. The remaining inverse alternative is again exact: persistent failure for the coherent reductions of P must put mass in the rational Borel B_P. It is important not to overstate this gain. Since the product modulus is larger than the complete Hodge defect range, positivity of (934) already asserts exact closing. The two-prime formulation is not a logically weaker conjecture. It is a better analytic organization of the same theorem: two sub-half-entropy variables are exposed to dispersion before their product crosses the 9/10 Arakelov threshold.

II.20.10 Why even a Ramanujan bound for the thin quotients is not enough

The loss can be seen in the elementary model of a regular nonbacktracking Schreier graph. Let N be the number of vertices and let β be the nonbacktracking growth rate. A Ramanujan estimate for the number $K_n(u,v)$ of length-n paths between two prescribed vertices has the shape

$$K_n(u,v) = \frac{\beta^n}{N} + O\Big(n\beta^{n/2}\Big). \tag{940}$$

Comparing the main term with the absolute error proves positivity only when

$$\frac{\beta^n}{N} \gg n\beta^{n/2}, \qquad \text{hence} \qquad n > (2+o(1))\log_\beta N.$$

At the information-theoretic radius $n = (1+\epsilon)\log_\beta N$, the expected count is N^ϵ, but the Ramanujan error allowed by (940) is $N^{(1+\epsilon)/2+o(1)}$. It is larger for every small ϵ. An L^2 estimate can still prove optimal lifting for almost all targets, which is exactly why the standard theorem has that quantifier.

The same barrier persists for the $(2,7,\infty)$ group after replacing the regular adjacency operator by its nonbacktracking transfer operator. The Hodge-balanced restriction is imposed by adding the two outer Cartan potentials. Thermodynamic large deviations show that the balanced window has full exponential growth, but an absolute spectral estimate still controls a prescribed congruence matrix coefficient only at the square-root scale.

Thus neither property τ, the conjectural Ramanujan bound, nor an almost-all optimal-lifting theorem proves Conjecture II.20.10. What is sufficient is the following genuinely stronger, target-sensitive assertion: for a fixed rational section x without a stabilizer, the Hodge-balanced nonbacktracking kernel from ∞ to x mod $\mathfrak{p}$ is nonzero at radius $(1+\epsilon)\log \mathrm{N}\,\mathfrak{p}$ for at least one split prime in infinitely many dyadic intervals. Equivalently, the only rational section which can lie in the exceptional holes of all these thin Schreier graphs is a periodic section.

This is where the full thin group, rather than its generators, must enter: the statement couples the congruence tower, the complete orbital-growth kernel, and the two Galois Cartan cocycles. The $1/10$ norm margin then performs the final algebraic step unconditionally.

II.20.11 A precise horizontal Hodge local limit statement

The target-sensitive assertion can be stated without referring to a word metric. For a canonical primitive cusp column $v = (a,c)^t \in G_7 e$, put

$$H_i(v) = \max\{|\sigma_i a|, |\sigma_i c|\}.$$

For $R > 0$ and $\delta > 0$, let $\mathcal{C}_R(\delta)$ consist of the cusp columns satisfying

$$H_0(v) \le e^{R/2}, \qquad H_1(v) \le e^{(1/10+\delta)R}, \qquad H_2(v) \le e^{(3/10+\delta)R}. \tag{941}$$

Write $B_R(\delta) = \#\mathcal{C}_R(\delta)$. The exponents are the central cusp exponent and the two correctly labelled Hodge-typical outer exponents.

For a good split prime ideal $\mathfrak{p}$, let $X_\mathfrak{p}$ be the projective orbit of $[e]$ under the reduction of G_7, and define the fixed-target kernel

$$K_{\mathfrak{p},R,\delta}(P) = \#\{v \in \mathcal{C}_R(\delta) : [v] = [P] \text{ in } X_\mathfrak{p}\}. \tag{942}$$

This counts cusp columns, not words; right multiplication by the parabolic stabilizer therefore introduces no artificial multiplicity.

Conjecture II.20.16 (Horizontal Hodge local limit). *Let $P \in \mathcal{O}_K^2$ be primitive and suppose that $[P]$ has trivial projective stabilizer in G_7. For every sufficiently small fixed $\eta, \delta > 0$, with*

$$(1+\eta)(9/10+2\delta) < 1, \tag{943}$$

put $R = (1+\eta)\log Q$. Then

$$\sum_{\substack{Q \le \mathrm{N}\,\mathfrak{p} \le 2Q \\ \mathfrak{p} \text{ good split}}} K_{\mathfrak{p},R,\delta}(P) = (1+o(1))B_R(\delta) \sum_{\substack{Q \le \mathrm{N}\,\mathfrak{p} \le 2Q \\ \mathfrak{p} \text{ good split}}} \frac{1}{\#X_{\mathfrak{p}}}, \tag{944}$$

and $B_R(\delta) = e^{R-o(R)}$.

Since $\#X_{\mathfrak{p}} = (\mathrm{N}\,\mathfrak{p})^{1+o(1)}$, the right side of (944) has order $Q^{1+\eta-o(1)}/\log Q$, and in particular is positive.

Theorem II.20.17 (Conditional classification from the horizontal local limit). *Conjecture II.20.16 implies the full heptagonal cusp/special fixed-point classification.*

Proof. Suppose that a primitive P has trivial projective stabilizer. The local limit supplies, for all sufficiently large Q, a good split prime $\mathfrak{p}$ in the indicated interval and a column $v \in \mathcal{C}_R(\delta)$ such that $[v] = [P] \bmod \mathfrak{p}$. Put $r = \det(P, v)$. Then $r \in \mathfrak{p}$. On the other hand, (941) gives

$$\log|\mathrm{N}\,r| \le (9/10+2\delta)R + O_P(1) = (1+\eta)(9/10+2\delta)\log Q + O_P(1) < \log \mathrm{N}\,\mathfrak{p}$$

for large Q. A nonzero element of $\mathfrak{p}$ has norm at least $\mathrm{N}\,\mathfrak{p}$, so $r = 0$. Thus $[P] = [v]$ is an actual cusp and has a parabolic stabilizer, contradicting the assumption. Hence every rational flag has a nontrivial projective stabilizer. The rational-Borel calculation classifies it as a cusp or a special hyperbolic fixed point. □

The preceding implication is correct, but it also reveals that the local limit conjecture is too close to the desired conclusion to serve as an independent mechanism. Write $\omega_Q(\mathfrak{b})$ for the number of good split prime divisors of $\mathfrak{b}$ with norm in $[Q, 2Q]$. Apart from the harmless projective-unit convention, there is the exact identity

$$\sum_{\substack{Q \le \mathrm{N}\,\mathfrak{p} \le 2Q \\ \mathfrak{p} \text{ good split}}} K_{\mathfrak{p},R,\delta}(P) = \sum_{v \in \mathcal{C}_R(\delta)} \omega_Q\big((\det(P,v))\big), \tag{945}$$

where a zero incidence is counted once for every prime in the outer sum. Under condition (943), every nonzero incidence has

$$|\mathrm{N}\det(P,v)| < Q$$

for all sufficiently large Q. Its contribution to the right side of (945) is therefore zero. Thus positivity of the left side is already equivalent to the existence of an actual cusp column on $[P]$. Averaging over primes has not produced a softer distribution problem.

This also explains why a standard large-sieve or transfer estimate cannot be extrapolated to (944). The congruence target lies beyond the complete norm support of every nonzero incidence. A theorem asserting that the orbit nevertheless hits that target has already proved exact vanishing. The conditional theorem above is a logically correct reduction, but the conjectured local limit should be regarded as a no-holes form of the classification, not as a standard analytic hypothesis suggested by expansion.

There is a familiar analytic strengthening of Conjecture II.20.16. Put $M_{\mathfrak{p}} = B_R(\delta)/\#X_{\mathfrak{p}}$. A pointwise square-root local limit would assert, away from stabilized rational sections,

$$K_{\mathfrak{p},R,\delta}(P) = M_{\mathfrak{p}} + O\big(M_{\mathfrak{p}}^{1/2}(\mathrm{N}\,\mathfrak{p})^{o(1)}\big). \tag{946}$$

At the chosen radius $M_{\mathfrak{p}} = Q^{\eta-o(1)}$, so the main term dominates. The error in (946) is square-root in the *local expected count*; the Ramanujan error in (940) is square-root in the *entire orbital ball*. This is the exact quantitative distinction, but the divisor identity above shows that calling the former a local limit does not make it an ordinary equidistribution statement.

There is an important correction to the most naive transfer-operator formulation. The two outer sizes are not scalar Hölder potentials on the Bowen–Series base. They are norms of matrix cocycles. To make them additive, enlarge the state by two moving projective directions $([u_1],[u_2]) \in \mathbf{P}^1(\mathbb{R})^2$ and put

$$\psi_i(\gamma,[u_i]) = \log \frac{\|\sigma_i(\gamma)u_i\|}{\|u_i\|}.$$

Then

$$\psi_i(\gamma_2\gamma_1,[u_i]) = \psi_i(\gamma_1,[u_i]) + \psi_i(\gamma_2,[\sigma_i(\gamma_1)u_i]),$$

so these weights become additive only on the skew product

$$(\zeta,[u_1],[u_2],y) \longmapsto \big(T\zeta,[\sigma_1(\gamma_\zeta)u_1],[\sigma_2(\gamma_\zeta)u_2],\gamma_\zeta y\big), \qquad y \in X_{\mathfrak{p}}.$$

The relevant operator is therefore a transfer operator on the symbolic coordinate, the two outer projective lines, and the congruence orbit, weighted by central return time and by the two functions ψ_i. Acceleration is also essential: the cusp gives unbounded branches, so the resulting roof and cocycle have tails which must be controlled.

The main pole and a two-variable local limit would give the predicted term. Square-root cancellation after decomposing the finite congruence coordinate into its nontrivial representations would give (946). A rational Borel stabilizer certainly supplies a closed loop with exceptional Hodge holonomy. The converse assertion—that every persistent exceptional matrix coefficient is supported by such a loop—is not formal. It is the remaining spectral-rigidity theorem. Thus the proposed conditional input is a statement about the whole thin group and its congruence tower, not a conjectural inequality for individual generators.

Comparison with existing congruence transfer estimates

The required estimate is substantially sharper, and more structured, than the published uniform congruence estimates. Mayer, Mühlenbruch, and Strömberg construct transfer operators for cofinite Hecke triangle groups, including their cusps [50]; this provides the correct symbolic starting point but does not add the two outer projective cocycles or prove uniform local limits in congruence level. Oh and Winter prove uniform exponential mixing for congruence covers in the convex-cocompact setting [52]; the cuspidal Hecke group considered here lies outside that hypothesis. Magee, Oh, and Winter prove congruence counting for Schottky semigroups with an error of the schematic form

$$\frac{\#B_R}{|G(\mathfrak{q})|} + O\big((\mathrm{N}\,\mathfrak{q})^C \#B_R e^{-\epsilon R}\big) \tag{$*$}$$

[53]. The polynomial level loss in $(*)$ permits levels only exponentially smaller than the orbital size; it does not reach the entropy-scale regime $\mathrm{N}\,\mathfrak{p} \asymp e^R$, and its absolute error is not square-root in the local expected count.

Consequently an entropy-level fixed-target local limit is not ordinary Ramanujan, Lindelöf, or Kuznetsov with familiar names substituted. At the closing scale it contains the exact vanishing assertion. A less circular analytic route must work below that scale and retain a second constraint—for example the real shrinking target in Proposition II.20.66. The viable parameter region has width $1/10$. What the thin group must then prove is a stopped, target-sensitive inverse theorem in that joint window, with a rational Borel as the only coherent obstruction.

A concrete full-thin-group architecture: pointed local–global dispersion

Thin-orbit local–global methods provide a more appropriate model than a pointwise expander bound. Bourgain and Kontorovich study a scalar matrix coefficient on a thin SL_2-orbit by decomposing a long element into macroscopic blocks and applying the circle method, bilinear estimates, and dispersion [68, 69]. Zhang proves a density-one representation theorem in the presence of parabolics [70]. These results are relevant here because they exploit multiplication of whole group blocks to pass beyond what a spectral gap alone can see.

For the fixed-point defect

$$D_g(P) = \det(P, gP)$$

and a block split $g = ab$, the exact identity is

$$D_{ab}(P) = \det(a^{-1}P, bP). \tag{947}$$

Thus the finite Fourier transform of $D_g(P)$ mod $\mathfrak{q}$ is a bilinear exponential sum in two complete macroscopic pieces of G_7. Fourier expansion of a smooth real shrinking target adds one archimedean frequency. The two Hodge cutoffs are imposed by carrying the moving outer projective directions in the block measures. This is an honest use of the thin orbit and its restriction-of-scalars representation, rather than of the six Rosen moves separately.

There is an exact favorable local calculation.

Proposition II.20.18 (Finite defect distribution)**.** *Let $k = \mathbb{F}_q$, let $0 \neq P \in k^2$, and set $D_g(P) = \det(P, gP)$. Then*

$$\#\{g \in \mathrm{SL}_2(k) : D_g(P) = 0\} = q^2 - q, \tag{948}$$

$$\#\{g \in \mathrm{SL}_2(k) : D_g(P) = t\} = q^2 \quad (0 \neq t \in k). \tag{949}$$

Consequently, for every nontrivial additive character ψ of k and every $\xi \neq 0$,

$$\sum_{g \in \mathrm{SL}_2(k)} \psi\big(\xi D_g(P)\big) = -q. \tag{950}$$

Proof. Choose $h \in \mathrm{SL}_2(k)$ with $hP = e_1$ and replace g by hgh^{-1}. For $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, one has $D_g(e_1) = c$. If $c \neq 0$, the q^2 choices of a, d determine $b = (ad-1)/c$. If $c = 0$, choose $a \neq 0$, put $d = a^{-1}$, and choose b arbitrarily, giving $q(q-1)$ matrices. Finally $\sum_{t \neq 0} \psi(\xi t) = -1$, which proves (950). □

The calculation extends without loss to prime powers.

Proposition II.20.19 (Prime-power defect distribution)**.** *Let*

$$R_m = \mathfrak{o}/\varpi^m\mathfrak{o}, \qquad q = \#(\mathfrak{o}/\varpi), \qquad Q = \#R_m = q^m,$$

and let $P \in R_m^2$ be a unimodular column. For $t \in R_m$, put

$$N_m(t) = \#\{g \in \mathrm{SL}_2(R_m) : \det(P, gP) = t\}.$$

Then

$$N_m(t) = \begin{cases} Q^2, & t \in R_m^\times, \\ Q^2(1-q^{-1}), & t \notin R_m^\times. \end{cases} \tag{951}$$

For an additive character ψ of R_m,

$$\sum_{g \in \mathrm{SL}_2(R_m)} \psi\big(\det(P, gP)\big) = \begin{cases} Q^3(1-q^{-2}), & \psi = 1, \\ -Q^3 q^{-2}, & \psi \neq 1 \text{ and } \psi|_{\varpi R_m} = 1, \\ 0, & \psi|_{\varpi R_m} \neq 1. \end{cases} \tag{952}$$

Proof. As before, conjugate P to e_1, so the defect is the lower-left entry c of $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$. For each primitive first column $(a, c)^t$, the solutions $(b, d)^t$ of $ad - bc = 1$ form an affine copy of R_m, and hence there are exactly Q of them. If c is a unit, all Q choices of a give a primitive column. If c is a nonunit, a must be a unit, giving $Q(1-q^{-1})$ choices. This proves (951).

Let

$$T(\psi) = \sum_{c \in R_m} \psi(c), \qquad U(\psi) = \sum_{c \in \varpi R_m} \psi(c).$$

Equation (951) gives

$$\sum_g \psi(\det(P, gP)) = Q^2\big(T(\psi) - q^{-1}U(\psi)\big).$$

If $\psi = 1$, then $T = Q$ and $U = Q/q$. If $\psi \neq 1$, then $T = 0$; moreover $U = Q/q$ precisely when ψ is trivial on ϖR_m, and otherwise $U = 0$. These are exactly the three cases in (952). □

For an arbitrary good ideal $\mathfrak{a} = \prod_{\mathfrak{p}} \mathfrak{p}^{m_{\mathfrak{p}}}$, the Chinese remainder theorem gives

$$\frac{\#\{g \in \mathrm{SL}_2(\mathcal{O}_K/\mathfrak{a}) : D_g(P) = 0\}}{\#\,\mathrm{SL}_2(\mathcal{O}_K/\mathfrak{a})} = \prod_{\mathfrak{p}^m \| \mathfrak{a}} \frac{1}{\mathrm{N}\,\mathfrak{p}^m + \mathrm{N}\,\mathfrak{p}^{m-1}} = \frac{1}{\#\mathbf{P}^1(\mathcal{O}_K/\mathfrak{a})}. \tag{953}$$

The Fourier transform also factors into the local alternatives in (952). In particular, every character which is genuinely ramified beyond the residue field at some prime has zero complete Fourier coefficient. A nontrivial character factoring through one residue field has normalized local coefficient

$$-\frac{1}{q^2-1}.$$

Thus higher prime powers create no new local resonance.

At every good split prime, strong approximation identifies the complete reduction of G_7 with this finite group. The rational Borel therefore has the exact local density $1/(q+1)$, while a nontrivial scalar Fourier coefficient has normalized size $O(q^{-2})$. This rules out local congruence bias as the obstruction. What is missing is uniformity after the finite quotient is coupled to a thin real cylinder and to the two outer projective cocycles.

There is an exact endpoint-uniform version which explains why a middle block is natural. Let $u, v \in R_m^2$ be unimodular and define the bridge defect

$$F_{u,v}(c) = \det(u, cv), \qquad c \in \mathrm{SL}_2(R_m).$$

Choose $h, k \in \mathrm{SL}_2(R_m)$ with $hu = kv = e_1$. Then the change of variables $c' = hck^{-1}$ gives

$$F_{u,v}(c) = \det(e_1, c'e_1). \tag{954}$$

Consequently $F_{u,v}$, for every pair of endpoints u, v, has exactly the value counts and Fourier transform in Propositions II.20.18 and II.20.19. In particular, complete bridge averaging destroys all finite graph correlation uniformly in the outer blocks: the normalized nontrivial residue-field coefficient is $-1/(q^2-1)$, and genuinely higher-conductor coefficients vanish. This is stronger than separate marginal equidistribution. It does not yet solve the problem because the actual bridge is a short thin segment, not the whole quotient, and the cancellation must remain valid after the real and Hodge restrictions are imposed.

The precise amount of cancellation in an incomplete bridge has an exact operator-theoretic description. This is useful because it prevents us from assigning a fictitious square-root gain to the middle block.

Proposition II.20.20 (Bridge Fourier transform equals congruence averaging)**.** *Let $V = k^2$, let $\xi \in k^\times$, and let ν be a probability measure on $\mathrm{SL}_2(k)$. On $\ell^2(V)$ define*

$$\begin{aligned}(\mathcal{F}_\xi f)(u) &= q^{-1}\sum_{v\in V}\psi(\xi\det(u,v))f(v),\\ (\pi(c)f)(w) &= f(c^{-1}w),\\ (\mathcal{T}_{\nu,\xi}f)(u) &= q^{-1}\sum_{c}\nu(c)\sum_{v\in V}\psi(\xi\det(u,cv))f(v).\end{aligned}$$

Then $\mathcal{F}_\xi$ is unitary and

$$\mathcal{T}_{\nu,\xi} = \mathcal{F}_\xi\,\pi(\nu). \tag{955}$$

Consequently

$$\big\|\mathcal{T}_{\nu,\xi}\big|_U\big\| = \big\|\pi(\nu)\big|_U\big\|$$

for every π-invariant subspace U. In particular, on the mean-zero functions on $V\setminus\{0\}$, a scalar bridge Fourier bound is exactly a congruence-expansion bound for the middle measure.

Proof. Additive orthogonality for the nondegenerate alternating form $\det(\cdot,\cdot)$ gives $\mathcal{F}_\xi\mathcal{F}_\xi^* = I$. In the definition of $\mathcal{T}_{\nu,\xi}$, make the change of variable $w = cv$. The inner sum becomes

$$q^{-1}\sum_{w\in V}\psi(\xi\det(u,w))f(c^{-1}w) = \big(\mathcal{F}_\xi\pi(c)f\big)(u).$$

Averaging over c proves (955); unitarity gives the norm identity. □

The off-diagonal correlation of two endpoint pairs has a still stronger exact formula. It identifies the complete inverse variety rather than merely bounding an operator norm.

Proposition II.20.21 (Complete bridge correlation and its singular cone)**.** *Let $k = \mathbb{F}_q$ have odd characteristic, let ψ be a nontrivial additive character, and let $u, u', v, v' \in k^2\setminus\{0\}$. Put*

$$K(u,u';v,v') = vu^tJ - v'u'^tJ, \qquad J = \begin{pmatrix}0 & 1\\ -1 & 0\end{pmatrix}.$$

Then

$$\det K(u,u';v,v') = -\det(u,u')\det(v,v'). \tag{956}$$

Moreover

$$\sum_{c\in \mathrm{SL}_2(k)} \psi\big(\det(u,cv)-\det(u',cv')\big) = \begin{cases} q(q^2-1), & K=0,\\ -q, & K\neq 0 \text{ and } \det K=0,\\ q\,\mathrm{Kl}_\psi(\det K), & \det K\neq 0, \end{cases} \tag{957}$$

where

$$\mathrm{Kl}_\psi(t)=\sum_{z\in k^\times}\psi(z+t/z).$$

In the last case the Weil bound gives

$$\left|\sum_c \psi\big(\det(u,cv)-\det(u',cv')\big)\right| \le 2q^{3/2}. \tag{958}$$

Proof. The phase is

$$\det(u,cv)-\det(u',cv')=\mathrm{tr}(Kc).$$

The determinant of the difference of the two rank-one matrices is the right side of (956); this follows by expanding the two-by-two determinant.

Write $\mathrm{tr}(Kg)=\alpha a+\beta b+\gamma c+\delta d$, so in these coefficient coordinates $K=\left(\begin{smallmatrix}\alpha & \gamma\\ \beta & \delta\end{smallmatrix}\right)$. For nonzero K, insert additive Fourier inversion for the equation $ad-bc=1$:

$$\begin{aligned}\sum_{g\in \mathrm{SL}_2(k)}\psi(\mathrm{tr}(Kg)) &= q^{-1}\sum_{t\in k}\psi(-t)\sum_{a,b,c,d\in k}\psi\big(t(ad-bc)+\alpha a+\beta b+\gamma c+\delta d\big)\\ &= q\sum_{t\in k^\times}\psi\big(-t-\det(K)/t\big)=q\,\mathrm{Kl}_\psi(\det K).\end{aligned}$$

For $\det K=0$ the Kloosterman sum is -1; for nonzero determinant Weil gives absolute value at most $2q^{1/2}$. If $K=0$, the phase is identically zero and the sum is the order of the group. □

The geometric meaning is decisive. The equality $K=0$ says that the two rank-one tensors are identical, and hence that both pairs of projective endpoint lines agree. The rank-one case $\det K=0$, $K\neq 0$, says that one of the two projective endpoint pairs agrees. If $[P]$ has trivial projective stabilizer, equality $[a^{-1}P]=[a'^{-1}P]$ with $a\neq a'$, or $[bP]=[b'P]$ with $b\neq b'$, is impossible. Thus, after the literal diagonal and the two semi-diagonals have been removed, every remaining complete-bridge correlation lies in the nonsingular case (958). The complete finite inverse theorem is therefore already exact: its singular locus is precisely endpoint repetition, which is the desired rational-Borel stabilizer.

The distinguished real place has the same dual variety. For a smooth compactly supported density w on $\mathrm{SL}_2(\mathbb{R})$, consider

$$\widehat{w}_\xi(K)=\int_{\mathrm{SL}_2(\mathbb{R})} w(c)e^{i\xi\,\mathrm{tr}(Kc)}\,dc.$$

The critical-point equation for the phase on $\det c=1$ is $K=\mu\,\mathrm{adj}(c)$. Away from $\det K=0$ its critical points are nondegenerate, and stationary phase gives the three-dimensional decay

$$|\widehat{w}_\xi(K)|\ll_w (1+|\xi|)^{-3/2} \tag{959}$$

on compact subsets of the nonsingular dual region; if the phase has no real critical point, repeated integration by parts is stronger. The loss of uniformity occurs exactly on $\det K = 0$, the same factored cone as at the finite place.

This coincidence sharply narrows the remaining analytic work. It is not necessary to classify an unknown correlation variety. One must transfer the complete Kloosterman and real stationary-phase cancellation to the short thin bridge, using the two long block variables, while controlling the neighborhoods of

$$\det(u, u') \det(v, v') = 0.$$

The product formula and the two Hodge cutoffs are naturally suited to that last task: a nonzero algebraic determinant cannot be simultaneously too small in all three embeddings. What is not yet proved is the off-diagonal completion estimate for the thin bridge.

There is nevertheless an immediate global consequence which is invisible in a calculation with the six letters. It removes the singular finite cone from an average over external primes.

Proposition II.20.22 (A genuine off-diagonal pair is singular at only boundedly many horizontal primes)**.** *Fix a primitive $P \in \mathcal{O}_K^2$. Let $a, a', b, b' \in G_7$ have distinguished displacement at most R, and put*

$$Z_{a,a';b,b'}(P) = \det(a^{-1}P, a'^{-1}P)\, \det(bP, b'P) \in \mathcal{O}_K. \tag{960}$$

Assume that $[P]$ has trivial projective stabilizer and that $a \neq a'$, $b \neq b'$. Then $Z_{a,a';b,b'}(P) \neq 0$, and there is a constant C_P, independent of the four group elements and of R, such that

$$\log\bigl|\mathrm{N}\, Z_{a,a';b,b'}(P)\bigr| \leq C_P R + O_P(1). \tag{961}$$

Consequently, for any fixed $r > 0$, at most $C_P/r + O_P(1/R)$ prime ideals of norm at least e^{rR} can divide $Z_{a,a';b,b'}(P)$. Among good split prime ideals with norm in a fixed dyadic interval $[e^{rR}, 2e^{rR}]$, the proportion for which the reduced complete bridge correlation is singular is $e^{-rR+o(R)}$.

Proof. If the first determinant in (960) vanished, then $[a^{-1}P] = [a'^{-1}P]$, so $a'a^{-1}$ would fix $[P]$. Triviality of the projective stabilizer and $a \neq a'$ exclude this. The same argument applies to b, b'.

The entries of every matrix of distinguished displacement at most R, and of its inverse, have size $e^{O(R)}$ in each of the three real embeddings. This follows either from the modular-embedding distance-decreasing property, or from the operator-norm bound in each Galois factor. Hence every conjugate of each determinant in (960) is $e^{O_P(R)}$, proving (961).

If m distinct prime ideals of norm at least e^{rR} divide the nonzero principal ideal (Z), then $e^{mrR} \leq |\mathrm{N}\, Z| \leq e^{C_P R + O_P(1)}$, and therefore $m \leq C_P/r + O_P(1/R)$. Chebotarev's theorem and the prime-ideal theorem give $e^{rR+o(R)}/R$ good split primes in the stated dyadic interval. At every prime not dividing Z, both endpoint determinants are nonzero, so (956) places the complete correlation in the nonsingular Weil case. Dividing a bounded number of exceptional primes by the size of the interval proves the last assertion. □

This is the first useful separation between a fixed off-diagonal quadruple and the horizontal prime average. It says that the algebraic inverse obstruction has already disappeared at almost every prime; one does not need a newly chosen factor extracted from a remainder. The primes are chosen externally and all finite valuations are controlled at once by the norm of the single global discriminant Z. What remains is an analytic incompleteness problem: the bridge measure occupies only $e^{\tau R}$ elements of a quotient of size about e^{3rR}, so the complete Kloosterman estimate cannot be inserted term by term. The two long blocks must provide the missing completion.

The same norm argument gives a useful averaged flattening theorem for an entire long block. It is worth stating because it is a genuine theorem about the full orbit, and because its sharp exponent also shows why one more use of Cauchy–Schwarz cannot finish the proof.

Proposition II.20.23 (Horizontal endpoint energy of a long thin block)**.** *Let $\mathcal{A}_R \subset G_7$ consist of distinct elements of displacement at most R, let $A = \#\mathcal{A}_R$, and assume that $[P]$ has trivial projective stabilizer. For a good split prime $\mathfrak{p}$, put*

$$\alpha_{\mathfrak{p}} = \frac{1}{A} \sum_{a \in \mathcal{A}_R} \delta_{[a^{-1}P \bmod \mathfrak{p}]} \quad \text{on} \quad X_{\mathfrak{p}} = \mathbf{P}^1(\mathcal{O}_K/\mathfrak{p}).$$

If $\mathcal{P}_R(r)$ is the set of good split primes with norm in $[e^{rR}, 2e^{rR}]$, then

$$\frac{1}{\#\mathcal{P}_R(r)} \sum_{\mathfrak{p} \in \mathcal{P}_R(r)} \|\alpha_{\mathfrak{p}}\|^2_{\ell^2(X_{\mathfrak{p}})} \le \frac{1}{A} + \frac{C_P}{r\,\#\mathcal{P}_R(r)}. \tag{962}$$

In particular, if $A = e^{rR+o(R)}$, the average collision energy is $e^{-rR+o(R)}$, the same exponential scale as the uniform measure on $X_{\mathfrak{p}}$.

Proof. Expanding the square gives

$$A^2\|\alpha_{\mathfrak{p}}\|_2^2 = A + \#\{(a,a') \in \mathcal{A}_R^2 : a \neq a',\ \mathfrak{p} \mid \det(a^{-1}P, a'^{-1}P)\}.$$

For $a \neq a'$, the determinant is nonzero: otherwise $a'a^{-1}$ would be a projectively nonidentity stabilizer of $[P]$. Its norm has logarithm $O_P(R)$, by the same modular-embedding estimate as in Proposition II.20.22. It is therefore divisible by at most C_P/r primes in $\mathcal{P}_R(r)$. Sum the displayed identity over $\mathfrak{p}$ and divide by $A^2\#\mathcal{P}_R(r)$. □

The preceding energy estimate becomes more informative after one uses the exact representation carried by the projective congruence orbit. This is a feature of the whole quotient of the thin group, not of its chosen generators. For a good split prime write

$$G_{\mathfrak{p}} = \mathrm{PSL}_2(\mathcal{O}_K/\mathfrak{p}), \qquad X_{\mathfrak{p}} = \mathbf{P}^1(\mathcal{O}_K/\mathfrak{p}), \qquad q = \mathrm{N}\,\mathfrak{p}.$$

Strong approximation identifies the image of G_7 with $G_{\mathfrak{p}}$. The action on $X_{\mathfrak{p}}$ is doubly transitive, and hence its permutation representation has the exact decomposition

$$\ell^2(X_{\mathfrak{p}}) = \mathbf{1} \oplus \mathrm{St}_{\mathfrak{p}}, \tag{963}$$

where $\mathrm{St}_{\mathfrak{p}}$ is the q-dimensional Steinberg representation. Thus there is only one nonconstant finite representation in the pointed lifting problem.

Proposition II.20.24 (The pointed discrepancy is one Steinberg coefficient)**.** *Let $\Omega \subset G_7$ be finite, let μ_Ω be its normalized counting measure, and put*

$$e_y = \delta_y - \frac{\mathbf{1}}{q+1} \qquad (y \in X_{\mathfrak{p}}).$$

For every primitive $P \in \mathcal{O}_K^2$, one has the exact identity

$$\frac{1}{\#\Omega}\#\{g \in \Omega : g\infty = [P \bmod \mathfrak{p}]\} = \frac{1}{q+1} + \left\langle \mathrm{St}_{\mathfrak{p}}(\mu_\Omega)e_\infty, e_{[P \bmod \mathfrak{p}]} \right\rangle. \tag{964}$$

In particular, all of the finite-quotient error is the single displayed matrix coefficient.

Proof. For $g \in G_{\mathfrak{p}}$,

$$\langle g\delta_\infty, \delta_y \rangle = \mathbf{1}_{\{g\infty = y\}}.$$

Subtract the orthogonal projection onto the constant functions and average over $g \in \Omega$. The orthogonal complement is precisely $\mathrm{St}_{\mathfrak{p}}$ by (963). □

This identity changes the analytic target. There is no need to prove a spectral statement simultaneously for every irreducible representation of a large finite quotient. Let $\Omega_R(\delta)$ be a central shell of Hodge-balanced elements,

$$R(g) \in [R, R+1], \qquad a_1(g) + a_2(g) \le (2/5 + \delta)R, \qquad \#\Omega_R(\delta) = e^{R - o(R)}.$$

If $q = e^{sR + o(R)}$ with $9/10 + \delta < s < 1$, the desired main count in (964) is $e^{(1-s)R - o(R)}$. A square-root counting error would be

$$\#\{g \in \Omega_R(\delta) : g\infty = [P] \bmod \mathfrak{p}\} = \frac{\#\Omega_R(\delta)}{q+1} + O_\epsilon\left(\left(\frac{\#\Omega_R(\delta)}{q}\right)^{1/2} q^\epsilon\right). \tag{965}$$

Because $s < 1$, its error is smaller than the main term. In normalized form it says

$$\left\langle \mathrm{St}_{\mathfrak{p}}(\mu_{R,\delta}) e_\infty, e_{[P \bmod \mathfrak{p}]} \right\rangle = O_\epsilon((q \#\Omega_R(\delta))^{-1/2} q^\epsilon) = o(q^{-1}).$$

This is strictly more precise than asking for an expander eigenvalue. An operator-norm estimate pays for the worst vector in a q-dimensional space and encounters the diagonal barrier above. Equation (965) asks only for the one relative coefficient between two Borel-fixed vectors. It is the natural output of a relative trace formula, or equivalently of a Steinberg-twisted transfer operator with the two Hodge weights retained.

There is also an exact inverse formulation. Suppose $[P]$ has no projectively nontrivial G_7-stabilizer. For every Hodge-balanced g in the displayed shell, the product formula gives

$$|\,\mathrm{N}\det(P, ge)| \le e^{(9/10+\delta)R + O_P(1)}.$$

At a prime with $q = e^{sR + o(R)}$, $s > 9/10 + \delta$, a congruence hit would therefore make this determinant zero and make P an actual cusp. Under the assumed trivial stabilizer there are no hits, so (964) forces the extremal identity

$$\left\langle \mathrm{St}_{\mathfrak{p}}(\mu_{R,\delta}) e_\infty, e_{[P \bmod \mathfrak{p}]} \right\rangle = -\frac{1}{q+1}. \tag{966}$$

Consequently the global theorem is reduced to the following representation-theoretic inverse statement: persistence of (966) for the coherent algebraic vectors $[P \bmod \mathfrak{p}]$, after imposing the Hodge-balanced shell, forces $G_7 \cap B_P(K) \neq \{1\}$. The exceptional subgroup is not guessed after the calculation. The vector $e_{[P \bmod \mathfrak{p}]}$ is fixed by the finite Borel $B_P(\mathbf{F}_q)$, and its characteristic-zero lift is exactly the rational Borel whose intersection with G_7 is the desired cusp or special stabilizer.

This is the sharpest current meaning of "use the thin group." The free product supplies the complete geodesic shells, strong approximation supplies $G_{\mathfrak{p}}$, double transitivity collapses the finite problem to one Steinberg representation, and the two modular embeddings select the Hodge-balanced part of that shell. A table of the six Rosen moves uses none of these four structures.

The exact relative-Lindelöf statement

The term *relative square-root error* can now be normalized without ambiguity. Let $\mathscr{H}$ be the accelerated Bowen–Series branches for G_7, and retain the two moving outer projective directions. For a good split prime define the Steinberg–Hodge transfer operator

$$(\mathscr{L}_{z,\mathbf{v},\mathfrak{p}}F)(\zeta,[u_1],[u_2]) = \sum_{h\in\mathscr{H}} |h'(\zeta)|^z \exp(v_1\psi_1(h,[u_1]) + v_2\psi_2(h,[u_2])) \\ \cdot \mathrm{St}_{\mathfrak{p}}(h)F(h\zeta,[\sigma_1(h)u_1],[\sigma_2(h)u_2]). \tag{967}$$

Here z is dual to central hyperbolic time and $\mathbf{v}$ is dual to the two outer Cartan cocycles. The parabolic acceleration is essential: it makes $\mathscr{L}$ nuclear in a right half-plane while retaining the cusp tail. Fast and slow transfer operators for Hecke triangle groups with arbitrary finite-dimensional unitary twists are constructed by Adam–Pohl [51]; the additional work here is the uniform growing Steinberg twist together with the two nonunitary Hodge variables.

The Hodge factors in (967) are themselves representations, not auxiliary scalar decorations. If $\rho_i = \sigma_i|_{G_7}$ and

$$\psi_i(g,[u]) = \log\frac{\|\rho_i(g)u\|}{\|u\|},$$

then the cocycle identity

$$\psi_i(gh,[u]) = \psi_i(g,[\rho_i(h)u]) + \psi_i(h,[u])$$

shows that

$$[\Pi_{i,v}(g)f]([u]) = \exp\bigl(v\psi_i(g^{-1},[u])\bigr)\, f([\rho_i(g^{-1})u]) \tag{968}$$

defines a projective principal-series representation. The central Jacobian factor is the analogous principal series for ρ_0. After the standard half-density normalization, the axis

$$\Re z = \frac{1}{2}, \qquad v_1, v_2 \in i\mathbb{R}$$

is unitary. Thus the group part of the full transfer operator is a convolution operator in the genuine tensor representation

$$\mathrm{St}_{\mathfrak{p}} \otimes \Pi_{0,z} \otimes \Pi_{1,v_1} \otimes \Pi_{2,v_2}. \tag{969}$$

The Hodge local limit is analytic continuation away from this unitary axis, and its mean $2/5$ is the derivative of the corresponding pressure.

This tensor form identifies the algebraic exceptional set. The K-rational line P gives the Galois endpoint

$$P^{\Sigma} = ([\sigma_0 P],[\sigma_1 P],[\sigma_2 P]) \in (\mathbf{P}^1)^3,$$

whose stabilizer in $\mathrm{Res}_{K/\mathbb{Q}}\mathrm{SL}_2$ is the $\mathbb{Q}$-group $\mathrm{Res}_{K/\mathbb{Q}} B_P$. Consequently

$$G_7 \cap \mathrm{Stab}(P^{\Sigma}) = G_7 \cap B_P(K).$$

The Galois Zariski closure SL_2^3 rules out replacing (969) by a hidden diagonal factor between two embeddings. Hence the exceptional clause in the relative Lindelöf conjecture is the common stabilizer of the actual tensor endpoints, rather than an exceptional subgroup inserted by analogy.

Let Φ_∞ and $\Phi_{P,\mathfrak{p}}$ be smooth symbolic endpoint states whose finite components are respectively e_∞ and $e_{[P \bmod \mathfrak{p}]}$. The Laplace transform of the pointed discrepancy is the relative resolvent coefficient

$$\mathcal{Z}_{P,\mathfrak{p}}(z,\mathbf{v}) = \left\langle (I - \mathscr{L}_{z,\mathbf{v},\mathfrak{p}})^{-1}\Phi_\infty, \Phi_{P,\mathfrak{p}} \right\rangle. \tag{970}$$

This is the precise object on which a Lindelöf assertion is needed. The relevant normalization is not merely $\mathcal{Z} \ll q^{\epsilon}$. The finite endpoint vectors are Borel-fixed newvectors in the Steinberg representation, and the required local factor is $q^{-1/2}$.

There is a noncircular way to state the exceptional spectrum. Let $\mathcal{E}_{P,\mathfrak{p}}(z, \mathbf{v})$ be the sum of the polar parts in the relative resolvent contributed by nonidentity elements of

$$G_7 \cap B_P(K).$$

For a parabolic stabilizer this is the usual cusp term. For a special hyperbolic stabilizer it is the geometric series in the primitive holonomy,

$$\frac{A_{\gamma,\mathfrak{p}}(\mathbf{v})e^{-z\ell(\gamma)}}{1 - A_{\gamma,\mathfrak{p}}(\mathbf{v})e^{-z\ell(\gamma)}},$$

up to the two orientations and the smoothing factor. The exact numerator is the relative matrix coefficient of the primitive element; what matters here is that it is explicitly determined once $\gamma \in G_7 \cap B_P(K)$ is given. Put $\mathcal{E}_{P,\mathfrak{p}} = 0$ when the projective stabilizer is trivial.

Conjecture II.20.25 (Borel-relative Steinberg–Hodge Lindelöf)**.** *Fix a primitive $P \in \mathcal{O}_K^2$. Uniformly for good split primes, for $\mathbf{v}$ in the Fourier neighborhood used to localize*

$$a_1(g) + a_2(g) \le (2/5 + \delta)R,$$

and on the square-root line $\Re z = 1/2 + \epsilon$, the smoothed relative resolvent satisfies

$$\mathcal{Z}_{P,\mathfrak{p}}(z, \mathbf{v}) - \mathcal{E}_{P,\mathfrak{p}}(z, \mathbf{v}) \ll_{P,\epsilon,A} q^{-1/2+\epsilon}(1 + |\Im z|)^{-A} \tag{971}$$

for every $A > 0$, after the usual smoothing in the central time.

This is the standard relative-trace shape: all nondecaying geometric terms are indexed by actual Borel orbits, and the remainder has square-root size. It does not define an exceptional point by failure of the desired estimate. A nonidentity element of $G_7 \cap B_P(K)$ supplies a periodic Borel period and hence a genuine pole or resonance; the conjecture says that there are no additional main terms.

Proposition II.20.26 (Relative Lindelöf has exactly the closing strength)**.** *Conjecture II.20.25 implies (965) whenever $G_7 \cap B_P(K)$ is projectively trivial. In all cases it implies the heptagonal cusp/fixed-point classification.*

Proof. Smooth Perron inversion in z, followed by Fourier inversion in the two Hodge variables, moves the nonconstant contour from the growth pole $\Re z = 1$ to $\Re z = 1/2 + \epsilon$. The factor $q^{-1/2+\epsilon}$ in (971) gives the error

$$q^{-1/2+\epsilon}e^{(1/2+\epsilon)R} = \left(\frac{e^R}{q}\right)^{1/2} q^{O(\epsilon)}$$

for a shell of size $e^{R-o(R)}$. Standard large deviations around the pressure minimum show that imposing $a_1 + a_2 \le (2/5 + \delta)R$ removes only an exponentially smaller portion; the same contour argument with the two Fourier variables gives (965).

Choose $9/10 + \delta < s < 1$ and $q = e^{sR+o(R)}$. The main term e^R/q dominates this error. Hence there is a congruence hit. The product-formula estimate above makes the corresponding determinant zero, so P is a cusp. If the relative resolvent has a nonzero exceptional term, its definition gives $G_7 \cap B_P(K) \neq \{1\}$, and the rational-Borel classification gives a cusp or a special hyperbolic fixed point. □

This calculation also locates exactly what familiar spectral hypotheses do and do not provide. A uniform congruence-transfer estimate of Magee–Oh–Winter type has schematic error

$$q^C e^{(1-\eta)R};$$

at $q = e^{sR}$ it would require $\eta > (C+1)s$, outside the content of a fixed spectral gap. Even an operator-norm Ramanujan estimate gives $e^{R/2+o(R)}$, which is still larger than the main term when $s > 1/2$. The relative Lindelöf statement is stronger by precisely $q^{-1/2}$. This is the finite local Borel-period factor, not an improvement of the archimedean Ramanujan exponent.

Equivalently, the remaining analytic question is a fixed-newvector problem. Plancherel over all $q+1$ projective targets can produce this factor on average, which explains the available almost-all lifting theorems. The coherent vectors $e_{[P \bmod \mathfrak{p}]}$ require the pointwise relative estimate, and the inverse clause says that its only persistent concentration is the characteristic-zero Borel B_P. This is the exact extra content which an ordinary transfer-operator spectral gap omits.

Bruhat–Fourier form of the Steinberg factor

The local factor $q^{-1/2}$ is not numerology. It is already visible in the exact finite representation once one diagonalizes the whole parabolic subgroup. This uses the rank-one Bruhat structure, rather than an inequality for a preferred generator.

Proposition II.20.27 (The cusp Fourier basis and the Kloosterman Weyl matrix)**.** *Let $\mathfrak{p}$ be a good completely split prime, let $k = \mathcal{O}_K/\mathfrak{p} = \mathbf{F}_q$, and choose a nontrivial additive character $\psi : k \to \mathbf{C}^\times$. Realize St_q as $\ell^2_0(\mathbf{P}^1(k))$. Put*

$$\phi_0 = \sqrt{\frac{q+1}{q}}\left(\delta_\infty - \frac{\mathbf{1}}{q+1}\right)$$

and, for $\xi \in k^\times$, put

$$\phi_\xi(x) = q^{-1/2}\psi(\xi x) \quad (x \in k), \qquad \phi_\xi(\infty) = 0.$$

Then $\{\phi_\xi : \xi \in k\}$ is an orthonormal basis of St_q. For

$$u_t = \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}, \qquad w = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix},$$

one has

$$\mathrm{St}_q(u_t)\phi_0 = \phi_0, \qquad \mathrm{St}_q(u_t)\phi_\xi = \psi(-\xi t)\phi_\xi, \tag{972}$$

$$\langle \mathrm{St}_q(w)\phi_\xi, \phi_\eta\rangle = \frac{1}{q}\sum_{x\in k^\times} \psi\left(-\frac{\xi}{x} - \eta x\right) \quad (\xi, \eta \neq 0). \tag{973}$$

Moreover,

$$\max_{\xi,\eta\in k} |\langle \mathrm{St}_q(w)\phi_\xi, \phi_\eta\rangle| \le \frac{2}{\sqrt{q}}. \tag{974}$$

Proof. The nonzero-frequency functions are orthonormal by additive orthogonality, have sum zero on k, and are orthogonal to ϕ_0. Together with ϕ_0 they give q orthonormal vectors, the dimension of the Steinberg representation. Translation by t proves (972). On $k^\times$, the Weyl element acts by $x \mapsto -1/x$, which gives (973). Weil's bound for the nondegenerate Kloosterman sum gives $2q^{-1/2}$.

It remains to check the row or column indexed by zero. The vector ϕ_0 is the normalized centered delta function at ∞, and $\mathrm{St}_q(w)\phi_0$ is the normalized centered delta function at 0. Their inner product is $-1/q$. Its inner product with any ϕ_ξ, $\xi \neq 0$, has absolute value $\sqrt{q+1}/q$. These values also satisfy (974). □

The same basis gives a sharp quantitative statement for incomplete cusp averages. For a probability measure ν on k, set

$$D_\nu = \sum_{t\in k} \nu(t)\,\mathrm{St}_q(u_t), \qquad E(\nu) = \sum_{t\in k} |\nu(t)|^2,$$

and let P_* be orthogonal projection onto $\mathrm{span}\{\phi_\xi : \xi \neq 0\}$, the complement of the Borel-fixed line.

Proposition II.20.28 (Two incomplete cusp blocks)**.** *For probability measures ν, μ on k,*

$$\|P_* D_\nu\, \mathrm{St}_q(w)\, D_\mu P_*\|_{\mathrm{HS}} \le 2\sqrt{qE(\nu)E(\mu)}. \tag{975}$$

In particular, if ν, μ are uniform on sets of sizes N, M, respectively, the right side is $2\sqrt{q/(NM)}$.

Proof. In the basis above,

$$D_\nu \phi_\xi = \widehat{\nu}(\xi)\phi_\xi \quad (\xi \neq 0), \qquad D_\nu \phi_0 = \phi_0.$$

By (974), the squared Hilbert–Schmidt norm on the nonzero-frequency subspace is at most

$$\frac{4}{q}\left(\sum_{\eta\neq 0} |\widehat{\nu}(\eta)|^2\right)\left(\sum_{\xi\neq 0} |\widehat{\mu}(\xi)|^2\right).$$

Additive Parseval gives $\sum_\xi |\widehat{\nu}(\xi)|^2 = qE(\nu)$, and similarly for μ. □

Thus two incomplete parabolic ranges yield genuine contraction once $NM \gg q$. The estimate is an uncertainty principle: a measure and its Weyl transform cannot both remain concentrated away from the Borel-fixed mode. It also gives another exact barrier. A parabolic range of size N costs central displacement $2\log N + O(1)$. Two such ranges inside total displacement R satisfy $NM \le e^{R/2+O(1)}$, whereas the closing prime has $q = e^{sR}$, $s > 9/10$. Hence a single pair of cusp averages is far below the threshold in (975). The missing gain must come from off-diagonal cancellation across many changes of frame, not from making one parabolic block longer.

The components omitted by P_* have rank at most two and pass through $\mathbf{C}\phi_0$. They are therefore not an uncontrolled analytic error: they are precisely the Borel channel which must become the stabilizer alternative. A successful multiblock argument must contract the P_*-part cumulatively and show that persistent mass in the ϕ_0-channel exactifies to one rational Borel.

For a completely split prime, the powers of the heptagonal parabolic T reduce onto the entire group $U(k) = \{u_t : t \in k\}$, because $\lambda \bmod \mathfrak{p} \neq 0$. Thus an accelerated cusp block is diagonal in the basis of Proposition II.20.27; changing cusp frame inserts a Kloosterman-flat Weyl matrix. Indeed, the transition between two distinct flags lies in $B_1 w B_2$; the two Borel factors only add unipotent phases and permute the nonzero Fourier frequencies, so they preserve the bound (974). Conjugating gives the same description at every cusp in the full G_7-orbit.

This proves that the finite local $q^{-1/2}$ factor required by (971) is available at each *complete* change of cusp frame. It does not by itself prove the relative resolvent bound. Between two flat Weyl matrices lie incomplete diagonal sums

$$D_\nu(\xi) = \sum_n \nu(n)\psi(-\xi n\lambda),$$

where ν is selected simultaneously by central stopping time and the two outer Hodge cocycles. Products of unitary diagonal matrices and flat matrices can have norm one; estimating each factor separately would again lose the correlation.

The proposition nevertheless identifies the inverse mechanism much more tightly. The only Fourier vector fixed by the complete cusp unipotent is ϕ_0, the Borel-fixed line. Therefore a failure to realize Kloosterman flatness through many frame changes must keep returning mass to the successive conjugates of this one line. Exact repetition of such a line says that the intervening group element lies in its rational Borel. The analytic task is now a quantitative version: if the Hodge-stopped product persistently has the extremal coefficient $-1/(q+1)$, its mass must concentrate in a chain of Borel-fixed modes; the global product formula must then exactify that chain to $G_7 \cap B_P(K) \neq \{1\}$.

Route II.20.29 (Outcome of the Bruhat–Fourier reduction)**.** The square-root local factor is proved, not conjectured: it is Weil cancellation in the Weyl matrix of the Steinberg representation. The unproved part is the transfer of this complete local cancellation through the incomplete, Hodge-correlated diagonal cusp sums. This distinguishes the genuinely deep step from the elementary finite calculation and replaces a vague appeal to thin-group expansion by a specific stopped Kloosterman inverse theorem.

There is an exact operator consequence. Let $\mathcal{B}_R$ be a second long block, with endpoint probability $\beta_{\mathfrak{p}}$, and let $\nu_{\mathfrak{p}}$ be a middle-block probability on the finite quotient. If

$$I_{\mathfrak{p}} = \langle \alpha_{\mathfrak{p}}, \pi(\nu_{\mathfrak{p}})\beta_{\mathfrak{p}} \rangle, \qquad \rho_{\mathfrak{p}} = \|\pi(\nu_{\mathfrak{p}})|_{\ell^2_0(X_{\mathfrak{p}})}\|,$$

then

$$\left| I_{\mathfrak{p}} - \frac{1}{\#X_{\mathfrak{p}}} \right| \le \rho_{\mathfrak{p}} \left\| \alpha_{\mathfrak{p}} - \frac{\mathbf{1}}{\#X_{\mathfrak{p}}} \right\|_2 \left\| \beta_{\mathfrak{p}} - \frac{\mathbf{1}}{\#X_{\mathfrak{p}}} \right\|_2 . \tag{976}$$

After averaging over primes and using Cauchy–Schwarz with (962), two blocks of size $e^{LR+o(R)}$ give an error of size

$$\sup_{\mathfrak{p}} \rho_{\mathfrak{p}}\, e^{-\min\{L,r\}R+o(R)}.$$

At the one-prime closing scale, write

$$r = s, \qquad L = \frac{s}{2}, \qquad \tau = 1 - s, \qquad 9/10 < s < 1.$$

The expected incidence is $e^{-sR+o(R)}$. Thus (976) can dominate its error only if

$$\rho_{\mathfrak{p}} \le e^{-(s/2+\epsilon)R}. \tag{977}$$

The middle block has length only $\tau R < R/10$. Even the formal best possible entropy-rate contraction $e^{-\tau R}$, and a fortiori a Ramanujan-rate contraction $e^{-\tau R/2}$, is much weaker than (977). Hence horizontal prime averaging proves optimal endpoint flattening in the exponential scale, but ordinary L^2 composition still loses the diagonal factor $e^{sR/2}$. This is the same obstruction as the three-block diagonal calculation, now derived directly from a whole-orbit arithmetic energy estimate.

The finite problem at these parameters has a particularly simple form. For a typical prime $q = \mathrm{N}\,\mathfrak{p} = e^{sR+o(R)}$, the two endpoint sets have sizes

$$|U| = |V| = q^{1/2+o(1)}, \qquad |C| = q^{(1-s)/s+o(1)}, \qquad 0 < \frac{1-s}{s} < \frac{1}{9},$$

and the return count is

$$\sum_{c \in C} |U \cap cV|. \tag{978}$$

Its random-model mean is $|C|$. This is not an ordinary upper-incidence problem. For example, the Warren–Wheeler energy estimate for arbitrary Möbius transformations [71] gives, in the balanced notation $N = |U| = |V|$,

$$I(U \times V, C) \ll N^{6/5}|C|^{3/5}E(C)^{1/10} + N^{1/2}|C|, \qquad E(C) = \#\{c_1c_2^{-1} = c_3c_4^{-1}\}.$$

Even at the minimal energy scale $E(C) \asymp |C|^2$, this is

$$\ll q^{3/5}|C|^{4/5} + q^{1/4}|C|,$$

far above the desired main term $|C|$; more fundamentally it is an upper-tail estimate and cannot force (978) to be nonzero. Published Möbius-incidence technology therefore does not supply the missing lower incidence. What is required is a thin-orbit variance or inverse theorem which uses that U, V, C are three correlated macroscopic pieces of one integral Galois orbit, and which identifies persistent avoidance with B_P.

For uniform ν, the operator $\pi(\nu)$ is projection onto the functions constant on the two $\mathrm{SL}_2(k)$-orbits $\{0\}$ and $V \setminus \{0\}$. This recovers the endpoint-independent coefficient $-1/(q^2-1)$. For a short bridge, however, Proposition II.20.20 says that the coefficient decays only as fast as the actual congruence random walk. Even a Ramanujan-sized estimate

$$\left\| \pi(\nu_{\mathcal{C}_R})\big|_{\ell^2_0} \right\| \le e^{-(\tau/2)R+o(R)}$$

cannot make a single bridge hit a prescribed projective class of mass e^{-rR}, because pointwise positivity would require $\tau/2 > r$, whereas here $\tau < 1 - s$ and $r = s/2 > 9/20$. The two long blocks are therefore not optional: their bilinear dispersion must amplify the bridge expansion and simultaneously average the real frequency.

The factorization also identifies the correct joint operator. If ρ is any unitary representation of G_7 on a Hilbert space $\mathcal{H}$, then

$$\mathcal{T}_{\nu,\xi,\rho} = (\mathcal{F}_\xi \otimes I_{\mathcal{H}}) \sum_c \nu(c)\big(\pi(c) \otimes \rho(c)\big), \tag{979}$$

and therefore

$$\|\mathcal{T}_{\nu,\xi,\rho}\| = \left\| \sum_c \nu(c)\big(\pi(c) \otimes \rho(c)\big) \right\|.$$

After the real shrinking cutoff is decomposed into frequencies, ρ is the corresponding boundary representation (or, in the accelerated coding, its transfer-operator model). Laplace transforms of the two Hodge cutoffs attach the two projective cocycle weights. Thus the actual middle operator acts simultaneously on

$$\text{(finite vectors)} \otimes \text{(real boundary data)} \otimes \text{(two outer Hodge directions)}.$$

This is the concrete representation of the whole thin group which the proof has to estimate.

It is important that the norm in (979) is not bounded by multiplying separate congruence and real spectral norms. Tensor products can retain correlations invisible in either marginal. In the graph example above, this retained tensor component is exactly the obstruction. The desired three-block estimate must show that the two long endpoint measures have negligible projection onto every such component, unless their incidence already gives an exact rational Borel return.

Put

$$r = \frac{\log \#X_{\mathfrak{p}}}{R}$$

and let β be the distinguished real shrinking rate. The natural per-prime count has exponent

$$1 - r - \beta. \tag{980}$$

The norm-closing calculation has exponent

$$\frac{9}{10} + 2\delta - r - \beta. \tag{981}$$

Consequently the same proposed theorem is statistically nonvacuous and arithmetically closing precisely in the interval

$$\frac{9}{10} + 2\delta < r + \beta < 1. \tag{982}$$

Only the sum $r + \beta$ matters to this budget. Moving entropy from the congruence condition to the real cylinder, or back, never enlarges the available width $1/10 - 2\delta$.

Conjecture II.20.30 (Pointed Hodge local–global dichotomy)**.** *Fix parameters in* (982)*. For a primitive $P \in \mathcal{O}_K^2$, one of the following holds:*

1. *$[P]$ has a projectively nontrivial stabilizer in G_7;*
2. *averaged over good split primes with $\log \mathrm{N}\,\mathfrak{p} = rR + O(1)$, the number of g in a central shell of radius R satisfying*

$$g[P] = [P] \pmod{\mathfrak{p}}, \qquad d_0(g[P], [P]) \ll e^{-\beta R}, \qquad a_1(g) + a_2(g) \le (2/5 + \delta)R$$

is $e^{(1-r-\beta-o(1))R}$ per prime.

Theorem II.20.31 (Conditional completion from pointed local–global)**.** *Conjecture II.20.30 implies the full heptagonal cusp/special fixed-point classification.*

Proof. If the first alternative holds, the rational-Borel proposition gives a cusp or a special hyperbolic fixed point. Suppose instead that the second alternative holds. Its positive lower bound supplies g and $\mathfrak{p}$. Put $D = D_g(P)$. Projective congruence gives $D \in \mathfrak{p}$. The chordal-distance identity for the wedge product and the central norm bound give

$$\log |\sigma_0 D| \le (1/2 - \beta)R + O_P(1).$$

The outer Cartan condition gives

$$\log |\sigma_1 D| + \log |\sigma_2 D| \le (2/5 + \delta)R + O_P(1).$$

Consequently

$$\log |\,\mathrm{N}\, D| \le (9/10 + \delta - \beta)R + O_P(1) < rR + O_P(1) = \log \mathrm{N}\,\mathfrak{p} + O_P(1).$$

The strict inequality follows from (982). For large R, a nonzero element of $\mathfrak{p}$ cannot satisfy it. Hence $D = 0$, so g projectively fixes $[P]$, contradicting the second alternative. Thus the first alternative always holds, and the rational-Borel proposition completes the classification. □

Unlike the entropy-level horizontal local limit, each of its two target restrictions can be placed inside its separate natural support; the contradiction uses their joint cost and the product formula.

The cited local–global theorems still do not prove the conjecture. Their conclusion concerns a density-one set of admissible scalar values, whereas zero is one prescribed singular value. A special hyperbolic line is locally admissible modulo every ideal and yet is not a cusp, so the singular series alone cannot describe the exceptional set. The necessary new clause is an inverse minor-arc theorem: persistent failure for one coherent rational section forces the two block measures in (947) into its rational Borel. The rational-Borel proposition then gives exactly the cusp or special alternative.

The quantifier failure is explicit in the cited literature, not merely a caution about proof technique. Zhang's theorem gives a power-saving exceptional set of admissible integers, and his introductory construction shows that after passing to a suitable finite-index subgroup one can make any prescribed finite list of integers absent from a fixed matrix entry while those integers remain locally admissible [70]. Finite index does not change the critical exponent. Thus parabolics, a critical exponent arbitrarily close to one, congruence expansion, and a density-one local–global theorem do not imply representation of one preassigned scalar value.

Bourgain–Kontorovich's beyond-expansion theorem has a different quantifier: it averages congruence errors for a scalar polynomial and raises its level of distribution from the spectral range to $7/24 - \epsilon$ in their normalization [69]. This is the relevant model for removing off-diagonal terms, but it neither selects a fixed scalar value nor retains two archimedean Hodge cocycles. Zhang's minor-arc theorem is likewise an ℓ^2-bound after summing over all output integers. Evaluating that bound at zero is exactly the step its exceptional set forbids. Hence the desired statement is not obtained by renaming either published estimate: its new content is the classification of the fixed exceptional fibre.

Route II.20.32 (Outcome of the thin local–global comparison)**.** The circle-method decomposition is the first tested architecture which uses the correct scalar observable, arbitrary macroscopic elements of the thin group, and a technique known to exceed the Ramanujan/spectral level in related problems. Its density-one output is insufficient here. The precise upgrade is pointed and arithmetic: classify a fixed exceptional value by a coherent K-rational Borel while retaining the two Hodge cocycles. This is materially more specific than asking for "better expansion," and it does not reduce to inequalities for the generators.

Audit: what a generator estimate can never provide

The objection that the preceding work uses too many inequalities for the six Rosen letters has a precise mathematical content. If $a_i(g) = \log \|\sigma_i(g)\|_{\mathrm{op}}$, then

$$\sigma_i(T^n) = \begin{pmatrix} 1 & n\lambda_i \\ 0 & 1 \end{pmatrix} \quad \Longrightarrow \quad a_i(T^n) = \log n + \log|\lambda_i| + o(1).$$

Consequently

$$\frac{a_i(T^n)}{a_0(T^n)} \longrightarrow 1, \qquad \frac{a_1(T^n) + a_2(T^n)}{a_0(T^n)} \longrightarrow 2. \tag{983}$$

No pointwise comparison valid for every word can therefore yield the Hodge-typical outer coefficient $4/5$ in Cartan normalization, or $2/5$ in hyperbolic-radius normalization. Replacing a run T^n by one accelerated symbol changes the notation and not (983).

This proves that the desired Hodge gain is intrinsically statistical. It comes from the Bowen–Margulis distribution of complete geodesic segments and the two Lyapunov exponents of the Hodge

cocycle. The six letters may be used to construct a Markov coding for that distribution; taking a worst case over the letters destroys the gain. Thus a more elaborate generator table cannot be the missing argument.

There are three genuinely group-wide inputs, with sharply different roles.

1. Because the distinguished representation is a lattice, a central shell of radius R has $e^{R+o(R)}$ elements, and a full-entropy subfamily has outer Cartan sum $(2/5 + o(1))R$.
2. Because the diagonal Galois image is Zariski dense in SL_2^3, strong approximation gives the full good finite quotients and the local density (953).
3. Because two arbitrary macroscopic elements multiply according to (947), dispersion can study the additive energy of the two complete block orbits. This third operation, not a letter inequality, is the possible source of an estimate beyond ordinary mixing.

The quantitative standard for the third input is severe.

Proposition II.20.33 (Necessary decay for a black-box mixing proof)**.** *Let Ω_R have cardinality $e^{R+o(R)}$, and let an event have predicted normalized mass*

$$m_R = e^{-(r+\beta+o(1))R}.$$

Suppose the only available conclusion is

$$\left| \frac{\#(\Omega_R \cap E_R)}{\#\Omega_R} - m_R \right| \le e^{-(\kappa - o(1))R}. \tag{984}$$

This estimate forces $\Omega_R \cap E_R \neq \varnothing$ only if $\kappa > r + \beta$. Hence it reaches the closing region (982) only if

$$\kappa > \frac{9}{10} + 2\delta. \tag{985}$$

Proof. The error in (984) must be smaller than its main term. Comparing the two exponents gives $\kappa > r + \beta$; the closing window gives the last assertion. □

This proposition rules out the phrase "use a spectral gap" unless the exponent is actually measured. Cutting the same convolution into more blocks does not improve a fixed decay rate per unit hyperbolic length. Dispersion is relevant precisely because it does not collapse the two macroscopic variables to one convolution norm.

Jialun Li's random-matrix theorem gives a useful, but strictly partial, full-group input [45]. It applies to a probability measure of finite exponential moment whose support is Zariski dense in a connected split semisimple real group. Taking $G = \mathrm{SL}_2(\mathbb{R})^3$ and a finitely supported generating measure on $\rho_{\mathrm{Gal}}(G_7)$, the Zariski-closure calculation verifies those hypotheses. The theorem supplies Fourier decay of the stationary flag measure, a uniform high-frequency gap for the Iwasawa-cocycle transfer operator, and an exponential renewal error. It therefore addresses the archimedean flag and Hodge frequencies of an unconditioned random walk by using the whole diagonal group.

It does not include the finite congruence representation with level e^{rR}, conditioning on one coherent rational boundary cylinder, or the Bowen–Margulis shell measure used here. More decisively, its positive decay exponent depends on the chosen walk and is not asserted to satisfy (985). It can supply a high-frequency minor-arc component of the pointed argument, but it cannot be substituted for the pointed estimate.

The exact inverse problem can be stated without referring to generators. For macroscopic sets $\mathcal{A}_R, \mathcal{B}_R \subset G_7$, define

$$\Phi_P(a, b) = \det(a^{-1}P, bP).$$

The dispersion calculation studies the fibre energy of this map after the congruence, real, and Hodge cutoffs. If Φ_P vanishes identically on a Cartesian subblock, then all the lines $[a^{-1}P]$ and $[bP]$ in that subblock coincide. Two distinct elements on either side give a nonidentity quotient fixing $[P]$. Thus the exact algebraic obstruction is already the desired rational Borel. The missing inverse theorem must show that anomalously large Fourier energy is quantitatively close to this exact configuration; the product formula must then turn that closeness into equality. This is the first formulation in these notes where every nontrivial operation is performed on arbitrary long elements of the thin group.

Route II.20.34 (Revised meaning of "exploit the thin group")**.** The thin group is not expected to supply a stronger inequality for each of its generators. Its useful structure is the simultaneous presence of full central entropy, Zariski-dense three-factor monodromy, large congruence images, and a bilinear multiplication law. The finite local Fourier transform is already completely favorable. The unresolved statement is now an inverse dispersion theorem for the stopped three-factor block measure: a persistent loss of the expected fixed target count must force concentration in the rational Borel of that target. This statement is genuinely about the whole thin orbit, and Proposition II.20.33 explains why neither a generator calculation nor a generic positive spectral gap can replace it.

New route: Fourier inversion on the Minkowski torus

The projective congruence model discards the additive structure of the scalar defect. Retain it for a full annular group shell. For the standard determinant-one lift $\widetilde{G}_7$, set

$$D_g(P) = \det(P, gP) \in \mathcal{O}_K.$$

Fix $0 < \delta < 1/20$. A Hodge-typical shell $\widetilde{\Omega}_R(\delta)$ has $e^{R-o(R)}$ elements and

$$|\sigma_0 D_g(P)| \ll e^{R/2}, \qquad |\sigma_1 D_g(P)| \ll e^{(1/10+\delta)R}, \qquad |\sigma_2 D_g(P)| \ll e^{(3/10+\delta)R}. \tag{986}$$

The lattice box in (986) has volume $e^{(9/10+2\delta)R}$. Thus the one-tenth surplus appears as the expected multiplicity of a *single algebraic integer*, with all finite valuations already assembled by its global norm.

Let $\mathbb{T}_K = K_\infty/\mathfrak{d}_K^{-1}$, and set

$$S_R(P;\xi) = \sum_{g \in \widetilde{\Omega}_R(\delta)} e^{2\pi i \operatorname{Tr}_{K/\mathbb{Q}}(\xi D_g(P))}.$$

Character orthogonality gives the exact global identity

$$\#\{g \in \widetilde{\Omega}_R(\delta) : D_g(P) = 0\} = \int_{\mathbb{T}_K} S_R(P;\xi)\, d\xi. \tag{987}$$

The box around zero whose three widths are the reciprocals of the three bounds in (986) has volume $e^{-(9/10+2\delta)R}$. On a sufficiently small such box every summand has positive real part. Its contribution to (987) is therefore

$$\gg e^{(1/10-2\delta)R-o(R)}. \tag{988}$$

An absolute bound making the complementary Fourier integral smaller would imply the classification, but the local calculation below shows that this is not the correct target. Rational torsion arcs have necessary alternating contributions. Their *signed* sum, not their absolute mass, gives the local density at zero.

There is exact evidence for the required cancellation at every good finite place. If $q = \#k_{\mathfrak{p}}$, strong approximation makes the full group reduction $\mathrm{SL}_2(k_{\mathfrak{p}})$. For $\bar{P} \neq 0$ and a nontrivial additive character ψ,

$$\frac{1}{|\mathrm{SL}_2(k_{\mathfrak{p}})|} \sum_{g \in \mathrm{SL}_2(k_{\mathfrak{p}})} \psi(\det(\bar{P}, g\bar{P})) = -\frac{1}{q^2-1}. \tag{989}$$

Indeed, $g\bar{P}$ is uniform on the q^2-1 nonzero vectors; the corresponding sum over all q^2 vectors is zero, and the deleted zero vector contributes one. In particular

$$\Pr(D_g(\bar{P}) = 0) = \frac{1}{q+1}, \qquad \Pr(D_g(\bar{P}) = a) = \frac{q}{q^2-1} \quad (a \neq 0).$$

This q^{-2} additive Fourier coefficient is lost in the projective Steinberg quotient, where the natural scale is only q^{-1}. It is the first calculation in this exploration where passing from a preferred word alphabet to the complete group quotient gives strictly more cancellation in each nontrivial additive mode. It is not a free net saving: Fourier inversion of the zero fibre sums q such modes and recovers the projective q^{-1} scale. Its advantage is that the modes fit one global Minkowski-torus transform, rather than a collection of unrelated prime incidences.

The prime-power transform is equally explicit. For a character of exact conductor $\mathfrak{p}^j$ on $\mathcal{O}_K/\mathfrak{p}^m$, its average over $D_g(P)$ is $-(q^2-1)^{-1}$ when $j = 1$, and zero when $j \geq 2$. Moreover

$$\Pr(D_g(P) = 0 \bmod \mathfrak{p}^m) = \frac{q^{1-m}}{q+1}.$$

Relative to the uniform scalar density q^{-m}, the local factor is $q/(q+1)$. The truncated singular series is therefore

$$\prod_{\mathrm{N}\,\mathfrak{p} \leq X} \frac{\mathrm{N}\,\mathfrak{p}}{\mathrm{N}\,\mathfrak{p}+1} \asymp_K (\log X)^{-1}.$$

Hence the rational major arcs lose only a logarithm, not any part of the exponential one-tenth surplus. The obstruction is transfer to the thin stopped shell, not a bad local density at zero.

There is a sharper way to organize this calculation which removes the absolute-minor-arc mistake. Choose $9/10 + 2\delta < u < 1$ and form the one global ideal

$$\mathfrak{A}_R = \prod_{\substack{\mathfrak{p} \text{ good} \\ \mathrm{N}\,\mathfrak{p} \leq uR}} \mathfrak{p}. \tag{990}$$

The prime-ideal theorem gives $\log \mathrm{N}\,\mathfrak{A}_R = uR + o(R)$. Therefore the bound (986) implies

$$D_g(P) = 0 \pmod{\mathfrak{A}_R} \quad \Longrightarrow \quad D_g(P) = 0 \tag{991}$$

for large R. Strong approximation makes reduction onto the complete product quotient simultaneous after a fixed finite set of places is deleted. Its exact zero-fibre density is

$$\rho(\mathfrak{A}_R) = \prod_{\mathfrak{p} \mid \mathfrak{A}_R} \frac{1}{\mathrm{N}\,\mathfrak{p}+1} = \frac{1}{\mathrm{N}\,\mathfrak{A}_R} \prod_{\mathfrak{p} \mid \mathfrak{A}_R} \frac{\mathrm{N}\,\mathfrak{p}}{\mathrm{N}\,\mathfrak{p}+1}. \tag{992}$$

Consequently a Hodge shell of size $M_R = e^{R-o(R)}$ has predicted fixed-fibre count

$$M_R \rho(\mathfrak{A}_R) = \frac{e^{(1-u)R-o(R)}}{\log R}, \tag{993}$$

which is exponentially large. A relative local limit for this one fibre therefore proves the classification by (991).

Finite Fourier inversion makes the sign issue exact. At a local factor, the signed sum over the trivial character and the $q-1$ nontrivial characters is

$$1+(q-1)\left(-\frac{1}{q^2-1}\right)=\frac{q}{q+1},$$

while the corresponding absolute sum is

$$1+(q-1)\frac{1}{q^2-1}=1+\frac{1}{q+1}.$$

Their products have sizes $(\log R)^{-1}$ and $\log R$, respectively. Thus the absolute minor-arc target suppresses the very cancellation which produces (992). The correct conditional estimate is the relative square-root bound

$$\begin{aligned}\Big|\#\{g\in\widetilde{\Omega}_R(\delta):D_g(P)=0\pmod{\mathfrak{A}_R}\}-M_R\rho(\mathfrak{A}_R)\Big|\\ \ll_{P,\epsilon}\big(M_R\rho(\mathfrak{A}_R)\big)^{1/2}(\mathrm{N}\,\mathfrak{A}_R)^{\epsilon}.\end{aligned}\tag{994}$$

For $\epsilon<(1-u)/(2u)$, its error is smaller than (993). This is the precise Lindelöf-shaped conditional completion produced by the global route.

There is an exact reason for expecting the inverse alternative to name B_P. If $P=(p,q)^t$, put $\ell_P=(-q,p)$. Then

$$D_g(P)=\ell_P gP,\qquad B_P(K)=\{g:\ell_P g\in K^{\times}\ell_P\}.$$

Hence $\xi\mapsto\xi\ell_P$ embeds the three-dimensional Minkowski torus in the dual of the six-dimensional standard representation, and $S_R(P;\xi)$ is the Fourier coefficient of its contragredient group orbit at the fixed integral character P. The rational normal line of this frequency subtorus has stabilizer exactly B_P. This is why an inverse theorem for persistent Fourier concentration would produce the desired characteristic-zero subgroup rather than an unspecified proper algebraic subgroup.

The archimedean side is also a whole-group representation. The standard K-representation, viewed over $\mathbb{Q}$, has dimension six. Its rational Zariski closure is $\mathrm{Res}_{K/\mathbb{Q}}\,\mathrm{SL}_2$, so it is irreducible over $\mathbb{Q}$, but over $\mathbb{R}$ it splits into the three Galois planes. The later theorem of He–de Saxcé [49] handles exactly this semisimple reducible real representation. In the present specialization, the algebra generated by the connected Zariski closure is

$$E=M_2(\mathbb{R})^3.$$

There are no compact factors. The nonzero frequency P has a nonzero component in each Galois plane, so $PE=(\mathbb{R}^6)^*$ and the exceptional space $W_0=(PE)^{\perp}$ in their theorem is zero. A large Fourier coefficient therefore forces the starting point $J_P(\xi)=\xi\ell_P$ close, in the anisotropic Lyapunov distance, to a rational torsion point of controlled denominator.

After transfer to the Hodge equilibrium shell, the three coordinate widths predicted by that distance are

$$e^{-\vartheta R/2},\qquad e^{-\vartheta R/10},\qquad e^{-3\vartheta R/10},$$

whose product on $J_P(\mathbb{T}_K)$ is $e^{-9\vartheta R/10}$. Thus the existing theorem has the correct restriction-of-scalars geometry. It does not yet give the result needed here because it concerns an independent random walk and a coefficient threshold t with the condition $n\ge C\log(\|P\|/t)$; it does not provide

the signed, relative error (994) for the Bowen–Margulis/Hodge-conditioned shell. Li's theorem [45] supplies a related archimedean Fourier-decay and renewal mechanism. The remaining operation is therefore transfer and summation at the entropy scale, not construction of a restriction-of-scalars inverse theorem.

Route II.20.35 (Outcome and exact remaining lemma)**.** This route gives the exact conditional completion (994). The characteristic-zero inverse theorem now exists and has the correct three Hodge scales. The remaining estimate is the relative fixed-fibre local limit for the Bowen–Margulis group shell, conditioned by the two Hodge exponents. Its finite main term is exactly (992), and its inverse alternative must turn any extra persistent resonance into the rational Borel B_P. The next theoretical task is to extend the semisimple Fourier localization from convolution powers to the Hodge-tilted Markov transfer operator and prove the signed error (994). This is sharper than the earlier absolute minor-arc target and removes the former real-irreducibility gap.

A correction: standard thin-group structure is not the rigidity

It is still too generous to say that the word *thin* explains the proposed conclusion. Four logically different layers have appeared:

1. The presentation $C_2 * C_7$ and the Rosen matrices provide a coordinate system. Estimates for these generators are local bookkeeping.
2. Zariski density, strong approximation, and expansion are genuine theorems about arbitrary long elements. They determine the good finite quotients and their local densities.
3. The modular embeddings, the split Euler defect $2e_0 - e_1 - e_2 = 3/7$, and McMullen's two compression functors are special to the heptagonal Hodge representation. They create the strict $1/10$ global norm margin.
4. The unproved step couples the second and third layers: a coherent rational flag cannot occupy the exceptional set in every quotient and at every real scale unless its rational Borel already meets the group.

Only the last two layers could prove the classification. The first is far too special, as a proof mechanism, and the second is too generic. Strong approximation sees $D_g(P) \equiv 0 \pmod{\mathfrak{a}}$; the theorem asks for $D_g(P) = 0$. The latter is singular because it says that the orbit has entered a subgroup variety, namely B_P. Ordinary thin-orbit counting represents a density-one collection of scalar values and is allowed to discard this one value. Thus the correct proposed use of the actual group is a Hodge-weighted *inverse* theorem for the multiplication map

$$(a, b) \longmapsto \det(a^{-1}P, bP),$$

not another collection of inequalities for S, T, or the six Rosen branches.

This also gives a clean test for any proposed advance. If its hypotheses and conclusion continue to make sense after the two modular embeddings are removed, then it is only generic thin-group technology and cannot by itself yield the classification. A successful argument must use the two outer Hodge coordinates to cross the product-formula threshold and must use multiplication by arbitrary macroscopic group elements to classify concentration at the zero fibre.

A balanced two-block, or birthday, formulation

The entropy can be divided so that neither individual block is asked to solve the full closing problem. Fix

$$\frac{9}{10} + 2\delta < s < 1, \qquad r = \beta = \frac{s}{2}, \tag{995}$$

and factor a radius-R element as $g = ab$, with both blocks of radius $R/2 + O(1)$. Let $\mathcal{A}_R, \mathcal{B}_R$ be Hodge-balanced families of such blocks. Their natural sizes are

$$\#\mathcal{A}_R = \#\mathcal{B}_R = e^{R/2+o(R)}.$$

Choose the two families from compatible one-sided cylinders in the positive nonbacktracking coding. Every concatenation ab then has radius $R + O(1)$ and is projectively nonidentity. This condition is essential. If one instead took inverse-symmetric balls, the pairs $b = a^{-1}$ would contribute $e^{R/2+o(R)}$ copies of the identity, far more than (996), without proving anything.

Partition the possible lines by two labels: reduction in $X_{\mathfrak{p}}$, where $\log \#X_{\mathfrak{p}} = rR + o(R)$, and a distinguished-real interval of diameter $e^{-\beta R}$. The finite marginal has $e^{rR+o(R)}$ classes and the real marginal has $e^{\beta R+o(R)}$ intervals. By (995), each number is strictly smaller than the entropy $e^{R/2}$ in a single block. Thus neither marginal is at an entropy-size level for the half block.

The joint partition has

$$e^{(r+\beta+o(1))R} = e^{(s+o(1))R}$$

cells. It is larger than either block, so one should not ask one block to fill it. The pair space has $e^{R+o(R)}$ elements, however, and the birthday prediction is

$$\frac{\#\mathcal{A}_R \#\mathcal{B}_R}{e^{(s+o(1))R}} = e^{(1-s+o(1))R} \tag{996}$$

collisions.

The half-entropy threshold has an exact spectral interpretation. It is useful to record it because it shows both what a Ramanujan conjecture would prove and what it would leave untouched.

Proposition II.20.36 (Ramanujan-sized marginal flattening)**.** *Let X_R be a finite state space with $\#X_R = e^{rR+o(R)}$, and let P_R be the normalized transition operator produced by a block of central radius $R/2$. Assume the entropy-one Ramanujan estimate*

$$\|\mathsf{P}_R|_{L^2_0(X_R)}\| \le e^{-R/4+o(R)}. \tag{997}$$

If δ_x is a point mass and u_R is uniform measure, then, in density norm relative to u_R,

$$\left\| \frac{\mathsf{P}_R \delta_x}{u_R} - 1 \right\|_2 \le e^{-(1/4-r/2)R+o(R)}. \tag{998}$$

In particular, one half-block is L^2-equidistributed whenever $r < 1/2$.

Proof. The density of a point mass relative to uniform measure has mean-zero part of norm $e^{rR/2+o(R)}$. Apply (997). □

Applied to the projective congruence orbit, this proposition places the finite marginal of (995) exactly in the Ramanujan range. An analogous optimal L^2 estimate for the Bowen–Series boundary operator, after coarsening the real line into $e^{\beta R}$ intervals, has the same threshold $\beta < 1/2$. These are separate conditional marginal statements; neither asserts joint equidistribution.

The distinction is not technical. Let both marginal sets have M labels. Give the first block equal mass on the graph

$$\{(j,j) : j \in \mathbb{Z}/M\mathbb{Z}\}$$

and the second equal mass on the shifted graph

$$\{(j,j+1) : j \in \mathbb{Z}/M\mathbb{Z}\}.$$

Both finite marginals and both real marginals are exactly uniform, yet the two joint supports are disjoint. Repeating every graph point with the same multiplicity makes each source have the required $e^{R/2+o(R)}$ mass. Thus even conjecturally optimal Ramanujan bounds for the two marginals do not prove one collision.

More explicitly, for a joint cell C, put

$$A_C = \#\{a \in \mathcal{A}_R : [a^{-1}P] \in C\},$$
$$B_C = \#\{b \in \mathcal{B}_R : [bP] \in C\}.$$

The exact thin-group statistic is

$$\mathcal{C}_R(P) = \sum_C A_C B_C. \tag{999}$$

The required lower bound is

$$\mathcal{C}_R(P) \geq e^{(1-s-o(1))R} \tag{1000}$$

unless $[P]$ already has a nontrivial stabilizer.

Indeed, one term in (999) places $[a^{-1}P]$ and $[bP]$ in the same projective congruence class and within real chordal distance $O(e^{-\beta R})$. Therefore

$$\det(a^{-1}P, bP) = D_{ab}(P) \in \mathfrak{p},$$

while its distinguished absolute value gains $e^{-\beta R}$. The outer determinant estimate is the sum of the two blockwise Hodge costs. Since $r + \beta = s > 9/10 + 2\delta$, the product-formula closing proof in Theorem II.20.31 gives $D_{ab}(P) = 0$.

This is more than a reparametrization. For a Ramanujan nonbacktracking walk, pointwise positivity from one radius-R ball has the familiar factor-two barrier. The collision (999) pairs two half-balls; their two L^2 errors multiply, which is the birthday mechanism that reaches a marginal quotient of size below $e^{R/2}$. The symmetric choice $r = \beta < 1/2$ places both separate marginals on that side of the threshold.

There is one remaining correlation which separate marginal estimates do not see. Two sparse subsets of a product may each have nearly uniform finite and real marginals but occupy disjoint graphs. Here that would mean that the congruence label and the real interval of a block remain locked together through all scales. The inverse theorem required for (1000) is exactly that such a coherent graph correlation forces the block orbit into the rational Borel of $[P]$. This is a substantially narrower target than a level-one pointwise local limit: the separate half-entropy estimates are Ramanujan-sized marginal problems, and only their arithmetic coupling must cross the full closing threshold.

Route II.20.37 (Outcome of the birthday split)**.** The macroscopic product $g = ab$ reveals a use of thinness which cannot be expressed by generator inequalities. Each half-orbit has enough entropy for Ramanujan-sized separate flattening modulo $\mathfrak{p}$ and at the real resolution required by (995); the two halves together have enough entropy to collide in the joint space. The proof obligation is now a two-source inverse theorem: absence of the predicted collisions forces a persistent congruence–real graph, and that graph must be the rational Borel. The Hodge variables enter as blockwise large-deviation cutoffs and add correctly under multiplication.

A short bridge which tests graph rigidity

The standard thin-orbit dispersion ensembles use three pieces rather than two: two long bilinear blocks and a shorter middle block which moves a possible structured obstruction. That device fits the present incidence exactly. Choose

$$0 < \tau < 1 - s \tag{1001}$$

and compatible positive block families with radii

$$R(a) = R(b) = \frac{1-\tau}{2}R + O(1), \qquad R(c) = \tau R + O(1).$$

The exact identity becomes

$$D_{acb}(P) = \det(a^{-1}P, cbP). \tag{1002}$$

The two outside block families have size $e^{(1-\tau)R/2+o(R)}$. Since

$$r = \beta = \frac{s}{2} < \frac{1-\tau}{2},$$

each outside block still has more entropy than either separate marginal. For one fixed bridge the birthday prediction is

$$e^{(1-\tau-s+o(1))R};$$

averaging over the $e^{\tau R+o(R)}$ bridges restores the total $e^{(1-s+o(1))R}$.

The bridge is not inserted to add entropy. It tests whether a sparse joint graph is genuine structure. If the first and second block histograms avoid each other after multiplication by every $c \in \mathcal{C}_R$, then the alleged congruence–real relation survives exponentially many changes of frame. In a dispersion proof, repeated Cauchy–Schwarz and a flattening theorem should convert this persistence into approximate invariance under the middle block measure. A sum–product inverse theorem then places that measure near a proper algebraic stabilizer.

At the finite places the desired effect is not heuristic. Apply (954) with $u = a^{-1}P$ and $v = bP$. A complete quotient bridge has the same favorable scalar Fourier transform for every pair (a, b). The thin group problem is to transfer this endpoint-uniform cancellation to the short family $\mathcal{C}_R$, using the two long outer families as the bilinear amplifiers. This is precisely where multiplication by many long elements of G_7, rather than a calculation with its generators, must enter.

Proposition II.20.20 also states the limitation exactly. The short bridge contributes the congruence spectral norm and nothing better. Hence the missing estimate cannot be proved by declaring the middle block equidistributed. It must be a three-linear estimate in which the long a- and b-families lower the effective ℓ^2-norm of the endpoint data before the bridge operator is applied.

The size of this obstruction can be computed without a model or a numerical experiment. Let $\mathscr{Z}_R$ be the joint partition by a projective residue class and a real interval of length $e^{-\beta R}$. Then

$$M_R := \#\mathscr{Z}_R = e^{(s+o(1))R}, \qquad s = r + \beta,$$

whereas one outside block has

$$N_R = e^{(L+o(1))R}, \qquad L = \frac{1-\tau}{2}.$$

For its endpoint histogram n_Z, let $p_Z = n_Z/N_R$, let $u_Z = 1/M_R$, and put $f = p/u$. With the $L^2(u)$-norm one has the exact support bound

$$\|f - 1\|_2^2 = \frac{M_R}{N_R^2}\sum_Z n_Z^2 - 1 \geq \frac{M_R}{N_R} - 1 = e^{(s-L+o(1))R}. \tag{1003}$$

Equality, to exponential order, is the most favorable case in which the block points occupy distinct joint cells. Uniformity of either marginal does not improve this bound.

Suppose one now uses only a bridge operator estimate

$$\|\mathsf{T}_{\mathcal{C}_R}|_{L^2_0}\| \le e^{-(\kappa\tau-o(1))R}.$$

For two endpoint densities the standard Cauchy–Schwarz argument bounds the normalized correlation error by

$$e^{-(\kappa\tau-o(1))R}\|f_A-1\|_2\|f_B-1\|_2.$$

Even if both histograms attain the minimum in (1003), this proof forces the error below the main term only if

$$\kappa\tau > s - \frac{1-\tau}{2}. \tag{1004}$$

Here $s > 9/10$ and $\tau < 1-s$. The right side is greater than $2/5$, while the bridge has length less than $R/10$; thus the required contraction per unit bridge length is greater than four. A Ramanujan-scale bridge has $\kappa = 1/2$. Even the formal rate $\kappa = 1$ is nowhere near (1004).

The same loss appears before introducing an operator. The collision energy of one histogram contains the diagonal contribution N_R, whereas the predicted number of distinct birthday collisions is

$$\frac{N_R^2}{M_R} = e^{(1-\tau-s+o(1))R}.$$

Their ratio is

$$\frac{N_R}{N_R^2/M_R} = \frac{M_R}{N_R} = e^{(s-(1-\tau)/2+o(1))R}. \tag{1005}$$

Thus an ordinary second moment is dominated by identical endpoint pairs, which yield the identity element and prove nothing. The required dispersion estimate must remove this diagonal before Cauchy–Schwarz and obtain cancellation from the two distinct long block variables. This is an off-diagonal trilinear problem, not a stronger spectral gap for the short bridge.

This calculation identifies a precise false start. Positivity of an averaged self-correlation, for instance by replacing the bridge measure with $\nu^* * \nu$, gives only the diagonal term in (1005). It cannot produce a projectively nonidentity stabilizer. Any successful inverse theorem must either prove the off-diagonal main term directly or show that failure of that term creates an exact zero incidence between two distinct endpoint blocks.

There is a second exact fact, this time at the distinguished real place. It is elementary, but it prevents a misleading description of the inverse alternative.

Proposition II.20.38 (There is no invariant congruence–real graph)**.** *Let X be a finite transitive G_7-set, for example a projective congruence orbit, and give $\mathbf{P}^1(\mathbb{R})$ its usual topology. If*

$$Z \subset X \times \mathbf{P}^1(\mathbb{R})$$

is nonempty, closed, and G_7-invariant for the diagonal action, then $Z = X \times \mathbf{P}^1(\mathbb{R})$. In particular there is no equivariant map $X \to \mathbf{P}^1(\mathbb{R})$.

Proof. The projection of Z to X is nonempty and invariant, hence is all of X. Choose $x_0 \in X$ and put $H = \mathrm{Stab}_{G_7}(x_0)$. Transitivity makes H finite index in G_7. Hence H is again a cofinite, non-elementary Fuchsian group and its limit set is the whole circle $\mathbf{P}^1(\mathbb{R})$. The action of such a group on its limit set is minimal. The fibre Z_{x_0} is a nonempty closed H-invariant subset, so it is the whole circle. Translating this fibre by G_7 proves the assertion for every point of X. □

Thus it is not literally correct to say that a persistent graph becomes the rational Borel. If persistence reached exact invariance, the graph would disappear into the full product. The required quantitative inverse theorem must act earlier: either dispersion produces the desired collision, or concentration produces a near-return to the zero-incidence variety which the global product formula makes exact. The rational Borel then appears as the stabilizer of that exact incidence, not as a proper invariant graph in the finite–real product.

The joint operator is a congruence-cover spectrum

The tensor product in (979) has a standard representation- theoretic meaning. It is not an arbitrary tensor product assembled for this problem.

The ambient Hilbert modular geometry supplies another exact use of the thin group. Let

$$\Lambda = \mathrm{PSL}_2(\mathcal{O}_K) \subset \mathrm{PSL}_2(\mathbb{R})^3$$

and let

$$\widetilde{C}_7 = \{(z, F_1(z), F_2(z)) : z \in \mathbb{H}\} \subset \mathbb{H}^3$$

be the lifted modular graph.

Proposition II.20.39 (The thin group is the full ambient stabilizer)**.** *Up to the harmless central sign,*

$$\mathrm{Stab}_\Lambda(\widetilde{C}_7) = G_7. \tag{1006}$$

Proof. Equivariance of the modular embeddings gives the inclusion from right to left. Put $\Delta = \mathrm{Stab}_\Lambda(\widetilde{C}_7)$. Projection to the first real factor is faithful on Δ: if that factor is the identity, then an element of either outer factor fixes the open image of the nonconstant map F_i, and hence is the identity. The projected group is discrete. Indeed, convergence to the identity in the first factor, together with

$$F_i(\delta_0 z) = \delta_i F_i(z),$$

forces convergence to the identity in both outer factors; discreteness of Λ then applies.

Thus Δ is a discrete Fuchsian group containing the cofinite group G_7, so $[\Delta : G_7] = n < \infty$. The orbifold Euler characteristic of the $(2, 7, \infty)$ orbifold is

$$\chi(G_7\backslash\mathbb{H}) = -\frac{5}{14}.$$

If $n > 1$, the cusped orientable orbifold $\Delta\backslash\mathbb{H}$ would have

$$-\chi(\Delta\backslash\mathbb{H}) = \frac{5}{14n} \le \frac{5}{28} < \frac{1}{4}.$$

The only orientable finite-area hyperbolic orbifold with a cusp and $-\chi < 1/4$ has signature $(2, 3, \infty)$, for which $-\chi = 1/6$. This is not $5/(14n)$ for an integer n. Hence $n = 1$. □

The pointed congruence covers also come directly from the ambient threefold. This uses strong approximation, not a choice of symbolic alphabet.

Proposition II.20.40 (The pointed cover is an ambient connected pullback)**.** *Let $\mathfrak{a}$ be a good ideal for which the reduction of G_7 is the full reduction of Λ. Let $\Lambda_{B,\mathfrak{a}}$ be the inverse image in Λ of the Borel stabilizing a chosen point of $\mathbf{P}^1(\mathcal{O}_K/\mathfrak{a})$, and put*

$$G_{7,B,\mathfrak{a}} = G_7 \cap \Lambda_{B,\mathfrak{a}}.$$

Then the fibre product of the triangle curve with the Borel-level Hilbert modular cover is connected and is naturally

$$G_{7,B,\mathfrak{a}}\backslash\mathbb{H}.$$

Its degree over $G_7\backslash\mathbb{H}$ *is*

$$[G_7 : G_{7,B,\mathfrak{a}}] = [\Lambda : \Lambda_{B,\mathfrak{a}}] = \#\mathbf{P}^1(\mathcal{O}_K/\mathfrak{a}). \tag{1007}$$

Proof. The connected components of the fibre product are indexed by

$$G_7\backslash\Lambda/\Lambda_{B,\mathfrak{a}}.$$

Strong approximation says that G_7 acts transitively on $\Lambda/\Lambda_{B,\mathfrak{a}} \simeq \mathbf{P}^1(\mathcal{O}_K/\mathfrak{a})$, so this double-coset set has one element. The stabilizer of that component is the displayed intersection, and orbit–stabilizer gives (1007). □

The two outer Hodge bundles are also the exact normal directions to the modular graph.

Proposition II.20.41 (Negative logarithmic normal bundle)**.** *On the compactified orbifold curve, the logarithmic normal bundle of* C_7 *in the Hilbert modular threefold splits as*

$$N^{\log}_{C_7/X_K} \simeq \mathcal{L}_1^{-2} \oplus \mathcal{L}_2^{-2}. \tag{1008}$$

Its two degrees are

$$\deg \mathcal{L}_1^{-2} = -\frac{1}{14}, \qquad \deg \mathcal{L}_2^{-2} = -\frac{3}{14}. \tag{1009}$$

Consequently $H^0(C_7, N^{\log}_{C_7/X_K}) = 0$*: the modular graph has no nontrivial infinitesimal logarithmic deformation in* X_K*.*

Proof. The logarithmic tangent bundle of the threefold restricted to the graph is

$$\mathcal{L}_0^{-2} \oplus \mathcal{L}_1^{-2} \oplus \mathcal{L}_2^{-2}.$$

The first coordinate of the graph map is the identity. Quotienting its tangent line therefore identifies the two normal coordinates with the last two summands: explicitly, send (v_0, v_1, v_2) to $(v_1 - dF_1(v_0), v_2 - dF_2(v_0))$. This proves (1008). The degree values follow from

$$\deg \mathcal{L}_1 = \frac{1}{28}, \qquad \deg \mathcal{L}_2 = \frac{3}{28}.$$

A line bundle of negative orbifold degree has no nonzero holomorphic section, so neither summand contributes to H^0. □

This is qualitative rigidity, not yet the required separation estimate. A relative trace proof needs an arithmetic effective form: if a distinct level translate shadows the modular graph, its normal displacement must define a nonzero approximate section whose finite congruence vanishing and archimedean smallness are incompatible with (1009). Constructing that global arithmetic normal displacement, rather than estimating prime divisors one at a time, is the precise Arakelov task.

The numerical window 1/10 also has an exact degree interpretation.

Proposition II.20.42 (The one-tenth surplus is a Hodge degree)**.** *Put*

$$\mathcal{S} = \mathcal{L}_0 \otimes \mathcal{L}_1^{-1} \otimes \mathcal{L}_2^{-1}.$$

Then

$$\deg \mathcal{S} = \frac{5-1-3}{28} = \frac{1}{28} > 0, \tag{1010}$$

and

$$\frac{\deg \mathcal{L}_0 + \deg \mathcal{L}_1 + \deg \mathcal{L}_2}{2 \deg \mathcal{L}_0} = \frac{9}{10}, \qquad \frac{\deg \mathcal{S}}{2 \deg \mathcal{L}_0} = \frac{1}{10}. \tag{1011}$$

Proof. Insert the three degrees $5/28, 1/28, 3/28$. The first ratio is the sum of the three column-growth exponents in hyperbolic-radius normalization; the second is its complement to the entropy-one orbital growth. □

Thus the closing surplus is the degree of one fixed Hodge line, not a numerical accident of the local estimates. What is still absent is a canonical morphism from the finite incidence divisor into $\mathcal{S}$ at the stopped return. Constructing such a morphism would turn the product-formula exponent calculation into a global Arakelov slope inequality; choosing a new local frame for every return does not.

Proposition II.20.43 (Tensor induction for the pointed congruence cover)**.** *Let $\Gamma = G_7$, let $X = \Gamma/H$ be a finite transitive congruence orbit, and let π_X be its permutation representation. For every unitary representation ρ of Γ,*

$$\pi_X \otimes \rho \simeq \mathrm{Ind}_H^\Gamma(\rho|_H). \tag{1012}$$

Proof. This is the tensor identity for induced representations:

$$(\mathrm{Ind}_H^\Gamma 1) \otimes \rho \simeq \mathrm{Ind}_H^\Gamma(\mathrm{Res}_H^\Gamma \rho).$$

Since $\mathrm{Ind}_H^\Gamma 1 = \ell^2(\Gamma/H)$, the left side is $\pi_X \otimes \rho$. □

For $X = X_{\mathfrak{p}}$, the subgroup H is the inverse image of a Borel in the finite quotient. Thus the joint finite–real spectrum is the spectrum of the corresponding finite cover of the triangle orbifold, with the boundary representation as coefficients. This also explains Proposition II.20.38: an invariant tensor would be an invariant object on that cover.

There is a useful ambient interpretation. The map

$$C_7 = G_7 \backslash \mathbb{H} \longrightarrow \mathrm{SL}_2(\mathcal{O}_K) \backslash \mathbb{H}^3, \qquad z \longmapsto (z, F_1(z), F_2(z)),$$

places the triangle curve in the Hilbert modular threefold. Proposition II.20.40 identifies the pointed congruence cover as the connected pullback of the ambient Borel-level cover. The two outer Hodge cocycles are the metrics on the two automorphic line bundles restricted to this curve.

This does not by itself identify the intrinsic spectrum of the curve cover with the ambient Hilbert modular spectrum. Such an identification requires a relative trace comparison. What Proposition II.20.39 does prove is that, on the geometric side of such a comparison, the stabilizer contribution is exactly the G_7-orbital kernel. Every other ambient translate moves the lifted curve.

The word "relative" needs one qualification here. The modular graph is not a homogeneous orbit of a positive-dimensional subgroup of $\mathrm{PSL}_2(\mathbb{R})^3$. Indeed, a complex one-dimensional symmetric subgroup orbit which projects nontrivially to the first factor is, after conjugacy, a diagonal copy of

$\mathbb{H}$; each nonconstant projection to another factor is then a Möbius isometry. Its three automorphic tangent degrees would consequently be equal. On the modular graph those degrees are instead

$$\deg \mathcal{L}_0 : \deg \mathcal{L}_1 : \deg \mathcal{L}_2 = 5 : 1 : 3.$$

Equivalently, the two period maps are strict Schwarz–Pick contractions, not Möbius maps. Thus a symmetric-subgroup relative trace formula cannot simply be quoted. The appropriate object is a restriction pretrace identity for a nonhomogeneous algebraic curve.

That identity can nevertheless be written exactly. Put

$$j(z) = (z, F_1(z), F_2(z)),$$

let $\mathcal{F}$ be a fundamental domain for G_7, and let k be a smooth compactly supported point-pair kernel on $\mathbb{H}^3$. For compactly supported smooth functions a, b on the triangle curve, set

$$\mathcal{I}_k(a,b) = \int_{\mathcal{F}}\int_{\mathcal{F}} a(z)\overline{b(w)} \sum_{\lambda\in\Lambda} k(j(z), \lambda j(w))\, d\mu(z) d\mu(w). \tag{1013}$$

Absolute convergence permits the exact splitting

$$\mathcal{I}_k(a,b) = \mathcal{I}_k^{\mathrm{thin}}(a,b) + \mathcal{I}_k^{\mathrm{tr}}(a,b), \tag{1014}$$

$$\mathcal{I}_k^{\mathrm{thin}}(a,b) = \int_{\mathcal{F}}\int_{\mathcal{F}} a(z)\overline{b(w)} \sum_{\gamma\in G_7} k(j(z), j(\gamma w))\, d\mu(z) d\mu(w), \tag{1015}$$

$$\mathcal{I}_k^{\mathrm{tr}}(a,b) = \int_{\mathcal{F}}\int_{\mathcal{F}} a(z)\overline{b(w)} \sum_{\lambda\in\Lambda\setminus G_7} k(j(z), \lambda j(w))\, d\mu(z) d\mu(w). \tag{1016}$$

The first equality uses more than the inclusion $G_7 \subset \Lambda$: Proposition II.20.39 says that the first sum contains *every* translate preserving the modular graph. Hence (1015) is the pretrace kernel of the complete thin group, while every term in (1016) is genuinely transverse.

There is a level version with exactly the finite variable required in the classification. Replace Λ by the inverse image of a residue Borel at a good split prime $\mathfrak{p}$, and take the matrix-valued kernel between the two finite vectors attached to ∞ and $[P \bmod \mathfrak{p}]$. Proposition II.20.40 says that the restricted curve is the connected pointed cover, not a union of smaller components. On the mean-zero fibre the finite test is exactly

$$e_{[P \bmod \mathfrak{p}]} \otimes e_\infty^* \quad \text{on} \quad \mathrm{St}_{\mathfrak{p}}; \tag{1017}$$

there are no further nonconstant finite representations. After choosing the radial kernel and Mellin weights imposing the two Hodge windows, the thin summand is precisely the pointed count in (964).

The local conductor saving in this comparison is already a theorem on the open Bruhat part. Whenever a term changes the two residue cusp frames, its nonconstant finite matrix passes through $B_1 w B_2$, and Proposition II.20.27 bounds every nontrivial cusp-Fourier coefficient by $2q^{-1/2}$. The qualification is essential: an ambient translate can be transverse in characteristic zero while lying in the same residue Borel modulo $\mathfrak{p}$. Such a term remains in the locally unflattened line $\mathbb{C}\phi_0$. Concentration in that line at one prime proves only a congruence Borel condition, not membership in $B_P(K)$.

This separates the remaining theorem into two precise statements. The open-Bruhat transverse sum must retain the proved local $q^{-1/2}$ after the two Hodge weights are imposed. The closed-Bruhat transverse sum must satisfy an inverse theorem: coherent persistence in the Borel-fixed lines

across growing levels and archimedean stopping scales forces the characteristic-zero intersection $G_7 \cap B_P(K)$. The second statement is exactly where the product formula and the negative normal degrees must enter. Generator estimates have no role in either part of the identity.

Spectrally, $\mathcal{I}_k$ is an ambient automorphic restriction period. Solving for $\mathcal{I}_k^{\mathrm{thin}}$ in (1014) shows why an ambient Lindelöf bound by itself is not enough: one needs square-root bounds both for the spectral restriction packet and for the transverse sum. The latter is the geometric Arakelov problem suggested by the negative normal bundle. In particular, the exact full-group route is

$$\boxed{\text{ambient restriction period} - \text{transverse curve translates} = \text{complete thin-group kernel.}} \tag{1018}$$

It uses the full stabilizer, the full congruence monodromy, and both Hodge-normal directions simultaneously.

This identifies the closest standard-shaped conjectural package. Ramanujan temperedness bounds each individual matrix coefficient at the square-root scale of the whole ball and is insufficient by (940). Lindelöf for the *packet of relative periods* predicts square-root cancellation after the spectral coefficients are summed. That is square-root in the local expected count, which is the stronger scale actually needed.

Conjecture II.20.44 (Relative Lindelöf for the heptagonal congruence tower)**.** *Use the three-block parameters (995) and (1001). Suppose $[P]$ has no projectively nonidentity stabilizer. Let $\mathcal{N}_R(P)$ be the smoothed three-block count with projective congruence equality, real resolution $e^{-\beta R}$, and total outer Hodge cost at most $(2/5+\delta)R$. Its zero-frequency term is*

$$\mathcal{M}_R(P) = e^{(1-s+o(1))R}.$$

The spectral expansion on the pointed congruence covers satisfies

$$\mathcal{N}_R(P) = \mathcal{M}_R(P) + \mathcal{E}_R(P), \qquad |\mathcal{E}_R(P)| \le \mathcal{M}_R(P)^{1/2} e^{o(R)}, \tag{1019}$$

uniformly after averaging over the prescribed good split primes and over the bounded Mellin parameters used to impose the Hodge window; the prime average here is normalized, so it does not change the displayed main scale.

Proposition II.20.45 (Relative Lindelöf completes the classification)**.** *Conjecture II.20.44 implies the full heptagonal cusp/special fixed-point classification.*

Proof. Since $s < 1$, equation (1019) gives

$$|\mathcal{E}_R(P)| = o(\mathcal{M}_R(P)).$$

Hence $\mathcal{N}_R(P) > 0$ for large R. Choose one counted compatible triple. Congruence equality makes its defect divisible by the chosen prime, real closeness supplies the factor $e^{-\beta R}$, and the Hodge window bounds the two outer absolute values. Because $r + \beta = s > 9/10 + 2\delta$, the product-formula calculation in Theorem II.20.31 forces the defect to vanish. Compatibility makes the resulting group element projectively nonidentity, contradicting the assumption on $[P]$. The rational-Borel proposition gives the two stated alternatives. □

Route II.20.46 (What is standard in this conditional route)**.** The tensor-induction identity, the Hilbert modular realization, and the ambient-stabilizer proposition are theorems. Temperedness of the new congruence spectrum is a standard Ramanujan-type conjecture but is not sufficient.

The additional square-root estimate (1019) is the relative-period analogue of generalized Lindelöf. It has the correct standard square-root scale, but it is not yet a formal instance of automorphic Lindelöf: one must first construct the relative trace comparison and control the non-stabilizing translates. The genuinely problem-specific work is precisely that comparison and uniformity in the two nonunitary Hodge weights.

Question II.20.47 (Ambient relative trace comparison)**.** *Construct a level-varying kernel on the Hilbert modular threefold whose $C_7 \times C_7$-period has the following properties:*

1. *its stabilizer term, identified by Proposition II.20.39, is the smoothed three-block count $\mathcal{N}_R(P)$;*
2. *non-stabilizing translates contribute $O(\mathcal{M}_R^{1/2} e^{o(R)})$, or produce an incidence which the product formula exactifies; and*
3. *the spectral side is bounded by ambient Ramanujan and generalized Lindelöf period estimates.*

The curve has complex codimension two. Its two normal Hodge directions are exactly the two outer line bundles whose degrees give the $1/5$ and $3/5$ exponents. Thus an Arakelov intersection estimate for the non-stabilizing translates would control all finite places at once and is the natural global replacement for prime-by-prime dispersion.

There is an exact algebraic endpoint to this argument. Let $G = \mathrm{SL}_2$ over K, let B_P be the Borel fixing $[P]$, and let $m : G \times G \to G$ be multiplication. Then

$$\{(a, h) : D_{ah}(P) = 0\} = m^{-1}(B_P). \tag{1020}$$

The right side is irreducible: the map

$$G \times B_P \longrightarrow m^{-1}(B_P), \qquad (a, u) \longmapsto (a, a^{-1}u)$$

is an algebraic isomorphism. Thus the zero-incidence variety is not an unclassified exceptional component; it is exactly the rational Borel relation required by the theorem.

This motivates the following sharper replacement for the broad pointed local–global conjecture.

Conjecture II.20.48 (Three-block adelic bridge inverse theorem)**.** *Use the parameters in (995) and (1001), and impose blockwise Hodge cutoffs whose sums are at most $(2/5 + \delta)R$. For every primitive $P \in \mathcal{O}_K^2$, either*

1. *the number of compatible triples (a, c, b) for which $[a^{-1}P]$ and $[cbP]$ have the same reduction modulo $\mathfrak{p}$ and real distance $O(e^{-\beta R})$ is $e^{(1-s-o(1))R}$, averaged over the prescribed split primes; or*
2. *the middle-block dispersion measure concentrates, at every relevant local scale, on the relation $m^{-1}(B_P)$, and the inverse argument has one of the following exact outputs: either it supplies a compatible triple whose incidence defect is divisible by the chosen finite ideal and has the three archimedean bounds required by the closing inequality, or it supplies a projectively nonidentity element of $G_7 \cap B_P$ directly.*

Proposition II.20.49 (The bridge inverse theorem completes the classification)**.** *Conjecture II.20.48 implies the full heptagonal cusp/special fixed-point classification.*

Proof. In the first alternative choose one counted triple. Equation (1002), congruence divisibility, real distance, and the blockwise Hodge estimates give the same norm inequality as in Theorem II.20.31. Since $s > 9/10 + 2\delta$, the defect vanishes. The compatible positive language makes acb projectively nonidentity, hence it fixes $[P]$.

In the second alternative, apply the product formula to the supplied triple in the first subcase; it gives $D_{acb}(P) = 0$, and compatibility makes acb projectively nonidentity. In the other subcase the required element of $G_7 \cap B_P$ is already supplied. The rational-Borel proposition then gives a cusp or a special hyperbolic fixed point. Thus both alternatives give the desired classification. Notice that mere metric concentration near $m^{-1}(B_P)$ would not justify this step; the exact output in the statement is essential. □

Route II.20.50 (Outcome of adding the bridge)**.** The two long blocks account for the near-full entropy, while the short bridge tests the only obstruction left by separate marginal mixing. The statement to prove is now an adelic flattening/inverse theorem, not a word inequality. Proposition II.20.38 rules out an exact invariant graph; the quantitative theorem must turn failure of dispersion into an exact point of the rational Borel incidence relation (1020). The Hodge cocycles supply the archimedean exactification budget, and the finite quotient calculation supplies the local density. This is the closest current formulation to the three-piece ensembles used successfully in thin-orbit dispersion.

II.20.12 A second full-orbit target: rigidity of a large incidence fibre

There is a purely archimedean–additive use of the same 1/10 gap. It avoids congruence quotients and makes McMullen's joint compression act on the whole group at once.

Use the canonical scale of a primitive cusp column supplied by its positive primitive parabolic generator, and consider columns $v = he$ in a central ball of radius R whose two outer sizes satisfy

$$H_1(v) \le e^{(1/10+\epsilon)R}, \qquad H_2(v) \le e^{(3/10+\epsilon)R}, \qquad H_0(v) \le e^{R/2}.$$

The exponents are the two Hodge-typical values in displacement normalization. Standard orbital counting with large deviations predicts, and a suitable vector-valued transfer theorem should prove, that this set contains $e^{R-o(R)}$ distinct cusp columns.

For a fixed primitive P, put $r(v) = \det(P, v)$. Each embedding of this algebraic integer lies in the corresponding coordinate box:

$$|\sigma_0 r(v)| \ll_P e^{R/2}, \qquad |\sigma_1 r(v)| \ll_P e^{(1/10+\epsilon)R}, \qquad |\sigma_2 r(v)| \ll_P e^{(3/10+\epsilon)R}.$$

Geometry of numbers shows that this box contains

$$O_{P,\epsilon}\Big(e^{(9/10+2\epsilon)R}\Big)$$

algebraic integers. Therefore one exact value r is assumed by at least

$$e^{(1/10-3\epsilon-o(1))R}$$

canonical cusp columns.

The algebra of such a fibre has a useful exact form. If $\det(P, v) = \det(P, w) = r$, primitivity of P gives a unique $t \in \mathcal{O}_K$ such that

$$w - v = tP.$$

Moreover,

$$\det(v, w) = t \det(v, P) = -tr. \tag{1021}$$

If $v = ge$ and $w = g'e$, then

$$\det(v, w) = \det(e, g^{-1}g'e),$$

so the right side of (1021) is a matrix entry of an element of G_7. Thus, after writing the fibre as

$$v_t = v_0 + tP,$$

every pairwise difference $r(t-s)$ belongs to the full entry set $\mathcal{E}_7$. McMullen's joint compression then places the two Galois ratios of every difference in the compact countable joint ratio set, with convergence to $(0,0)$ along unbounded thick double-coset distance.

Proposition II.20.51 (The large fibre is at the random scale, not a rigidity scale)**.** *The preceding pigeonhole produces a fibre of the size predicted by uniform distribution. More precisely, the Hodge-balanced ball contains* $e^{R-o(R)}$ *cusp columns and its incidence box contains* $e^{(9/10+o(1))R}$ *possible algebraic integers. Its mean fibre size is therefore*

$$\frac{e^{R-o(R)}}{e^{(9/10+o(1))R}} = e^{(1/10-o(1))R},$$

exactly the lower bound obtained above. Consequently cardinality alone cannot force a stabilizer: a randomly distributed orbit would have fibres of this size. Any valid rigidity theorem must use additional structure of the full parabolic conjugacy class.

There is an exact trace identity, but it does not repair the problem. If $U_v = I + \lambda N_v$ is the positive primitive parabolic with column v, then

$$\operatorname{tr}(U_v U_w^{-1}) - 2 = \lambda^2 \det(v,w)^2. \tag{1022}$$

Hence, in one incidence fibre,

$$\operatorname{tr}(U_{v_t} U_{v_s}^{-1}) - 2 = \lambda^2 r^2 (t-s)^2.$$

This looks like a strong restriction on the difference set until one notices that (1022) holds for *every* pair of cusps. It merely rewrites the fact that the points were chosen from the cusp orbit. No independent trace constraint has been gained.

Route II.20.52 (Outcome of the incidence-fibre reformulation)**.** The geometry-of-numbers calculation is correct, but it consumes the whole 1/10 saving merely to reach the average occupancy of a *moving* incidence value r_R. The pairwise entry and trace formulas are automatic identities on the cusp orbit. Thus this route does not reduce the classification to a smaller rigidity theorem.

A viable variant would have to keep the incidence value fixed as $R \to \infty$, prove a fibre size exponentially larger than $e^{R/10}$, or exploit a structural restriction on which affine fibres can contain many thin-group parabolics. None follows from pigeonhole alone. This is a useful failure: it shows why the fixed residue flag in Proposition II.20.9 is essential. One may average over the modulus, but not over the target incidence value.

The hidden unipotent orbit in an incidence fibre

Although cardinality at the random scale proves nothing by itself, an incidence fibre is not an arbitrary random fibre. It has an exact description in the full group. For $P = (p,q)^t$, define

$$N_P = \begin{pmatrix} -pq & p^2 \\ -q^2 & pq \end{pmatrix}, \qquad n_P(s) = I + sN_P.$$

Then $N_P^2 = 0$, $n_P(s)$ fixes $[P]$, and, if $r = \det(P, v) \neq 0$,

$$n_P(t/r)v = v + tP. \tag{1023}$$

For a canonical cusp column w, its positive primitive parabolic is $U_w = I + \lambda N_w$. Since $N_{gw} = gN_wg^{-1}$, equation (1023) gives

$$U_{v+tP} = n_P(t/r)U_v n_P(-t/r). \tag{1024}$$

Thus the parameter set of one fibre can be written as

$$\{t \in \mathcal{O}_K : n_P(t/r)U_v n_P(-t/r) \text{ is a positive primitive parabolic of } G_7\}.$$

There is an equally useful description on the group itself. Use the rational frame (P, Q), with $\det(P, Q) = 1$. In that frame

$$D_g(P) = r \quad \Longleftrightarrow \quad g = \begin{pmatrix} a & (ad-1)/r \\ r & d \end{pmatrix}.$$

For $r \neq 0$, if

$$u(z) = \begin{pmatrix} 1 & z \\ 0 & 1 \end{pmatrix}, \qquad w_r = \begin{pmatrix} 0 & -1/r \\ r & 0 \end{pmatrix},$$

then

$$\{g \in \mathrm{SL}_2(K) : D_g(P) = r\} = U_P w_r U_P, \tag{1025}$$

because

$$u(a/r)w_r u(d/r) = \begin{pmatrix} a & (ad-1)/r \\ r & d \end{pmatrix}.$$

The zero fibre is different: it is the rational Borel B_P. Thus a nonzero incidence fibre is a double-horospherical surface, not a generic hyperplane section of the ambient group. Large intersections with such surfaces are exactly the kind of phenomenon for which affine-sieve and horospherical counting methods are designed; their existence does not force the zero fibre.

This also measures what a generic "escape from subvarieties" theorem would have to accomplish. The balanced shell has $e^{R+o(R)}$ elements and at most $e^{(9/10+o(1))R}$ possible incidence values. To make pigeonhole alone force $r = 0$, every nonzero surface (1025) would need to contain fewer than $e^{(1/10-o(1))R}$ shell elements. In the usual notation

$$\#(\Omega_R \cap V) \leq e^{(1-\kappa+o(1))R},$$

this demands $\kappa > 9/10$. Qualitative exponential escape supplies only some positive κ, and the double-unipotent geometry gives no reason for the required near-total loss. Therefore Zariski density plus escape from proper algebraic subvarieties cannot replace the pointed inverse theorem. This is the intersection of the complete thin parabolic conjugacy class with one algebraic unipotent-conjugacy orbit.

Lemma II.20.53 (Algebraic equality subgroups of an incidence coefficient)**.** *Let k be a characteristic-zero field, let $Q, e \in k^2$ be linearly independent, and put*

$$F_Q(g) = \det(Q, ge) \qquad (g \in \mathrm{SL}_2).$$

If a positive-dimensional connected algebraic subgroup H satisfies $F_Q|_H = F_Q(1)$, then H is contained in either the unipotent subgroup fixing e or the unipotent subgroup fixing Q.

Proof. The connected positive-dimensional proper subgroups of SL_2 are tori, unipotent groups, and subgroups of Borels. A torus writes $e = e_+ + e_-$ in its two weight lines, and $F_Q(ae_+ + a^{-1}e_-)$ cannot be a nonzero constant in a. A Borel contains a torus, so it is excluded for the same reason; the whole SL_2 is excluded because its orbit of e is Zariski open in $k^2 \setminus \{0\}$.

It remains to consider a unipotent group N_L, with fixed line L. If $e \in L$, it fixes e. Otherwise

$$n(s)e = e + s\ell$$

for a nonzero $\ell \in L$, after rescaling s. Constancy of $F_Q(n(s))$ says $\det(Q, \ell) = 0$, so $L = kQ$ and N_L fixes Q. □

The first subgroup in the lemma is the right parabolic which does not change the cusp column ge. The second is the group n_Q in (1023). Thus the elementary algebraic equality classification is complete: no torus or semisimple subgroup can lie in a nonzero incidence level. A torus reappears only as the *normalizer* of n_Q, acting by unit dilation as in (1026). The quantitative inverse theorem must show that excess entropy forces one of these exact algebraic configurations.

The desired exceptional alternatives are built into this formulation. If $n_P(s) \in G_7$ for some $s \neq 0$, then $[P]$ is a cusp. If $hP = \varepsilon P$ for a hyperbolic $h \in G_7$, then

$$hn_P(s)h^{-1} = n_P(\varepsilon^2 s), \tag{1026}$$

because $hN_Ph^{-1} = N_{hP} = \varepsilon^2 N_P$. Conversely, every projectively nonidentity element of G_7 normalizing n_P fixes its unique boundary point $[P]$, and the rational-Borel lemma gives the cusp or special alternatives.

This isolates a second sufficient conjecture which uses the actual thin group and no prime-by-prime analysis.

Conjecture II.20.54 (Power-saving unipotent-intersection rigidity)**.** *Suppose $[P]$ has trivial projective stabilizer in G_7. There is $\delta_P > 0$ such that, uniformly in nonzero incidence values r, every fibre*

$$\{v \in \mathcal{C}_R(\epsilon) : \det(P, v) = r\}$$

has at most $O_{P,\epsilon}(e^{(1/10-\delta_P)R})$ elements for all sufficiently small fixed ϵ.

Proposition II.20.55 (Conditional completion from unipotent intersection)**.** *Assume the Hodge volume law $\#\mathcal{C}_R(\epsilon) = e^{R-o(R)}$ and Conjecture II.20.54. Then every rational flag has a projectively nontrivial stabilizer, and hence is a cusp or a special hyperbolic fixed point.*

Proof. For fixed primitive P, the incidence values of the balanced columns lie in a Minkowski box containing $e^{(9/10+2\epsilon+o(1))R}$ algebraic integers. Choose $\epsilon > 0$ with $2\epsilon < \delta_P/2$. Pigeonhole gives a fibre of size at least

$$e^{(1/10-2\epsilon-o(1))R},$$

If its incidence is zero, $[P]$ is already a cusp. Otherwise this contradicts the conjectural upper bound if the stabilizer is trivial. The rational-Borel lemma then supplies the classification. □

Route II.20.56 (What remains deep in the unipotent formulation)**.** Equations (1023) and (1024) are elementary matrix identities. The substantive assertion is the uniform power saving for the intersection of a thin parabolic conjugacy class with the moving K-unipotent orbit. Ambient unipotent-orbit rigidity in the Hilbert modular threefold does not immediately give it: every K-rational n_P already defines an arithmetic unipotent subgroup there. One must use the intersection

with the Teichmüller monodromy itself, and equality must force the normalizer relation (1026). This is precisely the place where a global Arakelov intersection estimate, rather than a list of prime-by-prime arguments, could enter.

Two existing orbit-counting frameworks stop well before the required statement. Aoun proves exponential escape of a non-elementary random walk from each fixed proper algebraic subvariety [42]. Applied schematically to the conjugacy curve in (1024), this gives some unspecified loss e^{-cR} from the full orbital count. Here one needs the much larger loss $e^{-(9/10+\delta)R}$, uniformly as both r and the base parabolic move. The theorem therefore supplies neither the exponent nor the uniform family required by the pigeonhole argument.

Higher-rank local mixing for Anosov subgroups gives precise counts on affine symmetric spaces [43]. The diagonal heptagonal representation is not Borel–Anosov: powers of its parabolic generator have only logarithmic Cartan growth. Moreover, the orbit in (1024) is a unipotent conjugacy curve, not one of the affine symmetric subgroups covered by that counting theorem. Thus the promising input would have to be a *relative*, cuspidal analogue with an equality classification for the unipotent normalizer. This comparison narrows the target; it does not furnish the power saving.

Tilting the Hodge window creates a genuine entropy excess

The preceding power-saving conjecture asks for more than the raw pigeonhole argument can justify. There is, however, a way to move the argument off the mean without choosing a prime or a preferred generator. For small fixed $t > 0$, let $\mathcal{C}_R^-(t)$ be the set of canonical primitive cusp columns $v \in G_7 e$ satisfying

$$H_0(v) \le e^{R/2}, \qquad H_1(v)H_2(v) \le e^{(2/5-t)R}, \qquad H_i(v) \le e^{R/2} \quad (i = 1, 2). \tag{1027}$$

The first exponent is the central hyperbolic size. The value $2/5$ is the typical sum of the two outer Hodge exponents. Thus t tilts the *whole parabolic orbit* slightly toward stronger outer contraction; it does not modify individual words after they have been found.

The thermodynamic input needed for this tilt is the following local statement about the full projective Hodge skew product.

Conjecture II.20.57 (Quadratic Hodge tilt)**.** *There are $t_0 > 0$ and a rate function $I : [0, t_0] \to \mathbb{R}_{\ge 0}$ such that*

$$\#\mathcal{C}_R^-(t) = \exp\big((1 - I(t) + o_t(1))R\big), \qquad I(t) = O(t^2) \quad (t \downarrow 0). \tag{1028}$$

The quadratic order is what an analytic, nondegenerate pressure function at its equilibrium slope predicts. It is not being inferred from a scalar digit potential: the operator must retain both moving outer projective directions, as in (1046). Establishing analytic pressure for the parabolically accelerated, infinite-branch matrix cocycle is a real theorem still to be proved.

Proposition II.20.58 (Pressure reduction for the quadratic tilt)**.** *Let τ be the return time of the accelerated Bowen–Series map and let Ψ be the logarithmic norm cocycle of $\rho_{12} = \rho_1 \otimes \rho_2$ on its moving projective direction. Suppose the projective transfer operator*

$$(\mathscr{M}_{s,q} f)(z, [u]) = \sum_{h \in \mathscr{H}} |h'(z)|^s e^{q\Psi(h,[u])} f\big(hz, [\rho_{12}(h)u]\big) \tag{1029}$$

has, near $(s, q) = (1, 0)$, a simple leading eigenvalue separated from the rest of its spectrum on a Hölder space, and that the resulting equilibrium state has exponential tails in τ. Then Conjecture II.20.57 holds.

Proof. Write $\Lambda(s,q)$ for the leading eigenvalue. Analytic perturbation and the implicit-function theorem give an analytic curve $s = s(q)$ on which $\Lambda(s(q),q) = 1$, with

$$s(0) = 1, \qquad s'(0) = \frac{\int \Psi \, d\mu_0}{\int \tau \, d\mu_0} = \frac{2}{5}.$$

The second derivative is the asymptotic variance of $\Psi - \frac{2}{5}\tau$, divided by $\int \tau \, d\mu_0$. The cocycle is not cohomologous to a multiple of τ: periodic words ST^n, with n large enough that all three conjugates are hyperbolic, have outer-to-central spectral ratios tending to a value different from $2/5$. Hence that variance is positive.

The Legendre transform of $s(q)$ gives the orbital large-deviation rate for Ψ/τ. At its minimum $2/5$ it has a nondegenerate quadratic Taylor expansion, so the rate at $2/5 - t$ is $I(t) = O(t^2)$. The Gibbs cylinder estimate and renewal in the suspension then give

$$\#\{v : H_0(v) \le e^{R/2},\ \log(H_1(v)H_2(v)) \le (2/5 - t)R\} = e^{(1 - I(t) + o_t(1))R}.$$

The individual bounds $H_i(v) \le e^{R/2}$ remove only an exponentially smaller large-deviation set near $t = 0$. This is (1028). □

Most of the hypotheses in this reduction have a direct group-theoretic check. The Zariski closure of $(\rho_1, \rho_2)(G_7)$ is SL_2^2, so its tensor representation is irreducible and contains proximal elements. On a parabolic branch of size n, $|h'| \asymp n^{-2}$ and $\|\rho_{12}(h)\| \ll n^2$; consequently (1029) has a finite moment in a two-sided neighborhood of $(1,0)$. The scalar accelerated Hecke operator is already available in the literature [50]. The remaining analytic verification is the spectral gap after adjoining the moving ρ_{12}-projective direction. This is substantially narrower than assuming the counting law itself.

Now fix a primitive $P \in \mathcal{O}_K^2$, and put $r(v) = \det(P, v)$. For $v \in \mathcal{C}_R^-(t)$,

$$|\sigma_0 r(v)| \ll_P e^{R/2}, \qquad |\sigma_1 r(v)\, \sigma_2 r(v)| \ll_P e^{(2/5 - t)R}, \qquad |\sigma_i r(v)| \ll_P e^{R/2}. \tag{1030}$$

The nonzero algebraic integers in this region number

$$\exp\big((9/10 - t + o(1))R\big). \tag{1031}$$

Here is a global proof which avoids a boundary estimate near the coordinate axes. For $0 \ne r \in \mathcal{O}_K$ in (1030),

$$|\,\mathrm{N}_{K/\mathbb{Q}}(r)| = |\sigma_0 r\, \sigma_1 r\, \sigma_2 r| \ll_P e^{(9/10 - t)R}.$$

There are $O_K(X)$ integral ideals of norm at most X, so this gives at most $e^{(9/10 - t + o(1))R}$ possibilities for the principal ideal (r). The generators of one such ideal differ by a unit. Under the logarithmic embedding, the units form a rank-two lattice in $x_0 + x_1 + x_2 = 0$. The three individual upper bounds and the product bound in (1030), together with the fixed value of $\log|\,\mathrm{N}(r)|$, confine the possible unit translate to a polygon of diameter $O_P(R)$. It contains $O_P(R^2)$ unit-lattice points. The ideal count times this polynomial proves (1031). Thus all finite places enter at once through the norm, and all archimedean ambiguity enters through one bounded region in the Arakelov unit plane.

Equations (1028) and (1031) therefore force either an incidence zero, or a nonzero fibre of size at least

$$\exp\big((1/10 + t - I(t) - o_t(1))R\big). \tag{1032}$$

This is exponentially larger than the *untilted* mean $e^{R/10}$. It is, however, exactly the mean occupancy of the tilted shell: the population exponent is $1-I(t)$ and its incidence-volume exponent is $9/10-t$. Their difference is $1/10+t-I(t)$. The linear gain therefore does not yet represent exceptional concentration.

The following conjecture was initially proposed to turn that comparison into a proof. The audit below shows that it is an overstrong target, not a consequence of thinness.

Conjecture II.20.59 (Overstrong incidence-entropy bound)**.** *If $[P]$ has trivial projective stabilizer in G_7, then, uniformly in $0 \neq r \in \mathcal{O}_K$,*

$$\#\left\{ v \in G_7 e : \begin{array}{c} H_0(v) \le e^{R/2},\ H_1(v)H_2(v) \le e^{2R/5}, \\ H_i(v) \le e^{R/2}\ (i=1,2),\ \det(P,v) = r \end{array} \right\} \le \exp\big((1/10 + o_P(1))R\big). \tag{1033}$$

This conjecture bounds every tilted subfibre by the mean of the larger untilted window. That is not the natural shell-dependent mean.

Theorem II.20.60 (Formal implication of the overstrong bound)**.** *Conjectures II.20.57 and II.20.59 imply that every point of $\mathbf{P}^1(K)$ is either a cusp of G_7 or the fixed point of a special hyperbolic element.*

Proof. Fix $[P]$, and suppose its projective stabilizer is trivial. Choose a small fixed $t > 0$ for which $t - I(t) > 0$. Pigeonhole applied to $\mathcal{C}_R^-(t)$ and (1031) gives either a column with $\det(P,v) = 0$, in which case $[P] = [v]$ is a cusp, or a nonzero fibre with the lower bound (1032). Since the tilted window is contained in the window of (1033), this contradicts Conjecture II.20.59. Hence the stabilizer is nontrivial. The rational-Borel proposition identifies its projectively nonidentity element as parabolic or hyperbolic, giving the two stated alternatives. □

Route II.20.61 (Why the Hodge tilt does not create rigidity)**.** Let $\mathfrak{a} = (r)$, and on a dyadic Hodge shell put

$$\epsilon_i \asymp \frac{|\sigma_i r|}{H_i(v)}.$$

The natural adelic hyperplane law is

$$\Pr\{\det(P,v) \in \mathfrak{a},\ d_i([v_i],[P_i]) \ll \epsilon_i\ (i=0,1,2)\} \asymp \frac{\epsilon_0\epsilon_1\epsilon_2}{\mathrm{N}\,\mathfrak{a}}. \tag{1034}$$

The product formula gives

$$\frac{\epsilon_0\epsilon_1\epsilon_2}{\mathrm{N}(r)} \asymp \frac{1}{H_0(v)H_1(v)H_2(v)}.$$

For the tilted shell this is $e^{-(9/10-t+o(1))R}$. Multiplying by $\#\mathcal{C}_R^-(t) = e^{(1-I(t)+o(1))R}$ reproduces (1032) exactly. Thus even perfect adelic equidistribution predicts the supposedly large fibre.

An abstract random model makes the logical failure transparent. Take $e^{(1-I(t))R}$ labelled inputs and distribute them uniformly over $e^{(9/10-t)R}$ labelled outputs. Every fibre then has the exponent $1/10+t-I(t)$, with no stabilizer or algebraic structure at all. The tilt has selected a smaller output space after counting; it has not made one fibre exceptional relative to that selection.

Consequently Conjecture II.20.59 cannot be motivated by a Frostman or entropy argument. It would have to forbid the ordinary large-deviation correlation displayed in (1034), and is essentially another encoding of the desired rigidity. The valid alternatives identified before the tilt remain: keep r fixed through infinitely many scales, or fix the rational residue target before selecting the group element. The latter is exactly the horizontal optimal-lifting route.

The attempted ten-block route and its valid residue

There is nevertheless a useful full-group identity behind the failed entropy argument. Let μ_R be normalized counting measure on a thin Hodge shell of cusp columns, and let

$$\pi_P(v) = \det(P, v) \in K.$$

If a return is cut into ten macroscopic blocks, $g = g_1 \cdots g_{10}$, the incidence can be exposed one block at a time through the symplectic identity

$$\det(P, g_1 \cdots g_{10} e) = \det(g_1^{-1} P, g_2 \cdots g_{10} e). \tag{1035}$$

Hence a block acts simultaneously on a moving dual line and a moving primal line. This identity uses arbitrary macroscopic elements of the thin group; it is not a calculation with the Rosen generators.

What failed was the proposed entropy bookkeeping. In the tilted shell the incidence image has entropy at most $(9/10 - t + o(1))R$, not $(9/10 + o(1))R$. Nine increments of size $R/10$ are therefore impossible in the random adelic model from Route II.20.61. Pressure can make the tilted shell typical for a Gibbs measure, but it cannot enlarge the incidence image of that shell. Thus the ten-block identity remains useful, while the claimed ninth entropy increment was exactly where the argument silently inserted the desired conclusion.

There is already a rigorous, but nonoptimal, consequence of one midpoint cut.

Proposition II.20.62 (The central point–line bound)**.** *Let $\mathcal{V}_R$ be a collection of cusp columns for which each $v \in \mathcal{V}_R$ has a chosen factorization*

$$v = abe, \qquad a, b \in G_7, \qquad d(o, ao), d(o, bo) \le R/2 + o(R).$$

If $[P]$ has trivial projective stabilizer, then, uniformly in $0 \neq r \in \mathcal{O}_K$,

$$\#\{v \in \mathcal{V}_R : \det(P, v) = r\} \le e^{(2/3 + o(1))R}.$$

Proof. Put $X = \{a^{-1}P\}$ and $Y = \{be\}$, using the distinguished real embedding. Lattice orbital growth gives $|X|, |Y| \le e^{R/2 + o(R)}$. For fixed nonzero $\sigma_0 r$, every $x \in X$ defines the affine line

$$L_x = \{y \in \mathbb{R}^2 : \det(x, y) = \sigma_0 r\}.$$

Two such lines coincide only if their defining vectors x coincide. The map $a \mapsto a^{-1}P$ has bounded multiplicity because any collision gives a projective stabilizer of $[P]$. Likewise the chosen pair (x, y) determines $v = abe$ up to a bounded central ambiguity. The Szemerédi–Trotter theorem therefore bounds the number of relevant pairs by

$$O\big(|X|^{2/3}|Y|^{2/3} + |X| + |Y|\big) = e^{(2/3 + o(1))R}.$$

□

If the chosen block decomposition has lengths $R/2 + o(R)$, the proposition is a genuine estimate for complete blocks of the thin group. Its exponent $2/3$ is much weaker than the shell-dependent mean fibre exponent $1/10 + t - I(t)$, so it does not approach classification. Its value is diagnostic: one real point–line incidence theorem uses the full words but discards the two conjugate embeddings and every finite place.

Why scalar congruence does not rescue the tilted argument

Let $\nu_{t,R}$ be normalized counting measure on a Gibbs-typical subset of $\mathcal{C}_R^-(t)$. Choose a good squarefree ideal $\mathfrak{q}$, independently of the incidence value, with

$$\log \mathrm{N}\,\mathfrak{q} = (9/10 + o(1))R.$$

The tempting estimate was

$$\sup_{a \in \mathcal{O}_K/\mathfrak{q}} \nu_{t,R}\{v : \det(P, v) \equiv a \pmod{\mathfrak{q}}\} \le (\mathrm{N}\,\mathfrak{q})^{-1+o(1)} \tag{1036}$$

It has the same defect as the archimedean entropy conjecture. The tilted incidence region contains only $e^{(9/10-t+o(1))R}$ Minkowski cells. When reduction modulo $\mathfrak{q}$ resolves those cells, its occupied residue set therefore has at most that size, and the natural largest-atom scale is $e^{-(9/10-t+o(1))R}$. The right side of (1036) is smaller by $e^{-tR+o(R)}$. No expansion theorem can manufacture the missing support; asking for this estimate asks the tilted measure to be uniform on residues it cannot reach.

Fourier inversion makes the thin-group content explicit:

$$\nu_{t,R}\{\det(P, v) \equiv a \pmod{\mathfrak{q}}\} = \frac{1}{\mathrm{N}\,\mathfrak{q}} \sum_{\xi \in (\mathcal{O}_K/\mathfrak{q})^\vee} \mathbf{e}_{\mathfrak{q}}(-\xi a) \sum_v \nu_{t,R}(v) \mathbf{e}_{\mathfrak{q}}(\xi \det(P, v)). \tag{1037}$$

After a block split $v = abe$, the inner sum is a bilinear exponential sum in $a^{-1}P$ and be. The formula is correct and may still be useful for a fixed target. In the tilted-fibre argument, however, sum–product or flattening was being asked for $9R/10$ of entropy from a scalar projection whose actual support has only $(9/10 - t)R + o(R)$ entropy. Fourier inversion does not remove that information-theoretic obstruction.

The distinction from horizontal optimal lifting is decisive. There the rational point and its residue class are fixed first, and the question is whether the complete orbit hits that prescribed class at a radius below the norm-closing threshold. No incidence value is selected after the orbit has been counted. Formula (1037) can legitimately enter that fixed-target problem, especially after averaging over external ideals.

Known dimension results clarify why they do not repair the fibre argument. Hochman–Solomyak compute the exact dimension of the Furstenberg measure for a finitely supported algebraic $\mathrm{SL}_2(\mathbb{R})$ random walk [44]; Rapaport proves exact dimensionality and a Ledrappier–Young formula for higher-dimensional Furstenberg measures under strong irreducibility and proximality [46]. These concern stationary projective measures. They do not by themselves give min-entropy for the joint three-embedding affine coefficient, do not handle the parabolically accelerated infinite alphabet with the required conditioning, and do not classify equality as a K-rational normalizer. More fundamentally, their dimension is computed for the measure actually sampled. Applied after a Hodge tilt, a correct affine extension would see the reduced support exponent $9/10 - t$, and would reproduce the conditional mean rather than contradict it.

Rush's recent formula for the *Frostman* dimension of an $\mathrm{SL}_2(\mathbb{R})$ Furstenberg measure remains a useful warning [47]. Frostman dimension controls every small ball uniformly and includes the worst large-deviation alignment with a prescribed projective point. But the relevant uniform exponent is still the exponent of the conditioned measure. A joint three-embedding Frostman theorem for the tilted Gibbs state should have affine exponent $9/10 - t$, not $9/10$. Exceptional alignment could lower that exponent further; it cannot supply the missing tR of support.

The real-Lie-group sum–product theorem of He–de Saxcé is relevant to the inverse mechanism in a fixed-target argument [48]. It applies to representations with no trivial Jordan–Hölder quotient and says, roughly, that a set of group operators which stays away from proper stabilizers expands vectors at a discretized scale. For $G = \mathrm{SL}_2(\mathbb{R})^3$, the direct sum of the three standard modules has no trivial quotient, and G is perfect. The full Zariski closure of the Galois representation rules out confinement to a normal factor. Since $[K : \mathbb{Q}] = 3$, there is also no intermediate trace subfield which could support a diagonal-factor exception. These facts leave the stabilizer of a projective line, hence a Borel or its torus normalizer, as the relevant P-dependent obstruction.

It does not validate the discarded shell-independent atom bound: a sum–product theorem cannot expand a projection beyond its Hodge support. For a fixed residue target, two upgrades would still be necessary. First, the published conclusion fills a ball using bounded sums and products; it does not give a no-holes theorem for one prescribed matrix coefficient under a stopped group convolution. Second, closeness to a real stabilizer is not yet an exact K-rational stabilizer. A nonzero algebraic defect of height $e^{O(R)}$ can be e^{-cR}-small at one real place. Killing it requires simultaneous control of the other two real places and its finite ideal—exactly the adelic, or Arakelov, strengthening in the preceding paragraph. Thus the Lie-group sum–product theorem identifies the correct algebraic exceptional sets, while the sharp exponent and the global exactification remain to be supplied by a Hodge-balanced adelic fixed-target theorem.

Route II.20.63 (Outcome of the thin-group entropy route)**.** The block-duality identity and the midpoint point–line proposition are valid uses of arbitrary long elements of the thin group. The proposed nine-increment completion is not valid: its ninth increment exceeds the support entropy of the tilted shell. Neither pressure, finite Fourier inversion, Furstenberg dimension, Frostman dimension, nor sum–product changes that fact. A viable full-group theorem must instead keep one incidence value fixed through many scales, or prescribe the rational target before the word is chosen. The horizontal no-holes problem does the latter and is now the principal route.

Proposition II.20.64 (Why compression of shapes alone cannot prove fibre rigidity)**.** *Let $\alpha \in \mathcal{O}_K$ satisfy*

$$|\sigma_0\alpha| > 1, \qquad |\sigma_1\alpha| < 1, \qquad |\sigma_2\alpha| < 1,$$

and put $T = \{\alpha^n : n \geq 0\}$. The joint Galois-ratio set of nonzero differences $T - T$,

$$\left\{ \left(\frac{|\sigma_1(\alpha^n - \alpha^m)|}{|\sigma_0(\alpha^n - \alpha^m)|}, \frac{|\sigma_2(\alpha^n - \alpha^m)|}{|\sigma_0(\alpha^n - \alpha^m)|} \right) : n \neq m \right\},$$

has compact countable closure, and every sequence with $\min\{m, n\} \to \infty$ converges to $(0, 0)$. Nevertheless T is infinite.

Proof. For $n > m$,

$$\alpha^n - \alpha^m = \alpha^m(\alpha^{n-m} - 1).$$

For each fixed gap $k = n - m$, the two ratios are a fixed pair multiplied coordinatewise by

$$\left(\left| \frac{\sigma_1\alpha}{\sigma_0\alpha} \right|^m, \left| \frac{\sigma_2\alpha}{\sigma_0\alpha} \right|^m \right),$$

and hence tend to $(0, 0)$. If the gap also tends to infinity, the factors $\sigma_i(\alpha^k) - 1$ remain bounded while $|\sigma_0(\alpha^k) - 1|$ grows exponentially, giving the same limit. The only remaining points occur for bounded m or bounded k; diagonal extraction proves compactness and countability. □

The nonunit case is not hypothetical in the heptagonal field. The element $\varepsilon = \lambda^2 - 1$ is a unit with one dominant real embedding and two embeddings of absolute value less than one. Hence $\alpha = 2\varepsilon^k$ satisfies the hypotheses for every sufficiently large k, while $|\,\mathrm{N}\,\alpha| = 8$.

This model identifies the datum omitted by the ratio set. If α is a unit, the logarithmic vector of α^n is motion in the rank-two unit lattice. Such a progression is the local model for a special hyperbolic orbit and must remain as an equality case. If α is not a unit, then

$$\log|\,\mathrm{N}(\alpha^n)| = n \log|\,\mathrm{N}\,\alpha| > 0.$$

The same Galois shapes now carry a growing finite degree. McMullen's compression cannot distinguish these two cases because it divides every outer quantity by the central one.

Route II.20.65 (What the Pisot countermodel still proves)**.** Even if a future argument produces a genuinely exceptional *fixed* fibre, its rigidity cannot follow from the topology or multiplicativity of the compression ratios alone. It must retain radial Arakelov degree. The Pisot model proves that a purely archimedean shape estimate cannot distinguish unit holonomy, which is the desired special equality case, from nonunit radial escape. This is a separate obstruction from the random-scale defect in Proposition II.20.51: the first says the fibre is not exceptional, and the second says that compression would still forget essential finite-place data even if it were exceptional.

II.20.13 Why Hilbert-modular torus rigidity does not directly close the argument

There is a natural temptation to place the preceding shape–degree pair inside the full Hilbert modular space

$$\mathrm{SL}_2(\mathbb{R})^3/\,\mathrm{SL}_2(\mathcal{O}_K)$$

and invoke rigidity of the maximal split torus. Tomanov proves that for a totally real field of degree greater than two, a locally divergent maximal split-torus orbit is either closed or dense [54]. The closed case is indeed governed by a rational torus and therefore resembles the unit-holonomy alternative.

The hypotheses do not match the present trajectory. The lifted Teichmüller curve is

$$z \longmapsto \big(z, F_1(z), F_2(z)\big),$$

where the two outer coordinates are nonlinear holomorphic modular embeddings. It is not an orbit of an algebraic subgroup of $\mathrm{SL}_2(\mathbb{R})^3$, and its geodesic flow is not the restriction of one fixed three-dimensional split torus. During a deep cusp excursion the trajectory is asymptotic to a diagonal segment in the frame of that cusp. At the next excursion the frame is changed by a new element of G_7. Conjugating all of these segments into one fixed torus would require the cusp frames to preserve one rational Borel. By the rational-Borel reformulation above, that is already the desired cusp or special conclusion.

There is a second mismatch. Tomanov's theorem concerns the ambient arithmetic lattice $\mathrm{SL}_2(\mathcal{O}_K)$, whereas the relevant orbit is the Zariski-dense, infinite-index diagonal copy of G_7. Density statements for values on all of $\mathcal{O}_K^2$ do not give separation or recurrence for this thin subset.

Thus homogeneous torus rigidity supplies the right geometric analogy but not the missing implication. A usable theorem would have to be a *stopped torus-rigidity theorem with thin-group frame changes*: either the accumulated radial degree is strictly positive, or all frame changes lie in one rational Borel. This is a precise reformulation of the shape-plus-degree problem, not an application of the existing Hilbert-modular orbit-closure theorem.

II.20.14 The earlier, less efficient, congruence shrinking-target route

There is a group-wide use of varying congruence quotients which is different from the discarded 2-adic closure calculation and from tracking newly appearing prime divisors of successive remainders. Fix a primitive $P = (p, q)^t \in \mathcal{O}_K^2$, put $x = [P]$, and define

$$D_g(P) = \det(P, gP) \in \mathcal{O}_K.$$

This is the fixed-point defect: it vanishes exactly when g fixes x. If $gx \equiv x \pmod{\mathfrak{a}}$, then

$$D_g(P) \in \mathfrak{a}. \tag{1038}$$

Let

$$R(g) = d_{\mathbb{H}}(i, \sigma_0(g)i), \qquad a_i(g) = \log \|\sigma_i(g)\|_{\mathrm{op}}.$$

One has $a_0(g) \le R(g)/2 + O(1)$. Suppose, for ideals $\mathfrak{a}_n$ of unbounded norm, there are projectively nonidentity $g_n \in G_7$ for which

$$g_n x \equiv x \pmod{\mathfrak{a}_n}, \tag{1039}$$

$$R(g_n) \le A \log \mathrm{N}\,\mathfrak{a}_n + O(1), \tag{1040}$$

$$|\sigma_0(g_n)x_0^\dagger - x_0^\dagger| \le e^{-\beta R(g_n)}, \tag{1041}$$

$$a_1(g_n) + a_2(g_n) \le \theta R(g_n) + O(1). \tag{1042}$$

In an affine chart, the elementary identity

$$\sigma_0(D_g(P)) = q(cp + dq)\big(x - \sigma_0(g)x\big)$$

gives

$$\log |\sigma_0 D_g(P)| \le \left(\frac{1}{2} - \beta\right) R(g) + O_P(1).$$

At the outer places,

$$|\sigma_i D_g(P)| \le \|\sigma_i(P)\|^2 \|\sigma_i(g)\|_{\mathrm{op}}.$$

Multiplication of the three estimates proves the following unconditional criterion.

Proposition II.20.66 (Congruence–Hodge closing criterion)**.** *Under* (1042)*, if*

$$A\left(\frac{1}{2} - \beta + \theta\right) < 1, \tag{1043}$$

then $D_{g_n}(P) = 0$ for all sufficiently large n. Hence x is a cusp or a special hyperbolic fixed point.

Proof. Put $\kappa = 1/2 - \beta + \theta$. The archimedean estimates give

$$|\,\mathrm{N}\, D_{g_n}(P)| \le C_P \exp(\kappa R(g_n)).$$

If $\kappa \le 0$, this is eventually smaller than $\mathrm{N}\,\mathfrak{a}_n$. If $\kappa > 0$, the bound for $R(g_n)$ and (1043) give the same conclusion. If the determinant were nonzero, (1038) would give the reverse weak inequality. Thus it vanishes. The rational-Borel classification of a nontrivial stabilizer gives the last assertion. □

The entropy calculation is unusually sharp. There are $e^{R+o(R)}$ elements in a central hyperbolic ball. Membership in the Borel fixing $x \bmod \mathfrak{a}$ has expected density $(\mathrm{N}\,\mathfrak{a})^{-1}$, while a real target of radius $e^{-\beta R}$ has expected density $e^{-\beta R}$. At

$$R = A \log \mathrm{N}\,\mathfrak{a},$$

the expected count is exponentially large precisely when

$$A(1-\beta) > 1. \tag{1044}$$

The typical outer Hodge exponents give

$$\theta = \frac{1}{2}\left(\frac{1}{5} + \frac{3}{5}\right) = \frac{2}{5}.$$

Combining (1043) and (1044) leaves the nonempty interval

$$\frac{9}{10} - \frac{1}{A} < \beta < 1 - \frac{1}{A}. \tag{1045}$$

Its width is exactly $1/10$, the Asai contraction gap in hyperbolic-distance normalization. With $A = 2$, for instance, the permitted window is $2/5 < \beta < 1/2$.

Conjecture II.20.67 (Borel congruence shrinking-target dichotomy)**.** *For every $x \in \mathbf{P}^1(K)$, either its G_7-stabilizer is nontrivial, or, for some $A > 1$ and some β in (1045), there are ideals and elements satisfying (1042) with $\theta = 2/5 + o(1)$.*

The proposition proves the full classification from this conjecture. The proposed analytic input is a joint local limit theorem for three objects attached to the same full-group element:

1. its action on the projective congruence orbit;
2. its shrinking real return to the prescribed boundary point; and
3. the two outer Hodge cocycles.

Strong approximation and expander machinery address the first distribution, hyperbolic transfer operators address the second, and large deviations for the modular-embedding cocycles address the third. Their separate statements do not imply the required joint result at the entropy scale. A special cyclic stabilizer is a genuine concentration obstruction and must occur as the alternative, rather than be discarded as an error term.

Route II.20.68 (Outcome of the congruence route)**.** This route produces a new unconditional closing lemma and locates a quantitative $1/10$ interval in which a full-group counting theorem would finish the proof. It does not use the previously computed image modulo 2^n: the ideals tend to infinity and are selected before the group element. Nor does it label a prime born in a remainder. Once an external modulus divides the fixed-point defect, all finite information is compressed into $\mathrm{N}\,\mathfrak{a}$, and all archimedean information is combined by the product formula. The unresolved part is the joint shrinking-target theorem, not another estimate for an individual generator.

II.20.15 Exact transfer-operator target and its entropy obstruction

The phrase "joint local limit theorem" can be made precise. Accelerate the Bowen–Series coding of the presentation $G_7 \simeq C_2 * C_7$ across the parabolic branch, and denote its inverse branches by $\mathscr{H}$.

For a branch h, let $\ell(h,z) = -\log|h'(z)|$. The outer Hodge data are not scalar potentials on the base: they are the matrix cocycles $\sigma_i(h)$ acting on moving directions. Put

$$\psi_i(h,[u]) = \log\frac{\|\sigma_i(h)u\|}{\|u\|} \qquad (i=1,2).$$

Thus the state space must retain $(z,[u_1],[u_2])$, as well as the finite projective orbit

$$X_{\mathfrak{a}}(x) = G_7\cdot[P \bmod \mathfrak{a}]$$

The operator which contains all three distributions is

$$(\mathscr{L}_{s,\mathbf{v},\mathfrak{a}}f)(z,[u_1],[u_2],\xi) = \sum_{h\in\mathscr{H}} |h'(z)|^s e^{v_1\psi_1(h,[u_1])+v_2\psi_2(h,[u_2])} f\big(hz,[\sigma_1(h)u_1],[\sigma_2(h)u_2],h\xi\big). \tag{1046}$$

This is genuinely a whole-group object. The relation $R^7=1$, the parabolic acceleration, the congruence quotient, and the two Galois representations all enter the same spectral problem.

For $\epsilon>0$, define $N_x(R,\mathfrak{a},\beta,\epsilon)$ by the four conditions

$$\begin{aligned} R\le R(g)<R+1, &\qquad gx\equiv x \pmod{\mathfrak{a}},\\ |\sigma_0(g)x_0^\dagger - x_0^\dagger|\le e^{-\beta R}, &\qquad a_1(g)+a_2(g)\le(2/5+\epsilon)R. \end{aligned} \tag{1047}$$

A sufficient spectral theorem is

$$N_x(R,\mathfrak{a},\beta,\epsilon) \gg_{x,\epsilon} \frac{e^{(1-\beta)R}}{\#X_{\mathfrak{a}}(x)}, \tag{1048}$$

uniformly through a sequence on which the right side tends to infinity, unless x already has a cyclic stabilizer. A stronger and cleaner version would give a relative error $O(e^{-\eta R})$ whenever

$$\log\#X_{\mathfrak{a}}(x) \le (1-\beta-2\eta)R.$$

Equation (1048), followed by the congruence–Hodge closing criterion, proves the classification.

The boundary localization in this statement is not an innocuous test function. A target interval of radius $e^{-\beta R}$ fixes a symbolic cylinder of length about βR. Ordinary large deviations for the unrestricted operator do not control the Hodge cocycle after this prefix has been prescribed. At large β, proving that the prescribed prefix has typical outer growth is nearly the trace-field Oseledets rigidity statement already isolated above. This is why merely listing strong approximation, transfer operators, and large deviations would be circular.

The obstruction has an exact numerical form. Write

$$r = \frac{\log\#X_{\mathfrak{a}}(x)}{R}.$$

Positive expected count requires $r+\beta<1$. The product-formula closing estimate, with outer exponent sum $2/5$, requires $r>9/10-\beta$. Therefore every successful choice lies in

$$0<1-r-\beta<\frac{1}{10}. \tag{1049}$$

Changing the parameter $A=1/r$ only moves the difficulty. With A near one, very little of the boundary word is fixed, but reduction modulo $\mathfrak{a}$ must mix at almost the information-theoretic rate. With large A, the congruence quotient is small enough for ordinary expansion, but the real target prescribes almost nine tenths of the word. No scale choice creates more than the one-tenth margin supplied by the Asai gap.

Route II.20.69 (Outcome of the transfer-operator audit)**.** The full thin group gives the operator (1046); this is substantially more structure than inequalities for the six Rosen moves. It also shows exactly why a black-box mixing argument is insufficient. The required new input is a cylinder-uniform Hodge large-deviation estimate with congruence twisting and a periodic concentration alternative, at an entropy loss strictly smaller than 1/10. This is the analytic form of the same stopped boundary-rigidity problem, not a separate easy route.

II.20.16 Arithmetic holonomy bounds: an exact applicability audit

Arithmetic holonomy estimates [57] are attractive here because the three modular embeddings arise from one hypergeometric local system. Their present Diophantine form, however, starts with a fixed finite collection of holonomic functions, fixed coefficient-denominator types, and a fixed Apéry limit. Analytic continuation then bounds approximation to that one target. In the heptagonal closing problem the evaluated line and its denominator type both change with the return word: the target is $[Y(gP)]$, not a fixed period independent of g.

The CM orbit formulas give an even more decisive applicability test. The natural elliptic-center approximant is

$$\beta_g = \frac{E_g}{2D_g} \in K, \qquad x - \beta_g = \frac{A_g(P)}{2qD_g}. \tag{1050}$$

Hence the target and approximant already belong to the same field K. In the norm-one-torus coordinate, the corresponding target is the root of unity -1. Theorem 2.6 of the cited work bounds K-rational approximation to a fixed system of local Apéry limits; its irrationality content cannot force equality in (1050). For a target lying in the approximation field, the nonzero lower bound is an elementary height bound, while the equality case is precisely what has to be produced. This is not repaired by optimizing the irrationality exponent.

Consequently the existing holonomy inequality cannot simply be applied to the family of return lines. A valid application would first require a single holonomic normal function on the rational-flag bundle whose different analytic continuations produce all the incidences $\det(P, g^{-1}e)$, with denominator growth uniform in g. Constructing that normal function would itself supply the missing transport from the Hodge metric to the integral incidence line.

Route II.20.70 (Outcome of the arithmetic-holonomy audit)**.** Holonomy bounds could become a proof mechanism after the moving return lines are realized as branches of one fixed arithmetic differential system. The currently available fixed-target theorem does not perform that realization and therefore does not improve the 1/10 transfer budget or the stopped joint-compression problem. The useful concrete construction target is a holonomic extension of the joint compression functor with one rational endpoint, not a stronger irrationality measure for an already fixed period.

II.20.17 Algebraic leaves: why Bost's theorem proves the wrong kind of stabilizer

There is a tempting global Arakelov construction which deserves a separate audit. Let $\mathbb{V}$ be the rank-six restriction-of-scalars local system over the triangle orbifold and consider the projective flat bundle $\mathbf{P}(\mathbb{V})$. A line $[P] \in \mathbf{P}^1(K)$ determines a horizontal analytic leaf. Since Bost's algebraicity theorem for leaves combines all finite and infinite places by a slope inequality, it initially appears to be exactly the requested replacement for a prime-by-prime argument [67].

The monodromy calculation shows that it cannot give the desired classification. Let $\Gamma = G_7$, let

$\widetilde{C} = \mathbb{H}$, and let $\ell = [P]$ be a line in the Betti fiber. The horizontal leaf through ℓ is

$$\mathrm{Stab}_\Gamma(\ell)\backslash\widetilde{C} \longrightarrow \mathbf{P}(\mathbb{V}). \tag{1051}$$

If this leaf were an algebraic multisection of the projective bundle, its projection to the algebraic base curve would be generically finite. Consequently $\mathrm{Stab}_\Gamma(\ell)$ would have finite index in Γ. The Galois representation is Zariski dense, so it has no finite orbit on the projective line. Thus no line relevant here has an algebraic horizontal leaf.

This remains true in the desired special case. A special hyperbolic line has an infinite cyclic stabilizer. Its cover in (1051) is an infinite-degree cyclic cover of the triangle orbifold, not an algebraic multisection. Bost algebraicity would therefore be much stronger than the conclusion needed here and would incorrectly discard the special alternative.

Two hypotheses of the algebraicity theorem also fail before one reaches that conclusion.

1. A vector $P \in K^2$ is algebraic in the Betti monodromy basis. Its image in an algebraic de Rham fiber is obtained through the Betti–de Rham period torsor and is not thereby a rational point of the algebraic projective bundle. Treating these two rational structures as identical would assume an algebraicity of periods.

2. In characteristic p, the horizontal distribution on the projective connection is closed under p-th powers at $[P]$ precisely when the p-curvature preserves that line. A Gauss–Manin connection has nilpotent, not generally zero, p-curvature. Betti rationality of P supplies no preservation statement at almost every prime.

Route II.20.71 (Outcome of the algebraic-leaf audit)**.** Bost's method is relevant as a model for how an Arakelov slope sums local costs, but its algebraicity conclusion is not the required periodicity conclusion. It detects finite monodromy orbit, whereas the theorem sought here detects a cyclic stabilizer. The correct object must therefore be a *stopped real boundary leaf* or an incidence line along the geodesic flow, not the full complex horizontal leaf of the flat connection. This explains why the maximal-slope record construction can borrow Bost's all-place bookkeeping but cannot be replaced by his leaf theorem.

Route II.20.72 (August 6 synthesis: the actual bottleneck)**.** The current manuscript now takes the integral Asai obstruction, rather than the finite-level local limit, as its unconditional proof spine. For a noncuspidal $[P] \in \mathbf{P}^1(K)$, compact returns and discreteness of the integral Asai lattice force the sum of the two outer Hodge growth rates to be at least 1. The Bowen–Margulis values are $1/5$ and $3/5$, whose sum is $4/5$. Thus a counterexample is not an arbitrary nonperiodic Rosen orbit: it must lie in a definite $1/5$-exceptional Lyapunov stratum. The remaining theorem is to show that a trace-field ray in this stratum is periodic. Full algebraic Rosen genericity is more than is needed: the sufficient conjecture is only that every nonperiodic trace-field ray has outer upper-growth sum strictly below 1 along compact returns.

The global-modulus argument remains useful, but its logical status has been corrected. Let $M_R = e^{R-o(R)}$ be the size of the Hodge-balanced shell and let $|Y_R| = e^{uR+o(R)}$, $9/10 < u < 1$, be the projective adelic target. A return of the distinguished vertex $[P]$ modulo the one global ideal is an exact characteristic-zero stabilizer by the product formula. If the stabilizer is trivial, that return count is therefore identically zero for large R. A proposed asymptotic with positive main term is not an ordinary equidistribution input: it already contains the desired rigidity in inverse form.

The finite target is best viewed as a coherent Schreier tower. A bare Ramanujan or transfer-operator norm gives error $M_R^{1/2}$, so it beats the main term $M_R/|Y_R|$ only when $u < 1/2$. The Hodge norm comparison requires $u > 9/10$. What is needed is instead a based return at the

information-theoretic time $R \sim \log |Y_R|$, plus the theorem that a coherent rational vertex with anomalously large injectivity radius gives a rational Borel stabilizer. He–de Saxcé localizes large Fourier coefficients near rational torsion, and the complete local transform computes the signed torsion weights, but neither result proves this coherent resonance-to-Borel implication.

Winsor's height does not change this diagnosis. If

$$H_{\max}([p:q]) = \prod_i \max(|\sigma_i p|, |\sigma_i q|), \qquad H_\Sigma([p:q]) = \prod_i (|\sigma_i p| + |\sigma_i q|),$$

then

$$H_{\max} \le H_\Sigma \le 8 H_{\max}.$$

The logarithms differ by at most $\log 8$, and both are unit invariant. Thus the sum height is useful for displaying the orbit but cannot repair an unbounded failure of projective-height descent. The substantive information in Winsor's computation is different: all tested $q = 7$ expansions terminate or become periodic, there are at least twenty-four observed G_7-orbits, and the special alternative is much richer than the formerly conjectured single noncusp orbit.

The two viable next attacks are therefore genuinely the same problem in different language. The geometric version asks for a stopped maximal-slope theorem controlling both unit-shape coordinates of the explicit transported Asai tensor. The analytic version asks for a pointed entropy-time return/resonance dichotomy on the coherent Schreier tower. In either version, equality must produce a parabolic line or a unit eigenline. Prime-by-prime divisor bookkeeping, an ordinary spectral gap, and a generic scalar Roth bound do not address this implication.

Route II.20.73 (Exact comparison with Filip's integral-vector theorem)**.** Filip's Theorem 4.5.3 gives the closest existing Hodge-theoretic model for the desired argument. For a log-Anosov representation preserving an integral lattice, a nonzero integral vector in the strictly decreasing boundary subbundle forces the endpoint to be parabolic. At compact returns the vector would otherwise converge to zero, contradicting the positive lower bound for nonzero vectors in one fixed lattice.

This theorem does not apply directly to the heptagonal Asai vector. For

$$H = \mathrm{Res}_{K/\mathbf{Q}} \mathrm{SL}_2$$

the log-Anosov root comes from the distinguished factor. The four weights of the Asai endpoint tensor above its central contracting line are

$$(-1,+1,+1), \quad (-1,+1,-1), \quad (-1,-1,+1), \quad (-1,-1,-1).$$

The two outer entries are weights of the Levi factor. They make all four weights incomparable with zero in the distinguished-root order. The pairwise tensors improve the real sign to $-a_0 + a_i$, but restriction of scalars restores their three cyclic Galois components; the negative degree in one component is accompanied by positive degree in another.

This failure is necessary, not technical. If a special element satisfies $hP = \varepsilon P$, then

$$\mathrm{As}(h)Y(P) = \mathrm{N}(\varepsilon)Y(P) = \pm Y(P).$$

The normalized state

$$\mathbf{U}(X) = \frac{X^{\otimes 2} \otimes \tau X \otimes \tau^2 X}{v}$$

is likewise unchanged projectively under $X \mapsto \varepsilon X$. A strict-decay theorem would therefore erase the special alternative which the classification must retain.

The correct target is an equality-case extension. At a first norm record, the content of $\mathbf{U}(r_j, r_{j+1})$ is $(r_{j+1})^{-1}$, its distinguished norm is the pairwise energy, and its projective class determines the complete Rosen state. One must prove a uniform alternative:

$$\mathscr{E}(r_j, r_{j+1}) \le 6h(P) + B_7,$$

or a remainder has vanished, or two earlier projective states agree. The first case bounds all record norms and the fixed-support S-unit theorem closes the orbit. The second is a cusp. In the third case the intervening word fixes the initial rational line and supplies the special unit holonomy.

This audit changes the geometric work program. There is no reason to search for another representation with a more negative weight. The normalized Asai representation already has the correct finite lattice, real weight, transport scalar, and equality locus. The remaining task is stopped semipositivity of its rank-two maximal-slope polygon after quotienting by the unit lattice. The Green–Higgs factors give the interior curvature and the fixed divisor $3P_7$; the missing step is to prove that their toroidal height jump at a first record is nonnegative, with zero jump only on one recurrent projective state.

The existing positivity theorem for biextension height jumps does not supply that step directly. Its test object is a holomorphic algebraic curve mapping to a normal-crossings compactification. A trace-field endpoint is rational in the Betti local system, but its orbit is a horizontal real leaf of the projective flat bundle; it descends to an algebraic section of the triangle curve only after a monodromy stabilizer has been found. The record times are discrete points on that real leaf, not a holomorphic arc meeting the boundary. Thus the published theorem predicts the correct local sign but has the wrong test category. The needed extension is precisely positivity for a stopped horizontal real leaf with a maximal-slope integral lattice.

Route II.20.74 (Affine-framing audit of the normalized Asai state)**.** The latest formulation used

$$\mathbf{U}(X) = \frac{X^{\otimes 2} \otimes \tau X \otimes \tau^2 X}{v}, \qquad X = (u, v)^t,$$

whose content ideal is $(v)^{-1}$ and whose distinguished tensor norm is the pairwise energy. There is an exact limitation which has to be imposed on every proposed Hodge proof:

$$[\mathbf{U}(X)] = [X^{\otimes 2} \otimes \tau X \otimes \tau^2 X].$$

Division by v changes an affine lift, but it changes no projective point. Thus the reciprocal record ideal is not a function on the projective flat bundle. It is the lattice framing of the affine cone:

$$\left((v)^{-1}, \mathbf{U}(X) \in \mathbb{W} \otimes (v)^{-1}\right).$$

Clearing the fractional content by v removes the finite record degree, and the product formula makes the resulting change in the real norms cancel it exactly.

This invalidates one overly optimistic version of the height-jump route. A projective semipositivity theorem could control the direction $[\mathbf{U}(X)]$ and identify its recurrent equality locus. It could not bound $\log|\,\mathrm{N}\, v|$, because two affine lifts with different content determine the same point of projective space. The correct test object is a record-framed affine tensor line, or equivalently the maximal-slope coordinate line together with its saturated quotient map from the dual normalized Asai lattice.

Route II.20.75 (The record framing is a saturated quotient)**.** The preceding affine datum has an exact module-theoretic form. Put

$$N(X) = X^{\otimes 2} \otimes \tau X \otimes \tau^2 X, \qquad \mathcal{R}_v = v^{-1}\mathcal{O}_K, \qquad \mathbf{U}(X) = N(X)/v.$$

Since $N(X)$ is primitive, evaluation on $\mathbf{U}(X)$ is a surjection

$$q_X : \mathbb{W}_{\mathcal{O}}^{\vee} \twoheadrightarrow \mathcal{R}_v, \qquad \ell \longmapsto \ell(\mathbf{U}(X)).$$

For $X' = AX = (u', v')^t$, exact transport gives

$$q_{X'}(\ell) = \frac{v}{v'} q_X(\ell \circ \rho_{\mathbb{W}}(A)), \qquad \frac{v}{v'}\mathcal{R}_v = \mathcal{R}_{v'}.$$

Thus the record line is a saturated quotient, not a positive subline. The bottom arrow is an exact finite-lattice isomorphism, while its three real logarithmic norms sum to

$$\log|\,\mathrm{N}\, v| - \log|\,\mathrm{N}\, v'|.$$

The whole record jump is therefore already the archimedean height of the quotient transition.

This route succeeds as a diagnosis and fails as a closing inequality. A generic Bost slope inequality for a quotient, or the fact that $\mathcal{R}_{v'}$ is a maximal-degree coordinate at a first record, returns only this identity. It supplies no negative defect. What is needed is a theorem for the particular decomposable quotient selected simultaneously by the rational endpoint, the Hodge filtration, and the first-record rule. Calling this a "Hodge-selected quotient" prevents ordinary maximal-slope formalism from being mistaken for the desired estimate.

The three natural real norms of the same tensor make the remaining twist explicit. If $E_i(X) = \log \|\sigma_i \mathbf{U}(X)\|_{\infty}$, then exact expansion of the tensor gives

$$E_i = 3h(X) + \eta_i, \qquad S := E_0 + E_1 + E_2 = 12h(X) - \log|\,\mathrm{N}\, v|,$$

and

$$\Delta := 2E_0 - E_1 - E_2 = 2\eta_0 - \eta_1 - \eta_2, \qquad \mathscr{E}(X) = E_0 = \frac{S + \Delta}{3}.$$

Thus the natural Galois-symmetric lattice controls S, whereas the stopped theorem asks for one selected norm E_0. Their exact difference is the transverse Hodge character represented, up to the metric sign, by $\mathcal{N}^{-1}$; its square is $\mathcal{N}^{-2} \simeq \mathcal{O}(3P_7)$. This rules out another tempting shortcut: symmetric Asai semistability cannot simply be relabelled as the pairwise-energy estimate. The order-seven divisor must be coupled to the saturated record quotient in the same inequality.

Route II.20.76 (One fundamental Hodge line and its arithmetic shape coordinate)**.** The cyclic Picard ledger contains a simplification which was not used in the earlier maximal-slope formulation. The proved triviality

$$\mathcal{Q}_2 = \mathcal{L}_0\mathcal{L}_1\mathcal{L}_2^{-2} \simeq \mathcal{O}_C$$

implies actual orbifold-line identities

$$\mathcal{N}_1 \simeq \mathcal{N}_2^{\otimes 2}, \qquad \mathcal{N} \simeq \mathcal{N}_2^{\otimes 3}.$$

After compatible normalization, the Higgs ratios satisfy

$$\Xi_1 = \Xi_2^2, \qquad \Xi = \Xi_2^3.$$

Thus the divisors $2P_7$ and $3P_7$ are powers of the one fundamental divisor P_7. The two pairwise curvatures are not independent.

This gives an exact decomposition of the relative transport. If $\mathcal{A}_i = \Theta_i^{\mathrm{rat}} - \Theta_i^{\mathrm{Hod}}$, then

$$\mathcal{A}_1 + \mathcal{A}_2 = 3\mathcal{A}_2 + \Theta_{\mathrm{neut}}^{\mathrm{rat}} - B(gz) + B(z),$$

where B is bounded on a cusp-adapted reduced cross-section and

$$\Theta_{\mathrm{neut}}^{\mathrm{rat}} = \Theta_1^{\mathrm{rat}} - 2\Theta_2^{\mathrm{rat}}.$$

For the denominator multiplier $m = j_g(x)$, this last term is the distinguished logarithm of the explicit field element

$$\nu(m) = \frac{m\tau(m)}{\tau^2(m)^2} = \frac{\mathrm{Nm}_{K/\mathbb{Q}}(m)}{\tau^2(m)^3}.$$

If $a_i = \log|\sigma_i(m)|$, put

$$s = a_0 + a_1 + a_2, \qquad c = a_2 - a_0, qquadf = a_0 + a_1 - 2a_2 = \log|\sigma_0\nu(m)|.$$

Then the change of coordinates is invertible:

$$a_0 = \frac{s-f}{3} - c, qquada_1 = \frac{s+2f}{3} + c, qquada_2 = \frac{s-f}{3}.$$

The record rule concerns the radial coordinate s; the fundamental Hodge line concerns the curved coordinate c; and f is the flat coordinate. This is an exact linear identity, not a heuristic division of the error terms. It also exhibits the remaining difficulty: signed control at the distinguished real place does not control the positive parts at every real and finite place which make up a Weil height.

There is, however, a further algebraic identification. Define the fundamental arithmetic coordinate

$$\alpha_2(m) = \frac{\tau^2(m)}{m}.$$

Its distinguished real logarithm is exactly the curved rational character:

$$\log|\sigma_0\alpha_2(m)| = a_2 - a_0 = c.$$

Moreover one has the exact field identity

$$\nu(m) = \alpha_2(m)^{-2}\tau(\alpha_2(m))^{-1}.$$

Modulo rational scalars, α_2 is the endomorphism $\tau^2 - 1$ of the two-dimensional shape torus. Its determinant is $(\zeta^2 - 1)(\zeta - 1) = 3$, so it is a degree-three isogeny. Consequently the full adelic height of $\alpha_2(m)$ already controls both shape directions. The analytically flat character is not a second independent arithmetic variable; it is a fixed monomial in the Galois conjugates of the coordinate paired with the fundamental P_7-Hodge line. Its finite divisor is

$$(\alpha_2(m)) = (\tau^2 m)/(m),$$

which is exactly the ratio of the corresponding two integral Asai component ideals. Hence the curved Hodge norm, the finite ideal ratio, and the full Galois shape height are three realizations of

the same K-line, not three quantities joined by analogy. If the state denominator changes from v to v', then with $m = v'/v$ one has

$$\alpha_2(m) = \frac{\tau^2(v')/v'}{\tau^2(v)/v}.$$

So this coordinate is exactly the height jump of the integral Asai shape line across the record transition. This also resolves the earlier pre-record/post-record indexing problem: the same scalar v'/v occurs in the normalized tensor transport and in this shape jump.

At a completely split rational prime, if $e_j = v_{\tau^j\mathfrak{p}}(m)$, the three valuations of $\nu(m)$ are

$$(e_0 + e_2 - 2e_1,\ e_1 + e_0 - 2e_2,\ e_2 + e_1 - 2e_0).$$

Their sum is zero and they all vanish exactly when $e_0 = e_1 = e_2$. At an inert prime the coefficient $1 + 1 - 2$ makes the valuation vanish automatically. Hence $\nu(m)$ is the cyclic ideal Laplacian: it detects exactly the Galois-asymmetric part of the divisor of m, while the kernel is a rational radial ideal, up to the fixed ramified prime.

The remaining problem therefore has one fundamental Hodge coordinate, carried by $\mathcal{N}_2^{-1}$ and P_7, and one corresponding K-valued arithmetic coordinate $\alpha_2(m)$. Its one-place Hodge logarithm does not control its full height, but its Galois completion does. Powers of a special unit show why an all-place statement must retain the periodic equality case. With fixed finite support the derived coordinate $\nu(m)$ is an S-unit and the already proved S-unit closing theorem applies.

For comparison, the flat coordinate itself is also global and proper. On the two-dimensional shape torus

$$\mathbb{T}_K^{\mathrm{sh}} = \big(\mathrm{Res}_{K/\mathbb{Q}}\,\mathbb{G}_m\big)/\mathbb{G}_m$$

one has

$$[\nu(m)] = [\tau^2(m)^{-3}].$$

Thus ν induces the degree-nine isogeny $[-3] \circ \tau^2$. In particular, the height of $\nu(m)$ is equivalent to the full non-radial height of m; it sees both independent unit directions and every Galois-asymmetric ideal direction at once. The more economical fundamental coordinate gives the degree-nine isogeny

$$m \longmapsto \big(\mathrm{Nm}_{K/\mathbb{Q}}(m), [\alpha_2(m)]\big)$$

from $\mathrm{Res}_{K/\mathbb{Q}}\,\mathbb{G}_m$ to $\mathbb{G}_m \times \mathbb{T}_K^{\mathrm{sh}}$. Consequently its two target heights form a proper height on $K^\times$. Using ν instead gives an equivalent degree-twenty-seven map.

This yields an alternative closing criterion which asks only for a bounded subsequence of complete state multipliers. If, at infinitely many record states, the state multiplier $m = r_{j+1}/r_j$ has bounded norm height plus bounded full α_2-height, Northcott leaves only finitely many exact projective states. Two repeat, so determinism of the Rosen map gives a periodic state; termination gives the cusp alternative. On such a period, unit holonomy may enlarge individual real representatives, but the field norms repeat and create only finitely many records. The remaining task can therefore be formulated as a stopped bound for one proper global multiplier height, rather than a bound on a list of local supports.

There is no hidden change of dynamical variable here:

$$[r_j : r_{j+1}] = [1 : m], \qquad h_K(m) = h([r_j : r_{j+1}]).$$

The radial-plus-shape height furnished by the degree-nine isogeny is an ample re-expression of the ordinary projective height of the complete Rosen state. At compact geodesic returns the outer

compensation excess is the same state height up to the bounded metric comparison proved below. The first-record and compact-return subsequences are not identical, but the arithmetic quantity being controlled is now literally the same one.

This route therefore succeeds in replacing a vague rank-two Harder–Narasimhan problem by one fundamental arithmetic–Hodge coordinate. It does not by itself prove the stopped bound for the combined multiplier height. The division of labor is now sharp. The degree-nine and degree-twenty-seven isogenies are elementary global algebraic number theory. The concentration of curvature in P_7 is Hodge theory. The assertion that the particular multiplier sequence selected by a rational flag has a bounded record subsequence, unless that flag has a stabilizer, is the thin-group input. It is false for an arbitrary sequence in $K^\times$, since units can escape forever. It also cannot be read from the six generators one at a time. The required statement is a relative-domination theorem for the whole G_7-orbit and is consequently a genuine arithmetic–Hodge–thin group hybrid.

The remaining wall is visible in one three-number ledger. The Galois Hodge degrees of the α_2-line are

$$-\frac{1}{14}, \qquad \frac{1}{7}, \qquad -\frac{1}{14}.$$

They sum to zero. The two contracting conjugates can be paid for exactly by the one expanding conjugate, and a Weil height sees the positive part rather than the signed sum. This is why an ordinary Arakelov-degree or semistability argument returns equality. What remains to prove is the orbit-specific equality theorem: repeated compensation at first-record times forces recurrence of the rational flag and hence a nontrivial G_7-stabilizer. This is now the precise obstruction; it is no longer an ambiguity about which tensor, divisor, or record index is being used.

The exact program `hodge_fundamental_line_tests.py` checked the Galois cube identity on 5000 random elements of coefficient size 10^5; $\tau^3 = 1$, both expressions for $\nu(m)$, and the monomial identity in $\alpha_2(m)$ agreed identically, while the displayed logarithmic character agreed to $5.4 \cdot 10^{-15}$.

Route II.20.77 (Exact all-place Roth criticality)**.** There is a second, independent cancellation which explains why a standard optimal Roth theorem still gives no strict surplus. Choose $Q \in \mathcal{O}_K^2$ so that (P, Q) is unimodular. For a cusp row w, put

$$L_P(w) = wP, \qquad L_Q(w) = wQ, \qquad H(w) = \prod_z \|w\|_z,$$

where the product is over all places. If both evaluations are nonzero, then the product formula gives the identity

$$\prod_z \frac{|L_P(w)|_z |L_Q(w)|_z}{\|w\|_z^2} = H(w)^{-2}.$$

At a finite prime dividing the remainder $L_P(w)$, the corresponding factor becomes small. The product of the three real values of that same remainder grows by the reciprocal amount. Hence retaining every finite valuation does not improve the exponent: the two-form product is identically at the Roth threshold 2, for every row, before any dynamical estimate is made.

This is stronger than the previous observation that the elliptic-center target and approximants both lie in K. It tests the most favorable fixed pair of linear forms and still finds exact criticality. A strict gain must use the nonparallel Hodge weight $(-2, +1, +1)$, while the preceding route shows that the repeated central copy must remain attached to the affine record lattice. The unresolved statement is therefore not a better scalar irrationality measure. It is a stopped, framed slope inequality whose zero-defect projective direction is allowed to recur.

Route II.20.78 (Numerical consequence of the indexing audit)**.** The pre-record pairwise inequality and the post-record maximal-slope inequality are not the same statement. If r_{j+1} is a first norm record, the finite content of $\mathbf{U}(r_j, r_{j+1})$ is exactly $(r_{j+1})^{-1}$; the tensor at the preceding state contains $(r_j)^{-1}$. The post-record version is therefore the one compatible with the affine framing.

The exact program `rosen_pairwise_stopping_dynamics.py` was rerun on 1722 reduced inputs from the coefficient box $[-5,5]^3$ and rational denominators through 5. It followed 26765 exact transitions and found 1500 strict record transitions. The largest observed pre-record and post-record excesses over $6h(P)$ were respectively

$$8.6576775900\ldots, \qquad 8.6576775900\ldots,$$

and the exact cocycle identity held to $5.7 \cdot 10^{-14}$ in the displayed real evaluation. These bounded observations do not prove the inequality. Their purpose is narrower: they show no immediate numerical counterexample to the correctly indexed post-record target, while the balanced inverse-unit family already recorded above disproves stronger monotone-reservoir substitutes.

Route II.20.79 (Sublinear rational-Borel factorization and the Morse audit)**.** The common state height admits an exact ambient-group interpretation. For an integral matrix A, put

$$\mathfrak{h}_\Delta(A) = \sum_{i=0}^{2} \log \max\{1, \|\sigma_i(A)\|_\infty\}.$$

Given a primitive state X, Dirichlet unit balancing produces a unit ε_X such that all three real sizes of $Q = \varepsilon_X X$ are $\exp(h([X]) + O_K(1))$. A geometry-of-numbers reduction of the affine Bezout lattice

$$\{(a,b) \in \mathcal{O}_K^2 : Q_1 b - Q_2 a = 1\}$$

then completes Q to $H_Q \in \mathrm{SL}_2(\mathcal{O}_K)$ with

$$\mathfrak{h}_\Delta(H_Q) \ll_K h([X]) + 1.$$

Fix a frame H_P for the initial endpoint. For every return element $g \in G_7$, take $X = gP$, put $C = H_Q H_P^{-1}$, and set $b = C^{-1}g$. The calculation

$$bP = \varepsilon_X^{-1} P$$

gives the exact factorization

$$g = Cb, \qquad C \in \mathrm{SL}_2(\mathcal{O}_K), qquad b \in B_P(K) \cap \mathrm{SL}_2(\mathcal{O}_K), qquad \mathfrak{h}_\Delta(C) \ll_K h([gP]) + 1.$$

Thus compensation $\mathcal{D}_P(t) = o(t)$ means precisely that the Galois-diagonal thin orbit tracks the ambient rational Borel B_P with sublinear transverse error. Bounded error gives an actual intersection by Northcott. The new boundary problem is whether sublinear error does.

The Hodge modular embeddings prove a genuine prefix theorem in the other direction. If $d_i(g)$ is displacement in the i-th real factor and $L_{\mathrm{thick}}(g)$ is thick-part length of the segment from a fixed basepoint to its g-translate, integration of Schwarz–Pick gives

$$2d_0(g) - d_1(g) - d_2(g) \geq \kappa_7 L_{\mathrm{thick}}(g).$$

Equivalently the two blocks $\mathrm{Sym}^2 \rho_0$ and $\rho_1 \otimes \rho_2$ have a singular-value gap growing with relative thick displacement. This proves the relative-domination input for every prefix of the full group, rather than merely for closed words.

The standard Morse theorem for a Borel-Anosov subgroup of a higher-rank symmetric space cannot supply the desired intersection [39]. Full Borel-Anosov regularity fails here: the joint outer image is dense, so there are infinitely many elements for which both outer factors remain in a fixed compact set while the discrete central factor escapes. Hence the Galois-diagonal orbit is not uniformly regular in all three simple-root directions. Keeping only the central root restores the known log-Anosov property, but then the Morse boundary remembers only the central endpoint P, not its two rational Galois flags. That is information already present before the proof begins.

This route therefore produces two rigorous advances—the controlled Borel factorization and open-segment relative domination—and isolates a new precise target:

If return prefixes in G_7 track one K-rational Borel with sublinear all-place transverse height, then G_7 actually intersects that Borel nontrivially.

It is an orbit-specific relative-Morse statement for the modular-embedding curve. Applying a full-Anosov theorem would assume a regularity condition which is contradicted by the proved outer density.

Route II.20.80 (The strict Green–Higgs boundary branch)**.** The individual logarithmic-derivative factors

$$\Psi_{i,x}(z) = \frac{F_i'(z)(z - x_0^\dagger)}{F_i(z) - x_i^\dagger}$$

give a rigorous boundary alternative. Put $z_t = x_0^\dagger + ie^{-t}$. If, for either outer embedding,

$$F_i(z_t) \longrightarrow x_i^\dagger, \qquad \operatorname{Im} F_i(z_t)/e^{-t} \longrightarrow \delta_i \in (0, \infty),$$

then Julia–Carathéodory gives $F_i(z_t) - x_i^\dagger = i\delta_i e^{-t} + o(e^{-t})$. The Hodge norm of the corresponding rational outer line has logarithm $-t + O(1)$. The central rational line also has exponent $-t$, while Schwarz–Pick lets the remaining factor grow by at most $t + O(1)$. Hence the norm of the integral Asai vector is at most $e^{-t+O(1)}$. Compact-return lattice discreteness is impossible, so the endpoint is a cusp.

Equivalently, finite rational angular matching is the boundary equality

$$|\Psi_{i,x}(z_t)| \longrightarrow 1.$$

The angular hypothesis is stronger than necessary. At any compact return the integral Asai norm is bounded below, the central rational factor has logarithmic squared norm $-t + O(1)$, and each outer Busemann function is at most $t + O(1)$ by Schwarz–Pick. Therefore each outer rational-line norm is individually bounded below:

$$\frac{|F_i(z_t) - x_i^\dagger|^2}{\operatorname{Im} F_i(z_t)} \geq c_{P,C} > 0 \qquad (i = 1, 2).$$

Thus entry into arbitrarily deep horoballs based at either rational conjugate line already forces a cusp. Finite angular matching is just one way of entering such a horoball.

A counterexample must consequently satisfy a uniform *double-horoball mismatch* at every compact return. This is compatible with the special alternative: for a permitted special unit both rational outer eigenlines are repelling, not forward Hodge lines. The route proves this entire strict branch but does not classify persistent double-horoball mismatch. The missing assertion must either produce a bounded return of the common state height or align the two repelling flags into one periodic unit holonomy. The pointwise inequalities alone still telescope, so the stopped affine lattice framing remains necessary.

Route II.20.81 (Exact equality corridor after the Filip audit)**.** The comparison with Filip separates the final argument into a deep boundedness statement and an elementary closing statement. Let z_t be the geodesic ray to $x = [P] \in \mathbf{P}^1(K)$, and at compact return times put

$$\mathcal{D}_P(t) = \log \|\sigma_1(P)\|_{1,t} + \log \|\sigma_2(P)\|_{2,t} - t.$$

Since the distinguished endpoint line has Hodge norm $e^{-t+O_P(1)}$, the logarithmic norm of the integral Asai vector is

$$\log \|Y(P)\|_t = \mathcal{D}_P(t) + O_P(1).$$

If g_t reduces the return to a fixed compact set, all three Hodge norms there are uniformly comparable with fixed Euclidean norms. Since $g_t P$ remains primitive, the ordinary height formula gives the sharper identity

$$\mathcal{D}_P(t) = 6h([g_t P]) + O_{P,C}(1).$$

Here the Hodge quantities are logarithms of squared norms, so the compact comparison is with twice the logarithmic Euclidean norm. The earlier coefficient (3) mixed squared Hodge norm with unsquared Euclidean norm; the closing argument was unaffected, but the exact coefficient is (6). Dirichlet's unit theorem now gives a useful affine version. The three numbers

$$H_i(Q) = \log \max\{|\sigma_i Q_1|, |\sigma_i Q_2|\}$$

have sum $3h([Q])$, and the logarithms of units form a cocompact lattice in the sum-zero plane. Hence a unit ε_Q may be chosen so that

$$|H_i(\varepsilon_Q Q) - h([Q])| \le C_K \qquad (i = 0, 1, 2).$$

Thus $\mathcal{D}_P(t) = o(t)$ is equivalent to simultaneous subexponential size of all three vectors $\sigma_i(\varepsilon_t g_t P)$. This is the precise sense in which the boundary problem is sublinear tracking of one rational Hilbert-modular Borel after quotienting by its rank-two unit lattice. If $\mathcal{D}_P(t)$ is bounded above on infinitely many compact returns, Northcott leaves only finitely many reduced projective states. Two agree, so the intervening element of G_7 fixes x. The rational Borel calculation then gives a cusp or a special unit eigenline.

The same argument has a useful quantitative limit. Balanced primitive representatives of points of height at most H lie in a rank-six Minkowski box, so there are $O_K(e^{6H})$ such projective points. Since $\mathcal{D} = 6h + O(1)$, N distinct reduced returns force

$$\max_{n \le N} \mathcal{D}_P(t_n) \ge \log N - O(1).$$

Thus ordinary Northcott counting cannot close a merely sublinear orbit: logarithmic escape in the return index is already enough to avoid its pigeonhole. This calculation identifies the quantitative wall; it is not offered as the missing Hodge inequality.

This proves an exact unconditional alternative: every counterexample must satisfy

$$\mathcal{D}_P(t) \longrightarrow +\infty$$

along every compact-return sequence. The proposition is stronger than the earlier exponent lower bound. It identifies the least exceptional remaining corridor:

$$\mathcal{D}_P(t) \to +\infty, \qquad \mathcal{D}_P(t) = o(t).$$

In this corridor the outer exponent is exactly 1, but the reduced rational orbit escapes to infinity in Weil height too slowly for an exponent argument to detect. A complete exceptional-ray theorem must also exclude the linearly supercritical outer-growth strata; the point of this audit is that the boundary stratum has now been sharpened from nonnegative exponent to unbounded height drift.

Filip's theorem settles the strict negative-weight branch because the adapted norm tends to zero. It does not settle this corridor: after restriction of scalars the fundamental degrees are

$$-\frac{1}{14}, \qquad \frac{1}{7}, \qquad -\frac{1}{14}.$$

The signed total is zero, while a Weil height takes positive parts. Thus ordinary Arakelov degree or semistability permits the middle component to drift indefinitely. Conversely, once that drift is bounded, no new "equality version" of lattice discreteness is required; the preceding finite-lattice argument already closes the orbit. The new theorem must therefore establish bounded compensation, or periodic unit compensation, from the stopped thin-group trajectory.

This also fixes the role of ambient Hilbert-modular torus rigidity. The curve

$$z \longmapsto (z, F_1(z), F_2(z))$$

is not an orbit of one split torus: the outer coordinates are nonlinear modular embeddings, and every cusp excursion uses a different rational Borel frame. Identifying all those frames with one torus is equivalent to finding a recurrent Borel, which is already the desired stabilizer. The ambient theorem is therefore a model for the equality classification, not a missing citation which proves it.

The most economical remaining target is now the following. At first record times, prove bounded liminf of

$$h_{\mathbb{Q}}(\operatorname{Nm} m) + h_{\mathrm{sh}}([\tau^2(m)/m]), \qquad m = r_{j+1}/r_j,$$

unless the rational Borel recurs. The map $m \mapsto (\operatorname{Nm} m, [\tau^2(m)/m])$ is a finite isogeny, so this is a proper height and Northcott supplies the closing step. Hodge theory supplies the two negative degrees, algebraic number theory supplies the proper height, and the thin-group theorem must prevent the single positive conjugate from drifting in the equality corridor. These are three parts of one argument, not three competing explanations.

The equality corridor has a sharper unit-lattice form. Write $B_i(t) = \log \|\sigma_i(P)\|_{i,t}$, reduce by g_t, and choose a unit ε_t which balances $\varepsilon_t g_t P$ at the three real places. Since the height is $o(t)$, compact comparison gives

$$2\log|\sigma_0\varepsilon_t| = t + o(t), \qquad 2\log|\sigma_i\varepsilon_t| = -B_i(t) + o(t) \quad (i = 1, 2).$$

Double-horoball exclusion gives $B_i(t) \geq -O(1)$, Schwarz–Pick gives $B_i(t) \leq t + O(1)$, and equality compensation gives $B_1(t) + B_2(t) = t + o(t)$. Hence every normalized unit direction has an accumulation point

$$(1, -\theta_1, -\theta_2), \qquad \theta_i \geq 0, \quad \theta_1 + \theta_2 = 1.$$

This is an explicit compact segment in the logarithmic unit plane.

For an actual special period $hP = \varepsilon P$, oriented so that P is centrally attracting, horoball exclusion forces

$$|\sigma_0\varepsilon| > 1, \qquad |\sigma_1\varepsilon| < 1, \qquad |\sigma_2\varepsilon| < 1,$$

and its slope is

$$\theta_i = -\frac{\log|\sigma_i\varepsilon|}{\log|\sigma_0\varepsilon|}.$$

Thus the remaining equality theorem is no longer "control the units." The unit direction describes only the diagonal part of a possible neutral period. The theorem must first lock the complete rational Borel frame; only after that recurrence does the unit ray identify the semisimple holonomy. The next route gives the exact transition equation and records this correction.

Route II.20.82 (Exact Borel-frame transitions: what unit locking omits)**.** The controlled Borel factorization has an exact consecutive-return form. Write

$$g_n = C_n b_n, qquad b_n \in \mathscr{B}_P := B_P(K) \cap \mathrm{SL}_2(\mathcal{O}_K),$$

and put $a_n = g_{n+1} g_n^{-1}$. Then

$$a_n = C_{n+1} u_n C_n^{-1}, qquad u_n = b_{n+1} b_n^{-1} \in \mathscr{B}_P.$$

If H_P is an integral frame with first column P, the middle factor has the unique form

$$H_P^{-1} u_n H_P = \begin{pmatrix} \varepsilon_n/\varepsilon_{n+1} & \beta_n \\ 0 & \varepsilon_{n+1}/\varepsilon_n \end{pmatrix}, qquad \beta_n \in \mathcal{O}_K.$$

For the balanced columns $Q_n = C_n P = \varepsilon_n g_n P$, this gives

$$a_n Q_n = (\varepsilon_n/\varepsilon_{n+1}) Q_{n+1}.$$

Most importantly,

$$C_n \mathscr{B}_P = g_n \mathscr{B}_P,$$

and two such cosets agree if and only if the complete projective states $[g_n P]$ agree. The actual closing variable is therefore the arithmetic Borel coset, not its diagonal unit coordinate.

The entry β_n is gauge-dependent, not an additional invariant of the state. Replacing the chosen completion by

$$C_n' = C_n H_P \begin{pmatrix} 1 & s_n \\ 0 & 1 \end{pmatrix} H_P^{-1} \qquad (s_n \in \mathcal{O}_K)$$

does not change $C_n \mathscr{B}_P$, and changes β_n to

$$\beta_n + d_n s_n - d_n^{-1} s_{n+1}, \qquad d_n = \varepsilon_n/\varepsilon_{n+1}.$$

Consequently, bounding β_n is not the target. Its occurrence only shows that the diagonal unit does not determine a lifted transition; the complete coset is the gauge-invariant object which must recur.

This coset has a proper all-place height. If

$$\mathfrak{H}_P(g\mathscr{B}_P) = \inf_{b \in \mathscr{B}_P} h_\Delta(gb),$$

controlled unimodular completion and the product formula give

$$3h([gP]) - O_P(1) \le \mathfrak{H}_P(g\mathscr{B}_P) \le A_P h([gP]) + O_P(1).$$

At compact returns, where $\mathcal{D}_P(t) = 6h([gP]) + O(1)$, this makes the frame height comparable with the compensation excess. Thus the phrase "Borel tracking" has not discarded a hidden arithmetic coordinate: the missing assertion really is that this proper coset height cannot drift sublinearly along the specified thin-group return cocycle.

This corrects an overstatement in the unit-cone formulation. The unit cone identifies the Levi part of neutral holonomy. The lifted equation also has the gauge-dependent additive term β_n, and convergence of normalized unit directions does not make $C_n\mathscr{B}_P$ recur. The desired theorem is a *frame-locking* theorem: sublinear transverse height for this particular geodesic prefix cocycle must force repetition of the complete coset. Once that happens, the unit cone classifies the permitted semisimple holonomy.

Three standard black boxes were checked against this exact equation.

First, the Evertse–Schlickewei–Schmidt theorem controls a linear equation whose variables lie in one fixed finite-rank multiplicative group [29]. Here only the diagonal entry is a unit. The unipotent entry β_n and the entries of C_n range additively through $\mathcal{O}_K$; the transition equation supplies no closed linear S-unit equation in the diagonal entries alone. The fixed-support four-term equation remains valid after the remainder support has already been bounded, but it does not prove frame recurrence before that bound.

Second, the Ru–Vojta moving-target theorem requires hyperplanes in general position and a point sequence linearly nondegenerate over the coefficient field of those moving hyperplanes [21]. The relation $g_n \in C_nB_P$ is one exact moving hyperplane relation. It makes the point sequence degenerate over the moving-target field at the outset. The degenerate conclusion is precisely the already known containment in C_nB_P; it does not make the hyperplane independent of n. Thus the hypothesis $h(C_n) = o(h(g_n))$, although it is exactly the "slowly moving" height condition, is insufficient without several new independent incidence forms.

Third, relative-Anosov theory supplies boundary contraction only for a parabolic type in which the representation is divergent [40]. The full Galois representation is not Borel-divergent: outer density gives escaping central elements with bounded outer simple roots. The proposed block representation

$$\operatorname{Sym}^2 \rho_0 \oplus (\rho_1 \otimes \rho_2)$$

is not projective-divergent even on the peripheral sequence T^m: both blocks have top singular value of order m^2, so the first singular gap stays bounded. Extended geometrical-finiteness variants allow more peripheral behavior, but their conclusion is a boundary convergence statement, not recurrence of an arithmetic coset. They therefore do not supply frame locking.

The outcome is a sharper target, not a new name for unit control:

For the exact return cocycle $a_n = C_{n+1}u_nC_n^{-1}$, prove that $h_\Delta(C_n) = o(t_n)$ and double-horoball mismatch force two cosets $C_n\mathscr{B}_P$ to agree.

All three ingredients in that sentence are essential. Dropping the return-cocycle condition makes arbitrary drifting frames possible; dropping the all-place height loses the unit directions; keeping only the Levi factor does not determine the gauge-invariant state coset.

Route II.20.83 (Full Asai coordinates and weighted moving targets do not break the degeneracy)**.** The suggestion at the end of the preceding audit must be qualified: "more independent incidence forms" cannot mean more coordinates of the same transported Asai state.

Let $Y(P) = \sigma_0(P) \otimes \sigma_1(P) \otimes \sigma_2(P)$, and for $C \in \mathrm{SL}_2(K)$ let

$$I_C = \{\ell : \ell(\mathrm{As}(C)Y(P)) = 0\}.$$

Then, for $A \in \mathrm{SL}_2(K)$,

$$\ell(\mathrm{As}(A)Y(P)) = 0 \quad (\ell \in I_C) \quad \Longleftrightarrow \quad A \in CB_P(K).$$

Indeed, simultaneous vanishing says that the two Asai vectors span the same line; equality of their distinguished tensor factors says $[AP] = [CP]$. A basis of I_C has height $O(h_\Delta(C) + 1)$. Thus all seven hyperplanes annihilating the moving Asai line are slowly moving in the equality corridor, but their common zero set on the group is exactly the same moving Borel coset. They make the degeneracy more explicit; they do not make the target fixed.

The first invariant incidence which is genuinely transverse to one moving Borel compares two frames. Put

$$\Delta_{m,n} = \det(Q_m, Q_n) \in \mathcal{O}_K.$$

It vanishes exactly when $C_m \mathscr{B}_P = C_n \mathscr{B}_P$, and its absolute ideal norm is independent of the balancing units. If it is nonzero, then

$$0 \le \log |\operatorname{Nm} \Delta_{m,n}| \le 3h([g_m P]) + 3h([g_n P]) + 3 \log 2.$$

More exactly, with $M_{k,i} = \|\sigma_i(Q_k)\|_\infty$,

$$\sum_{i=0}^{2} \log \frac{M_{m,i} M_{n,i}}{|\sigma_i \Delta_{m,n}|} = 3h([g_m P]) + 3h([g_n P]) - \log |\operatorname{Nm} \Delta_{m,n}|.$$

Therefore an overcritical lower bound for this two-frame proximity forces recurrence. This is the simplest possible target for a Hodge-transverse jet; it also shows that ordinary product-formula height theory lands exactly at the critical boundary.

The strongest weighted moving-target formulation currently located has the same verdict [22]. Its right-hand coefficient contains

$$\max_W \frac{\alpha(W)}{\operatorname{codim} W}, \qquad \alpha(W) = \sum_{W \subset \operatorname{Supp} D_i} c_i.$$

If the central hyperplane is entered twice with total weight c, take W to be that hyperplane: the concentration quotient is at least c. Equivalently, if $D = mH$ is defined by L^m, then

$$\frac{1}{m} \lambda_D(x) = \lambda_H(x) + O(1).$$

The repeated proximity and its critical coefficient grow together. The new theorem also retains the algebraic nondegeneracy hypothesis, which the exact moving-Borel relation violates.

This connects the frame audit to the earlier jet calculations. A useful new incidence must be transverse to the ideal of the current Asai state and must compare at least two consecutive frames. The concrete target is an integral Hodge jet with four properties:

1. it is algebraic over K and has coefficient height $o(t_n)$;
2. it is not in the incidence ideal defining $C_n B_P$, even after restriction to the group orbit;
3. its stopped proximity receives the positive order-seven divisor P_7, producing a strict term rather than a repeated support;
4. its zero-defect locus is a repeated frame coset or an actual parabolic incidence.

The previous nearest-wall jet route had the correct archimedean divisor but did not construct the integral evaluation map with these properties. The present calculation proves that this integral Hodge-transverse jet is not optional if a moving-target theorem is to be the final engine.

Route II.20.84 (The generic transverse jet is the Gauss–Manin Wronskian)**.** The last sentence can now be made more precise. The transverse covariant already exists over the generic fibre. Choose a local Hodge vector ω_i in the i-th rank-two real-multiplication system and a logarithmic tangent vector ∂. Regard the two integral frame vectors as flat covectors $\ell_{m,i}, \ell_{n,i}$, using the polarization, and set

$$W_i^{m,n} = \det \begin{pmatrix} \ell_{m,i}(\omega_i) & \ell_{n,i}(\omega_i) \\ \ell_{m,i}(\nabla_\partial \omega_i) & \ell_{n,i}(\nabla_\partial \omega_i) \end{pmatrix}.$$

Alternating linear algebra gives the exact identity

$$W_i^{m,n} = \sigma_i(\Delta_{m,n}) \big(\omega_i \wedge \nabla_\partial \omega_i\big) = \sigma_i(\Delta_{m,n})\, \vartheta_i.$$

Consequently

$$\frac{W_1^{m,n} W_2^{m,n}}{(W_0^{m,n})^2} = \frac{\mathrm{Nm}(\Delta_{m,n})}{\sigma_0(\Delta_{m,n})^3}\, \Upsilon_0,$$

where $\Upsilon_0 = \theta_1\theta_2/\theta_0^2$ has divisor $3P_7$. The central Higgs field is nowhere zero. Thus, away from the order-seven point, the jet loses rank exactly when $\Delta_{m,n} = 0$, equivalently when the rational frame cosets repeat. The two outer rank losses at the order-seven point have orders two and one; their sum is precisely the Hodge defect, not an unwanted exceptional component.

There is also an exact denominator calculation. For

$$a_k = \frac{3}{4} - \frac{k}{14}, \qquad b_k = \frac{1}{4} - \frac{k}{14},$$

the period equation has coefficient

$$p_k(t) = \frac{1}{t} + \frac{-1 + k/7}{1 - t}$$

in front of y' after normalization. Abel's identity therefore gives

$$\mathcal{W}_k(t) = c_k t^{-1} (1 - t)^{k/7 - 1}.$$

The Galois order is $k = 1, 3, 2$, so

$$\frac{\mathcal{W}_3 \mathcal{W}_2}{\mathcal{W}_1^2} = c(1 - t)^{3/7} = c\, u^3, \qquad u^7 = 1 - t.$$

Thus the determinant of the logarithmic first-jet lattice has fixed algebraic denominator type on the orbifold chart. When the two integral frames are inserted, its only return-dependent finite factor is $(\Delta_{m,n})$, up to the fixed ideal of the chosen de Rham model.

The local analytic inequality can also be proved. Put $R_s = f_{n,s}/f_{m,s}$ and $H_i = R_i \circ F_i \circ R_0^{-1}$. On the central geodesic joining the two rational frame states, $\zeta = R_0(z)$ is purely imaginary. Schwarz–Pick gives

$$\left| \frac{J_i^{m,n}}{J_0^{m,n}}(z) \right| = \left| \frac{\zeta H_i'(\zeta)}{H_i(\zeta)} \right| \le \|DF_i(z)\|_{\mathrm{hyp}} < 1.$$

Thus the first jet has the desired strict local contraction. The remaining obstruction is not local: this rational two-frame geodesic is not known to shadow the stopped orbit segment from t_m to t_n. Equivalently, the rational combination is not yet known to overconverge through the corresponding chain of Schwarz triangles. Multiplying the displayed local inequality along the original orbit without such a tracking theorem would be invalid.

This settles two of the three missing assertions in the earlier arithmetic-holonomy audit: the holonomic functions and their first derivatives have now been identified, and the rank-drop alternative is the desired closing alternative, while the Wronskian determinant has the correct finite ideal. The remaining assertion is metric: show that conformal size minus the fixed denominator type in the Bost slope calculation is the integrated Schwarz–Pick defect

$$\int_{t_m}^{t_n} \bigl(2 - \|DF_1\|_{\mathrm{hyp}} - \|DF_2\|_{\mathrm{hyp}}\bigr)\, ds$$

up to a bounded endpoint term.

The dependence alternative is also completely explicit. If

$$f_{r,i}(t) = \ell_{r,i}(\omega_i(t)),$$

then $f_{m,i}$ and $f_{n,i}$ belong to one fixed rank-two Picard–Fuchs module and

$$\dim_{K(t)} \langle f_{m,i}, f_{n,i} \rangle = 2 \quad \Longleftrightarrow \quad \Delta_{m,n} \neq 0.$$

Write $\mathscr{G}(\varphi)$ for the boundary-energy numerator and $\tau(\mathbf{b})$ for the denominator term in the holonomy dimension bound. For two independent functions that bound reads schematically, but with the same constants as in the cited theorem,

$$2 \leq \frac{\mathscr{G}(\varphi)}{\log|\varphi'(0)| - \tau(\mathbf{b})}.$$

Thus a continuation branch satisfying

$$2\bigl(\log|\varphi'(0)| - \tau(\mathbf{b})\bigr) > \mathscr{G}(\varphi)$$

forces the two frame functions to be dependent and hence forces $\Delta_{m,n} = 0$. The differential system is fixed; the return word changes the integral sections and the continuation branch. The exact remaining task is to construct those three Galois continuation branches so that the surplus in this inequality is the integrated Hodge defect above, up to bounded endpoint terms.

A long return word alone does not produce this surplus. It is merely monodromy of the same system on the same hyperbolic universal cover, so the unmodified holonomy inequality only permits the known rank two. A valid application must prove that the stopped rational linear combination cancels the first analytic obstruction and overconverges through a larger domain. Showing that this added conformal radius is the integrated Hodge defect is the substantive assertion; choosing a word-dependent symbol $\varphi_{m,n}$ without proving overconvergence would be circular.

If this holds, the slope inequality would give

$$\mathcal{P}_{m,n} \geq 3h([g_m P]) + 3h([g_n P]) + \int_{t_m}^{t_n} \bigl(2 - \|DF_1\|_{\mathrm{hyp}} - \|DF_2\|_{\mathrm{hyp}}\bigr)\, ds - O_P(1).$$

On a ray with infinitely many compact returns the integral is unbounded. For $\Delta_{m,n} \neq 0$, however, the exact product formula gives the opposite critical upper bound

$$\mathcal{P}_{m,n} \leq 3h([g_m P]) + 3h([g_n P]).$$

Hence the two estimates force $\Delta_{m,n} = 0$ for some pair, which closes the rational Borel frame and proves the cusp/special alternative.

The arithmetic holonomy theorem of Calegari–Dimitrov–Tang [57] is relevant here as a slope architecture, not as a fixed-target irrationality estimate. It already treats holonomic functions and their derivatives, and it places denominator loss and conformal gain in one all-place inequality. What must be supplied from the double-heptagon geometry is exactly this stopped conformal-size comparison. No calculation with a finite list of group generators can verify it.

Route II.20.85 (Natural-extension repair returns the affine Borel cocycle)**.** One possible repair of the path mismatch is to use the true past flag. Write an initial unimodular frame as $H_P = (P, E)$, and for a reduced return put

$$Q_n = \varepsilon_n g_n P, \qquad R_n^{\mathrm{dyn}} = \varepsilon_n^{-1} g_n E.$$

This pair has determinant one and is the future–past frame of the actual natural-extension geodesic. Its all-place transverse proximity is exactly

$$3h([g_nP]) + 3h([g_nE]).$$

Thus it removes the geometric path mismatch, but introduces the second projective height.

Let R_n^{ctl} be the height-reduced second column used in the controlled Borel frame. There is a unique $s_n \in \mathcal{O}_K$ with

$$R_n^{\mathrm{dyn}} = R_n^{\mathrm{ctl}} + s_n Q_n.$$

If the controlled transition is

$$\begin{pmatrix} d_n & \beta_n \\ 0 & d_n^{-1} \end{pmatrix}, \qquad d_n \in \mathcal{O}_K^{\times},$$

then direct transport of the two second columns gives

$$s_{n+1} = d_n^2 s_n + d_n \beta_n.$$

The true past flag therefore restores exactly the additive variable which is removed when one passes to the rational Borel coset. This also pinpoints the limitation of the S-unit equation: d_n is multiplicative, while s_n and β_n range through the additive lattice. Once those latter terms have bounded multiplicative support, the standard S-unit theorem closes the orbit; the affine recurrence does not itself supply that boundedness.

The affine variable has an intrinsic height interpretation. If $M_n = (Q_n, R_n^{\mathrm{ctl}})$ is the controlled determinant-one frame, then $[R_n^{\mathrm{dyn}}] = M_n[s_n : 1]$. Applying the standard projective height estimate to M_n and M_n^{-1}, and using the controlled-completion bound, gives

$$h([s_n : 1]) = h([g_nE]) + O(h([g_nP]) + 1).$$

Thus the shear height is the height of the true past cusp modulo the already measured future-state height. If the sum of these two heights is bounded along infinitely many returns, Northcott repeats the ordered future–past pair and its quotient gives the required stabilizer.

The recurrence is also the coordinate form of a canonical rank-two bundle. On the projective flat bundle let $\mathscr{L} \subset V$ be the tautological line. Tensoring the tautological sequence by $\mathscr{L}$ gives

$$0 \longrightarrow \mathscr{L}^2 \longrightarrow \mathscr{L} \otimes V \longrightarrow \det V \simeq \mathcal{O} \longrightarrow 0.$$

A determinant-one second column is a splitting of the last arrow. For $b = \left(\begin{smallmatrix} d & \beta \\ 0 & d^{-1} \end{smallmatrix}\right)$, the action on the splitting coordinate is

$$db = \begin{pmatrix} d^2 & d\beta \\ 0 & 1 \end{pmatrix},$$

which sends $(s_n, 1)^t$ to $(s_{n+1}, 1)^t$. Thus the integral affine extension does not have to be invented: it is the universal tautological extension. The remaining task is to put the cyclic Wronskian/Hodge metric on its Galois completion and prove the stopped degree inequality.

The outcome is a genuine no-shortcut statement. The dynamic frame has the correct path and an uncontrolled endpoint height. The controlled frame has sublinear height and the wrong past endpoint. Arithmetic holonomy or a stopped Arakelov slope must amortize the affine endpoint cost; it cannot be deleted by choosing a different completion.

Route II.20.86 (The apparent eigenframe singularity cancels against the fixed-point discriminant)**.** The symmetric-power route in `curt12.tex` left an "eigenframe comparison" as part of its archimedean conjecture. This piece can be separated from the genuinely arithmetic difficulty.

For a hyperbolic real matrix A, choose determinant-one eigenvectors e_+, e_-, with eigenvalues $e^{\ell/2}, e^{-\ell/2}$. In the corresponding coordinates,

$$Q_A(xe_+ + ye_-) = \det(v, Av) = -2\sinh(\ell/2)xy.$$

Thus a nearly singular eigenframe comes with a small scalar $2\sinh(\ell/2)$ in the unnormalized integral fixed-point form. The two effects must be estimated together.

More precisely, with Bombieri norms and $n = d - 2r$, $k = n + 1$, multiplication by $(XY)^r$ has determinant entropy

$$\mathcal{E}(d, r) = \frac{1}{2}\sum_{j=0}^{n}\left(\log\binom{d}{r+j} - \log\binom{n}{j}\right).$$

If $B \in \mathrm{SL}_2(\mathbb{R})$ has singular values $e^{\kappa}, e^{-\kappa}$, conjugating the critical multiplication map by B changes its k-dimensional logarithmic Jacobian by at most $kd\kappa$. This follows because $\det \mathrm{Sym}^n(B) = 1$ and every singular value of $\mathrm{Sym}^d(B)$ lies between $e^{-d\kappa}$ and $e^{d\kappa}$.

For the balanced eigenframe, $\kappa = R/2$, where R is the distance from the Hodge basepoint to the axis. The exact formula

$$\sinh(D/2) = \cosh R \, \sinh(\ell/2)$$

then couples this basis-change cost with the scalar in Q_A. Hence the apparent blow-up when $\ell \to 0$ is not a separate obstruction; it was created by normalizing the eigenline factors and discarding their scalar separation. Backtracking is also favorable for an endpoint metric because $D \le L$.

This repairs a real loose end but does not prove the classification. The nearest-wall construction replaces the sum of the two outer weights by their maximum using a quotient of Hodge lines. That quotient is curved but is not an independent integral flat line. Its integral realization has to retain the affine Borel framing, and doing so restores the cocycle

$$s_{n+1} = d_n^2 s_n + d_n \beta_n.$$

The symmetric-power and Wronskian approaches therefore meet at the same remaining task: an Arakelov slope theorem for this affine extension.

Route II.20.87 (Actual-path contraction and the exact affine endpoint)**.** Using the true future–past frame does more than name the path mismatch. Let (P, E) be the two boundary endpoints of the oriented geodesic and put $R_i = f_{E,i}/f_{P,i}$. On the central geodesic, $d\log|R_0|$ is hyperbolic arclength. The two-frame Schwarz estimate therefore gives the unconditional differential inequality

$$\frac{d}{dt}\left(2\log|R_0| - \log|R_1| - \log|R_2|\right) \ge 2 - \|DF_1\|_{\mathrm{hyp}} - \|DF_2\|_{\mathrm{hyp}}.$$

Thus the full Hodge defect accumulates on the original stopped orbit; no shadowing of the geodesic between two future tails is needed.

For two future states with $\Delta_{m,n} = \det(Q_m, Q_n) \neq 0$, the rational vector

$$\widetilde{R}_{m,n} = Q_n/\Delta_{m,n}$$

is a determinant-one completion of Q_m. It differs from the dynamic past by one rational shear:

$$R_m^{\mathrm{dyn}} = \widetilde{R}_{m,n} + s_{m,n}Q_m, \qquad s_{m,n} = -\frac{A_{m,n}}{\Delta_{m,n}}, \quad A_{m,n} = \det(Q_n, R_m^{\mathrm{dyn}}).$$

Scaling the second column cancels between a Wronskian numerator and its period denominator, so

$$J_i(Q_m, \widetilde{R}_{m,n}) = J_i^{m,n}.$$

The exact discrepancy between the accumulated dynamic metric and the two-future integral Wronskian is consequently the single boundary factor

$$\log \left| \frac{f_{\widetilde{R},1} f_{\widetilde{R},2} f_{\mathrm{dyn},0}^2}{f_{\mathrm{dyn},1} f_{\mathrm{dyn},2} f_{\widetilde{R},0}^2} \right|.$$

Its denominator ideal is already $(\Delta_{m,n})$; there is no hidden path error or additional denominator.

The two coordinates are themselves a full thin-group invariant. If $k_{m,n} = g_m^{-1} g_n$ and $H_P = (P, E)$, write

$$H_P^{-1} k_{m,n} P = (a_{m,n}, b_{m,n})^t.$$

Then

$$b_{m,n} = \det(P, k_{m,n}P) = Q_{k_{m,n}}(P), \qquad \Delta_{m,n} = \varepsilon_m \varepsilon_n b_{m,n},$$

and $A_{m,n} = (\varepsilon_n/\varepsilon_m) a_{m,n}$. Hence the old fixed-point quadratic and the new Wronskian minor are the same invariant, applied to the quotient of two prefixes. The residual slopes lie in the pointed orbit

$$H_P^{-1} G_7 H_P \cdot [1:0],$$

not in the full ambient arithmetic projective line.

This changes the description of the remaining theorem. The path contraction is proved. What remains is an all-place endpoint-absorption statement: the signed Hodge norm of $-A_{m,n}/\Delta_{m,n}$ cannot absorb every accumulated defect unless one incidence coordinate vanishes or the affine frame recurs. Such a statement is false for the ambient Hilbert modular group, which realizes every primitive coordinate pair. It must use the ordered pointed orbit of the thin triangle group.

Route II.20.88 (Cusp normalization, the cross-ratio wall, and the correct quantitative target)**.** The previous route silently used two incompatible choices of the second column. An integral unimodular completion is convenient for Borel-height reduction, but its projective line need not be a cusp. For the actual geodesic choose instead a primitive cusp column C, put

$$\delta_P = \det(P, C), \qquad E = C/\delta_P,$$

and work in the fixed fractional lattice $\delta_P^{-1}\mathcal{O}_K^2$. Then $\det(P, E) = 1$, all height changes are $O_P(1)$, and $A_{m,n} = 0$ really is the cusp incidence. This corrects the normalization without changing the varying denominator $(\Delta_{m,n})$.

The endpoint factor has an exact invariant form. If $Z_i(t)$ is the period point and $U = R_m^{\mathrm{dyn}}$, $V = \widetilde{R}_{m,n}$, then

$$\mathrm{CR}_i(m,n) = \frac{[Z_i(t_n); \sigma_i U, \sigma_i V]}{[Z_i(t_m); \sigma_i U, \sigma_i V]}, \qquad [Z; U, V] = \frac{\det(Z, U)}{\det(Z, V)}.$$

The whole boundary loss is

$$2\log|\,\mathrm{CR}_0\,| - \log|\,\mathrm{CR}_1\,| - \log|\,\mathrm{CR}_2\,|.$$

After applying g_m^{-1}, the two algebraic endpoint lines are exactly $[E]$ and $[k_{m,n}P]$. Thus the remaining statement is a Galois-weighted cross-ratio theorem for consecutive quotients in the pointed orbit $\widehat{H}_P^{-1} G_7 \widehat{H}_P[1:0]$.

There is also an exact central diagnostic. Send E to 0, P to ∞, write $k_{m,n}P = [a:b]$, and set $v = \sigma_0(a/b)$. Along $z(t) = ie^t$,

$$\frac{d}{dt}\log\left|\frac{z(t)}{z(t)-v}\right| = \frac{v^2}{e^{2t}+v^2},$$

so the alignment deficit on $[t_m, t_n]$ is

$$\frac{1}{2}\log\frac{e^{2t_n}+v^2}{e^{2t_m}+v^2}.$$

Near-total central absorption therefore forces $|\sigma_0(a/b)| \geq e^{t_n - O_P(1)}$: the pointed orbit returns exponentially close to P in the visual metric. The outer two factors must satisfy the corresponding signed Hodge alignments simultaneously. This is the precise wall which ordinary approximation, ambient height, and the affine recurrence each see only partially.

One overstrong formulation has been discarded. It is unnecessary to prove that the endpoint factor is bounded. It is enough to find $\eta_P > 0$ such that

$$\mathfrak{B}_{m,n}(t_n) - \mathfrak{B}_{m,n}(t_m) \leq (1-\eta_P)\mathfrak{C}(m,n) + O_P(1).$$

The surviving fraction $\eta_P \mathfrak{C}(m,n)$ already contradicts the product-formula upper bound for a nonzero Wronskian minor. This weaker target is both sufficient and more plausible.

Schwarz–Pick by itself does not prove it. For $\psi_{u,v}(z) = \log|(z-u)/(z-v)|$,

$$\|d\psi_{u,v}(z)\|_{\mathrm{hyp}} = \frac{\mathrm{Im}(z)|u-v|}{|z-u|\,|z-v|} \leq 1.$$

Applied separately at the three embeddings this produces the wrong outer signs: it bounds the endpoint derivative by $2+\|DF_1\|+\|DF_2\|$, whereas the useful curvature is $2-\|DF_1\|-\|DF_2\|$. The missing theorem must recover those signs from the common algebraic pair (a,b) and its consecutive thin-orbit origin.

The sign problem has an exact decomposition. Let

$$L_0 = t_n - t_m, \qquad L_i = \int_{t_m}^{t_n} \|DF_i\|_{\mathrm{hyp}}\,dt, \qquad X_i = \log|\,\mathrm{CR}_i\,|, \qquad \alpha_i = L_i - X_i.$$

The Lipschitz inequality gives $\alpha_i \geq 0$, and direct substitution gives

$$\mathfrak{B}_{m,n}(t_n) - \mathfrak{B}_{m,n}(t_m) = \mathfrak{C}(m,n) - 2\alpha_0 + \alpha_1 + \alpha_2.$$

Thus the desired endpoint theorem is equivalent to

$$2\alpha_0 - \alpha_1 - \alpha_2 \geq \eta_P \mathfrak{C}(m,n) - O_P(1).$$

Here α_0 is the explicit central visual misalignment computed above, while α_1, α_2 are the backtracking of the two Hodge images relative to their Galois-conjugate endpoint geodesics. The precise thin-group theorem must say that the two outer paths cannot backtrack more efficiently than the central pointed orbit approaches recurrence. This is false for unrelated paths and is the reason repeated applications of scalar Schwarz–Pick return to the same impasse.

The equality calculation is complete. If a quotient fixes P and is $\left(\begin{smallmatrix} d & \beta \\ 0 & d^{-1} \end{smallmatrix}\right)$ in the frame (P, E), then its affine action $s \mapsto d^2 s + d\beta$ has fixed point $s_* = -d\beta/(d^2-1)$ when $d^2 \neq 1$. The line $[E + s_* P]$ is exactly the other eigenline. When $d^2 = 1$ and the projective element is nonidentity, it is parabolic. Thus the two neutral outcomes of the affine endpoint theorem are already exactly the special and cusp alternatives.

Status. The normalization error is repaired; actual-path contraction, the rational shear, the cross-ratio identity, the central visual formula, and the affine equality classification are proved. The single remaining assertion is the strict simultaneous Galois non-absorption estimate above. A targeted numerical test, if pursued, should measure the ratio

$$\frac{\mathfrak{B}_{m,n}(t_n) - \mathfrak{B}_{m,n}(t_m)}{\mathfrak{C}(m,n)}$$

on long exact stopped intervals and stratify by the visual sizes of $\sigma_i(a/b)$. Scalar remainder energies do not test this statement.

Route II.20.89 (Two-channel Hodge dichotomy and rank-two Arakelov fusion)**.** The preceding route formulated the residual problem as the scalar inequality

$$2\alpha_0 - \alpha_1 - \alpha_2 \geq \eta \mathfrak{C} - O(1).$$

That formulation is sufficient but asks for more sign control than the geometry naturally supplies. A second calculation shows that the outer backtracking defects do not need to be bounded separately.

Keep the dynamic pair $(Q, U) = (Q_m, R_m^{\mathrm{dyn}})$, and put $V = Q_n/\Delta_{m,n}$. Define

$$\mathfrak{A}^{\mathrm{ctl}}(t) = 2\log\left|\frac{f_{V,0}(t)}{f_{Q,0}(t)}\right| - \sum_{i=1}^{2} \log\left|\frac{f_{V,i}(t)}{f_{Q,i}(t)}\right|.$$

In the distinguished coordinate $Q = \infty$, $U = 0$, $V = v$, and $z(t) = ie^t$, one has

$$[\log|f_{V,0}/f_{Q,0}|]_{t_m}^{t_n} = L_0 - X_0 = \alpha_0.$$

At either outer embedding, the cross-ratio Lipschitz estimate for the endpoint pair $([Q], [V])$ bounds the absolute change by L_i. Therefore

$$\mathfrak{A}^{\mathrm{ctl}}(t_n) - \mathfrak{A}^{\mathrm{ctl}}(t_m) \geq 2\alpha_0 - L_1 - L_2 = \mathfrak{C}(m,n) - 2X_0.$$

It follows formally that

$$\max\{X_0, \Delta\mathfrak{A}^{\mathrm{ctl}}\} \geq \frac{1}{3}\mathfrak{C}(m,n).$$

This is proved, not conjectural. It says that geometry always puts a fixed fraction of the defect into one of two coordinates of the rational two-future frame.

The first coordinate is also an exact Green quantity. With $G(t) = \Lambda(Z_0(t); [Q_m], [Q_n])$, projective invariance and

$$\Lambda(ie^t; \infty, v) = \frac{1}{2}\log(1 + v^2 e^{-2t}) = -u_v(t)$$

give

$$X_0 = G(t_m) - G(t_n) \le G(t_m).$$

Thus the alternatives have a simple geometric meaning. Either the rational endpoint line shadows the true forward line and the two-frame Green proximity drops substantially, or it does not and the controlled Wronskian retains the Hodge curvature. Outer backtracking cannot erase both.

This calculation changes the precise location of the remaining gap. It is no longer necessary to prove the scalar alignment inequality above. The finite denominator of V is exactly $(\Delta_{m,n})$, and the Green and controlled-Hodge quantities are two coordinates of the same tautological affine extension

$$0 \longrightarrow \mathscr{L}^2 \longrightarrow \mathscr{L} \otimes V \longrightarrow \mathcal{O} \longrightarrow 0.$$

The required arithmetic theorem must put the three Galois copies of this rank-two extension into one adelic lattice, enter $(\Delta_{m,n})$ once, and let maximal slope choose the larger of the two channels. A convenient exact target is

$$\log|N\Delta_{m,n}| + \eta \max\{X_0, \Delta\mathfrak{A}^{\mathrm{ctl}}\} \le O_P(1)$$

on stopped intervals before affine-frame recurrence, apart from $A_{m,n}\Delta_{m,n} = 0$. Together with the one-third inequality this immediately bounds the number of compact returns and closes the orbit.

There is an additional algebraic realization of the Green channel. Put $L = K(\eta)$, where $\eta + \eta^{-1} = \lambda$, and for a column $q = (q_1, q_2)^t$ write $Z_q = q_1 - q_2\eta$. For two columns define

$$t(q, v) = \frac{Z_q \overline{Z_v}}{\overline{Z_q} Z_v}.$$

This point has relative norm one and absolute value one at every complex embedding. Direct expansion gives the exact recurrence factor

$$t(q, v) - 1 = \frac{(\eta - \eta^{-1})\det(q, v)}{\overline{Z_q} Z_v}.$$

If η_i denotes the corresponding upper-half-plane order-seven point at the i-th real embedding, then

$$\Lambda(\eta_i; [q_i], [v_i]) = \log\frac{2}{|1 - t_i|}.$$

Summing gives

$$\sum_i \Lambda(\eta_i; [q_i], [v_i]) = 3\log 2 - \frac{1}{2}\log 7 - \log|N\det(q, v)| + \frac{1}{2}\log|NQ_7(q)| + \frac{1}{2}\log|NQ_7(v)|.$$

Thus the full three-place Green coordinate is one norm-one CM-torus coordinate, and the determinant divisor occurs exactly once. On a compact return set this differs from the Green coordinate at the actual period basepoints by $O(1)$. This is a useful synthesis of the order-seven Hodge divisor with the two-future frame: the first channel is now completely algebraic, not just visual.

It also gives a clean Roth audit. Recurrence is $t = 1$, but both the target and the moving $t(q, v)$ lie in the same number field. Unless the moving torus points are first confined to a fixed

finitely generated multiplicative subgroup, ordinary Roth or the standard S-unit theorem has no strict exponent to use; the displayed product formula is exactly critical. Hence the CM conversion identifies the right global variable but does not by itself prove the stopped fusion.

The controlled channel has an equally concrete Hodge realization. For a rational flag x, put

$$\Psi_x(z) = \frac{F_1'(z)(z-x_0)}{F_1(z)-x_1}\frac{F_2'(z)(z-x_0)}{F_2(z)-x_2}.$$

For two flags x, y, the derivatives cancel and

$$\frac{\Psi_y}{\Psi_x} = \left(\frac{z-y_0}{z-x_0}\right)^2 \prod_{i=1}^{2} \frac{F_i(z)-x_i}{F_i(z)-y_i}.$$

The logarithm of its modulus is exactly the controlled cyclic potential. Every Ψ_x has the same zero divisor $3P_7$, independently of x. Hence the quotient is holomorphic and nowhere zero on the upper half-plane: it contains no second Hodge divisor. This proves a second part of the bookkeeping which was previously only heuristic. Separate scalar estimates would duplicate not only the determinant ideal but also the order-seven Hodge charge. The positive curvature belongs to the rank-two pencil $\langle \Psi_x, \Psi_y \rangle$; after quotienting, the controlled coordinate is harmonic.

The obvious slope shortcut was also checked. The ordinary module

$$M_{m,n} = \mathcal{O}_K Q_m + \mathcal{O}_K Q_n \subset \mathcal{O}_K^2$$

has determinant ideal $(\Delta_{m,n})$, but its determinant degree sees only the symmetric sum of the three Green functions. It does not see the cyclic coordinate $2r_0 - r_1 - r_2$. If instead one takes the projective line of $Q_n/\Delta_{m,n}$ and saturates it before applying the usual maximal-slope theorem, clearing the scalar $\Delta_{m,n}$ returns the line of Q_n and erases the affine record. Thus the standard two-column slope inequality keeps either the determinant or the cyclic framing, not both. This attempt fails before any thin-group estimate is used and confirms that the required object is genuinely the framed Galois-completed extension.

The integral extension itself can nevertheless be constructed without a conjecture. Determinant with Q_m gives

$$0 \longrightarrow \mathcal{O}_K Q_m \longrightarrow \mathcal{O}_K Q_m + \mathcal{O}_K Q_n \longrightarrow (\Delta_{m,n}) \longrightarrow 0.$$

Tensoring by $\mathcal{O}_K Q_m$ gives the integral model of the tautological affine extension. After dividing its one quotient ideal by $\Delta_{m,n}$, the distinguished rational lift is exactly

$$Q_m \otimes \frac{Q_n}{\Delta_{m,n}}.$$

Thus the lattice, filtration, rational splitting, and single finite divisor are all explicit. What remains is solely the metric assertion that the maximal slope of its Galois-completed Hodge metric captures the larger channel on consecutive thin-orbit intervals. This construction rules out another possible ambiguity: "rank-two fusion" no longer means that an integral extension still has to be guessed.

The maximal-slope coordinate can also be written exactly. Every rational splitting line is generated by

$$e(s) = Q_m \otimes \left(\frac{Q_n}{\Delta_{m,n}} + sQ_m\right), \qquad s \in K.$$

At the i-th real place put $\zeta_i(t) = -f_{Q_n/\Delta,i}(t)/f_{Q_m,i}(t)$. The Hodge norm of $e(s)$ is a fixed factor times

$$|\sigma_i(s) - \zeta_i(t)|.$$

Relative to the integral basis $Q_m \otimes Q_m, Q_m \otimes Q_n$, its finite content is the one fractional ideal

$$(s, \Delta_{m,n}^{-1}).$$

Maximal slope is therefore literally simultaneous approximation of the three moving Hodge splitting targets by one $s \in K$, with the determinant denominator inserted once. For the dynamically correct splitting $s = -A_{m,n}/\Delta_{m,n}$, this content is exactly $\Delta_{m,n}^{-1}$ up to the fixed cusp-normalization ideal, and $e(s) = Q_m \otimes R_m^{\mathrm{dyn}}$. The finite and Hodge data are carried by the same global line.

This identifies the strongest relevant Roth statement. Ordinary adelic Roth would address fixed algebraic targets at the critical exponent two. Here the three targets are moving period ratios of one Picard–Fuchs system, and their motion is the positive Hodge defect. What is needed is a relative holonomic Roth theorem that charges this motion while retaining the displayed content ideal. A better numerical constant in the classical fixed-target theorem does not address this problem.

Failed shortcut. Adding two separately normalized scalar heights is invalid. Both use the same rational splitting $Q_n/\Delta_{m,n}$, so the sum inserts the divisor $(\Delta_{m,n})$ twice. This recreates the critical multiplicity found in the weighted moving-target experiment and manufactures a false surplus. The maximum must be taken only after the two coordinates have been assembled in one rank-two Arakelov object.

What this route accomplishes. It proves an analytic one-third survival theorem and replaces the difficult comparison of three oriented cross-ratios by one precise integral-fusion problem. It does not yet prove that fusion inequality. The remaining theorem is now about the maximal slope of an explicit rank-two lattice along the consecutive pointed thin orbit; it is not a statement about individual prime factors, arbitrary rational approximation, or isolated Rosen generators.

Route II.20.90 (Intermediate-group audit: the natural unit enlargement is already arithmetic)**.** The proposal to interpolate between the triangle group and the Hilbert modular group was tested at three different levels.

First, there is no larger discrete Fuchsian group containing G_7. Indeed a discrete overgroup has finite index d, and its noncompact orbifold has Euler characteristic $-5/(14d)$. The lower bound $-\chi \geq 1/6$ for a noncompact orientable hyperbolic orbifold forces $d \leq 2$. Index two would require a signature with $-\chi = 5/28$; the only numerical possibility would be $(0; m, n; 1)$ with $1/m + 1/n = 23/28$, which has no integral solution. Thus a proper intermediate group is necessarily nondiscrete at the distinguished real embedding. Its closure is all of $\mathrm{PSL}_2(\mathbb{R})$: the Lie algebra of the identity component is nonzero and invariant under the Zariski-dense adjoint action of G_7, hence is all of $\mathfrak{sl}_2$. In particular, a proper intermediate group no longer acts through a quotient hyperbolic orbifold on which the Rosen stopping rule can be read.

Second, the exact finite-adelic closure does give one canonical comparison group

$$\Gamma_{\mathrm{cong}}(G_7) = \mathrm{SL}_2(\mathcal{O}_K) \cap \widehat{G}_7.$$

Calegari's calculation gives index 144 in $\mathrm{SL}_2(\mathcal{O}_K)$ for $n = 7$ [56, Theorem 1.6]. The integral model in that paper is conjugate to the present one by $\mathrm{diag}(\lambda, 1)$, which scales the two parabolic parameters from $\lambda, -\lambda$ to $\lambda^2, -1$; this is an integral conjugation because λ is a unit. The opposite parabolics generate the full group: writing $r = ST$, their product in one order is r^2, and the order-seven element r is recovered because 2 is invertible modulo 7. The closures of G_7 and its congruence hull are identical at every finite place. Yet every K-rational line has a nontrivial unipotent stabilizer in the hull: clear denominators in a rank-one nilpotent with that kernel and then take a power lying in the finite-index hull. This makes the comparison useful diagnostically but not as a descent theorem. A hull stabilizer is automatic and need not be a G_7-stabilizer.

The most natural attempt to promote that stabilizer to an intermediate group is also forced to be arithmetic. Let U_P be the unipotent K-group fixing P, let $\Lambda_P = U_P(K) \cap \Gamma_{\mathrm{cong}}(G_7)$, and set

$$H_P = \langle G_7, \Lambda_P \rangle.$$

In the threefold Minkowski embedding, H_P is a discrete Zariski-dense subgroup of $\mathrm{SL}_2(\mathbb{R})^3$. Moreover Λ_P is commensurable with a fractional ideal of K, so it is an irreducible lattice in the three-dimensional horospherical group $U_P(\mathbb{R})^3$. Benoist–Miquel's arithmeticity theorem therefore makes H_P an arithmetic lattice [9]. Since $H_P \subset \mathrm{SL}_2(\mathcal{O}_K)$, it has finite index. This gives a general form of the threshold which the diagonal-unit calculation below exhibits explicitly: adding the complete rank-three cusp stabilizer destroys thinness. Adding only one cyclic unipotent does not meet the theorem's hypothesis, but then the distinguished projection is already dense and there is neither a Fuchsian quotient nor a Rosen stopping rule. Thus it supplies no canonical middle geometry. There is not even a canonical rank-two additive compromise at the standard cusp: since $\mathrm{Tr}_{K/\mathbb{Q}}(\lambda) = 1$, one has

$$\mathcal{O}_K = \lambda\mathbb{Z} \oplus \ker(\mathrm{Tr}_{K/\mathbb{Q}}).$$

The triangle group already contains $T(\lambda\mathbb{Z})$, so adjoining the trace-zero rank-two translation lattice produces the full cusp lattice and is arithmetic by the preceding theorem.

Third, the most natural attempt to make the neutral unit direction into an actual group symmetry collapses the interpolation completely. Put

$$D = \mathrm{diag}(\lambda, \lambda^{-1}).$$

Since $\lambda^{-1} = -\lambda^2 + \lambda + 2$, this is an integral matrix. Conjugation gives

$$D^n T(\lambda) D^{-n} = T(\lambda^{2n+1}).$$

Exact reduction by $X^3 - X^2 - 2X + 1$ gives

$$\lambda = (0, 1, 0), \quad \lambda^3 = (-1, 2, 1), \quad \lambda^5 = (-3, 5, 4)$$

in the basis $1, \lambda, \lambda^2$. The determinant is 1, so these three parameters form an integral basis. Hence $\langle G_7, D \rangle$ contains all upper and lower elementary matrices. Vaserstein's elementary-generation theorem then gives

$$\langle G_7, D \rangle = \mathrm{SL}_2(\mathcal{O}_K)$$

[61]. This exact computation succeeded; it is not a heuristic density test. It proves that the unit which would remove the shape obstruction also removes thinness.

The remaining canonical rank-two cusp candidate can also be eliminated exactly. A bidirectional search in the matrix group over $\mathbb{Z}[\lambda]/(\lambda^3 - \lambda^2 - 2\lambda + 1)$, using the six letters $S^{\pm 1}, T(1)^{\pm 1}, T(\lambda)^{\pm 1}$, found the length-fourteen identity

$$\begin{aligned} T(-1)ST(-\lambda)ST(1)T(-\lambda)S^{-1}T(-\lambda) \\ \cdot T(-1)S^{-1}T(\lambda)T(-1)S^{-1}T(1) = \mathrm{diag}(\lambda, 2 + \lambda - \lambda^2) = D. \end{aligned}$$

The program `intermediate_unit_word_search.py` verifies the identity using exact coefficient triples. An independent Mathematica calculation reduced all four entries modulo the cubic polynomial and returned $\begin{pmatrix} \lambda & 0 \\ 0 & 2+\lambda-\lambda^2 \end{pmatrix}$; it also returned $\lambda(2 + \lambda - \lambda^2) - 1 = 0$. Consequently

$$\langle G_7, T(1) \rangle = \mathrm{SL}_2(\mathcal{O}_K).$$

This is stronger than a density or finite-quotient test. It proves that raising the standard cusp translation module from $\lambda\mathbb{Z}$ to the canonical rank-two module $\mathbb{Z}+\lambda\mathbb{Z}$ already generates the diagonal unit and hence the full arithmetic group.

There is one apparently rank-two additive enlargement which the first audit missed:

$$L_2 = 2\mathbb{Z} + \lambda\mathbb{Z}, \qquad H_2 = \langle G_7, T(2)\rangle.$$

It is strictly larger than G_7, because the latter's standard cusp stabilizer is $T(\lambda\mathbb{Z})$. It is strictly smaller than $\mathrm{SL}_2(\mathcal{O}_K)$, because $T(2)$ is trivial modulo two: the image remains the dihedral group D_7 rather than $\mathrm{SL}_2(\mathbb{F}_8)$. Its distinguished real projection is nevertheless dense, since $2\mathbb{Z}+\lambda\mathbb{Z}$ is dense in $\mathbb{R}$, and S conjugates the resulting upper one-parameter group to the lower one.

The wall calculation below decides its index: H_2 is finite index. The two visible translations do not remain the whole standard cusp intersection. A split wall shares a rational eigenline with a special element of G_7; their commutator creates a parabolic on that line, and conjugation by the special element expands its one parameter into a full cubic lattice. Thus H_2 is a proper arithmetic subgroup, not a thin intermediate group.

Calegari's tangent-space description gives the congruence closure of H_2 index 72. In the conjugate model $T(2)$ is $T(2\lambda)$; modulo four its tangent vector adds one of the two scalar directions absent from the triangle-group image, while its mod-two image is unchanged and the full level-four kernel is already present. This finite calculation explains properness and the canonical halfway congruence position; the commutator calculation, not the finite quotients, is what proves finite index.

Proposition II.20.91 (The split cross-forms are the H_2 wall square)**.** *Set $L = T(\lambda)$, $E = T(2)$, and retain the four forms $Q_{AC}, Q_{AD}, Q_{BC}, Q_{BD}$ of Proposition II.22.20. The exact identities*

$$Q_{LSE^{-1}SL} = -2Q_{AC}, \qquad Q_{SLSE^{-1}SLS} = -2\alpha^2 Q_{BD},$$
$$Q_{E^{-1}SE^{-1}L^3SE^{-1}S} = (3-6\lambda)Q_{AD}, \qquad Q_{ESEL^{-3}SES} = (3-6\lambda)Q_{BC}$$

hold. The four fixed-line pairs are

$$(e_A, e_C), \quad (e_B, e_D), \quad (e_A, e_D), \quad (e_B, e_C);$$

they are the four edges of the square on e_A, e_B, e_C, e_D. Opposite edges have eigenvalues

$$(\lambda\alpha^2)^{\pm 1} \quad \text{or} \quad (\lambda^{-1}\alpha^4)^{\pm 1},$$

up to simultaneous signs. The exponent vectors $(1,2)$ and $(-1,4)$ span an index-six sublattice of $\mathcal{O}_K^\times/\{\pm 1\} = \langle \lambda, \alpha\rangle$.

All four elements lie in $H_2 \setminus G_7$. More precisely, modulo four the images of G_7 and H_2 have orders 1792 and 3584, and the four words represent the nontrivial coset. Since H_4 has the same modulo-four image as G_7, none of these walls belongs to H_4.

Proof. Exact multiplication gives, for example,

$$LSE^{-1}SL = \begin{pmatrix} -1-2\lambda & -2\alpha^2 \\ -2 & -1-2\lambda \end{pmatrix}$$

and

$$E^{-1}SE^{-1}L^3SE^{-1}S = \begin{pmatrix} -4+12\lambda & -3+6\lambda \\ 3-6\lambda & 2-3\lambda \end{pmatrix}.$$

Substitution in $Q_M(X,Y) = cX^2+(d-a)XY-bY^2$ proves the first and third identities; conjugation by S and the displayed opposite word give the other two. The eigenvalue identities reduce to

$$1+2\lambda+2\alpha = \lambda\alpha^2, \qquad 2-3\lambda-(3-6\lambda)\alpha = \lambda^{-1}\alpha^4.$$

The determinant of the two exponent vectors is six.

Finally the level-four image of H_2 is the index-two extension of the G_7-image generated by E. Each displayed word contains an odd number of $E^{\pm1}$ letters. The exact program `intermediate_cross_form_search.py` checks all matrix and fixed-form identities, enumerates the two finite images, and verifies the four coset memberships. □

Proposition II.20.92 (The wall square arithmeticizes H_2)**.** *The proper subgroup*

$$H_2 = \langle G_7, T(2)\rangle \subsetneq \mathrm{SL}_2(\mathcal{O}_K)$$

has finite index. Consequently its intersection with the standard upper-unipotent group has rank three and strictly contains $T(2\mathbb{Z}+\lambda\mathbb{Z})$.

Proof. Let $v_C = (\alpha,1)^t$ and

$$N_{v_C} = \begin{pmatrix} -\alpha & \alpha^2 \\ -1 & \alpha \end{pmatrix}.$$

The special element

$$W = \begin{pmatrix} 1+\lambda^2 & \lambda \\ \lambda & 1 \end{pmatrix} \in G_7$$

and the positive representative

$$h = -LSE^{-1}SL = \begin{pmatrix} 1+2\lambda & 2\alpha^2 \\ 2 & 1+2\lambda \end{pmatrix} \in H_2$$

both fix Kv_C. Exact multiplication gives the nontrivial parabolic

$$[W,h] = I + 2\lambda^3\alpha^4 N_{v_C}. \tag{1052}$$

The eigenvalue of W on v_C is $u = \lambda^2+\lambda$. Hence

$$W^n[W,h]W^{-n} = I + 2\lambda^3\alpha^4 u^{2n} N_{v_C}.$$

Since K has prime degree three and $u^2 \notin \mathbb{Q}$, the three elements $1, u^2, u^4$ are a $\mathbb{Q}$-basis of K; their coefficient determinant in the basis $1,\lambda,\alpha$ is 29. Products and inverses of the three displayed unipotents for $n = 0,1,2$ therefore give a full additive lattice in the three-dimensional Minkowski horosphere fixing Kv_C. Benoist–Miquel's horospherical arithmeticity theorem makes H_2 finite index. Its standard cusp intersection must then have rank three.

The exact certificate records

$$2\lambda^3\alpha^4 = 32 + 58\lambda + 72\alpha$$

and verifies every entry of (1052); it is implemented in `intermediate_cross_form_search.py`. □

Corollary II.20.93 (Wall norms are adjacent H_2-cusp incidences)**.** *For each vertex $X \in \{A,B,C,D\}$, choose a nontrivial parabolic $k_X = I + t_X N_{e_X} \in H_2$, which exists by Proposition II.20.92. Then for every edge $XY \in \{AC, AD, BC, BD\}$,*

$$|\,\mathrm{N}\,Q_{XY}(P)|^2 = \frac{|\,\mathrm{N}\det(P,k_XP)|\,|\,\mathrm{N}\det(P,k_YP)|}{|\,\mathrm{N}(t_Xt_Y)|}.$$

Indeed $\det(P,(I+tN_e)P)=t\det(P,e)^2$*, and* $Q_{XY}=X(P)Y(P)$*. The wall target is therefore an adjacent-cusp height in the arithmetic cusp complex of* H_2*. This is stronger structure than the earlier formal Segre square, but it still retains the essential ordering problem: unrestricted arithmetic reduction may choose any cusp, whereas the Rosen step supplies only the four adjacent ones.*

Experiment II.20.94 (The same wall word at $q=11$)**.** Keeping λ symbolic gives

$$LSE^{-1}SL=\begin{pmatrix}-1-2\lambda & -2\lambda(\lambda+1)\\ -2 & -1-2\lambda\end{pmatrix},\qquad Q=-2\big(X^2-\lambda(\lambda+1)Y^2\big).$$

At $q=7$, the identity $\lambda(\lambda+1)=(\lambda^2-1)^2$ makes this precisely the split AC wall. At $q=11$, exact resultants give instead

$$\mathrm{N}(\lambda(\lambda+1))=-1,\qquad \mathrm{N}(\lambda^2+4)=5741.$$

Hence the first quantity is not a square, and the wall field $K_{11}(\sqrt{\lambda(\lambda+1)})$ is different from the adjacent torus field $K_{11}(\sqrt{\lambda^2+4})$: equality would make their ratio a square, although its rational norm is $-1/5741$.

This test succeeds as a discriminator. The intermediate group $\langle G_q,T(2)\rangle$ exists uniformly, but its shortest natural wall meets the Rosen adjacent axis only at the order-seven algebraic identity above. Thus the H_2 square is a genuine extra resonance available for the heptagonal proof; it is not the family conjecture and should not be expected to survive at $q=11$. The symbolic calculation is reproduced by `intermediate_cross_form_search.py`.

Proposition II.20.95 (The wall square is not locally monotone)**.** *Let*

$$\epsilon_N=(\lambda\alpha)^N,\qquad q_N=(\epsilon_N-1)/\lambda,\qquad p_N=1-\alpha^{-1}q_N.$$

Then (p_N,q_N) *is primitive,*

$$A(p_N,q_N)=\epsilon_N,\qquad D(p_N,q_N)=1,$$

and hence both split orientation minima have norm one. For all large N*, the next Rosen digit is* $+1$*, but at the next state*

$$\begin{aligned}B'&=-\alpha,\\ C'&=\alpha^{-1}\epsilon_N-2\alpha,\\ D'&=2(2\alpha-2\lambda-1)\epsilon_N+(2+2\lambda-3\alpha).\end{aligned}$$

Both $|\,\mathrm{N}\,C'|$ *and* $|\,\mathrm{N}\,D'|$ *tend to infinity.*

Proof. Primitivity follows because a common ideal divisor would divide the two unit incidences ϵ_N and 1. At the distinguished place,

$$p_N/q_N\longrightarrow-\alpha^{-1},\qquad -q_N/(\lambda p_N)\longrightarrow\alpha/\lambda=1.2469\ldots,$$

so the state is reduced and its digit is one for all large N. The split transport formula gives $B'=-\alpha D$, and direct substitution gives C' and D'. The unit $\lambda\alpha$ has distinguished absolute value greater than one and both outer absolute values less than one. Thus the distinguished factors of C',D' grow exponentially, while their outer factors tend to the nonzero constant terms. Their field norms tend to infinity. □

This is the precise obstruction to converting the wall square into a one-step descent. At a $+1$ digit one opposite pair is transported by units while the other receives the transverse shear $2\lambda r_j$; at -1 the roles reverse. One half of the square is controlled at a time. The exact audit `rosen_h2_wall_height_audit.py` also checks 17945 primitive admissible states from a deterministic 20000-sample run. The one-step logarithmic change ranged from -7.1337869633 to 6.7632968753. A small exact example is

$$p = (-1, 4, -4), \qquad q = 5$$

in the basis $1, \lambda, \alpha$: its digit is $+1$, its two current minima are $1, 1$, and its two following minima are $1, 239$. The parametric family proves that the increase is unbounded, so the computation is only an implementation check.

Route II.20.96 (Exact remaining statement: H_2 wall shadowing)**.** Let $\mathcal{H}_\square$ be the four elements in Proposition II.20.91. If g_j is the j-th Rosen prefix of a primitive endpoint column P, covariance gives

$$\det(P, g_j^{-1} h g_j P) = Q_h(g_j P).$$

Consequently the sufficient all-time split estimate is equivalent, up to the fixed scalar norms 8 and 27, to

$$\min_{h \in \mathcal{H}_\square} |\,\mathrm{N}\det(P, g_j^{-1} h g_j P)| \le C(P) \quad \text{for every } j. \tag{1053}$$

This is the cleanest current formulation of the missing theorem. It is an ordered nearest-wall statement for four conjugacy families in one nontrivial H_2 congruence coset. It is global rather than primewise, and Proposition II.20.95 proves that the initial Arakelov height cannot be replaced by a local monotonicity argument.

Proposition II.20.92 rules out one tempting shortcut. Since H_2 is finite index, every K-rational line has an H_2-parabolic stabilizer. Unrestricted Hilbert modular reduction would therefore recover only this automatic ambient stabilizer. The displayed inequality asks for much more: reduction by four prescribed walls, in the order in which they are conjugated by consecutive prefixes of the thin subgroup G_7. The remaining assertion is a relative shadowing theorem between the Rosen coding and the arithmetic cusp complex of H_2, not ordinary reduction theory for H_2.

The thin-group stammering theorem proves (1053) indirectly whenever one repeated block dominates the accumulated denominator height, because it closes the orbit outright. The remaining case is the balanced many-block regime: the transverse shears must be fused across successive choices of opposite wall while retaining the same integral divisor. This is the concrete content of the stopped rank-two Arakelov inequality.

The additive action remains informative despite arithmeticity. Multiplication on the left by $T(\ell)$, $\ell \in L_2$, sends a first column $(A, \Delta)^t$ to $(A + \ell\Delta, \Delta)^t$. It changes the additive representative $-A/\Delta$ but not the ray–Levi class $[\bar{A}] \in \mathscr{R}_{(\Delta)}$. Hence the rank-two cusp plane acts in the splitting fibre. Proposition II.20.91 proves that this fibre contains the four walls selected by the sufficient split estimate. Proposition II.20.92 adds the missing global fact: their holonomy produces a third cusp direction on another rational line and arithmeticizes the whole group. The obstruction is therefore not a lack of ambient reduction; it is retaining the consecutive G_7-order while borrowing that reduction.

This threshold is comparable to Hao's arithmeticity criterion for rational groups generated by root unipotents: there, too, one additional upper-triangular element beyond the visible cusp group forces a congruence group [63]. His theorem does not include the present cubic triangle group, but

it identifies the same structural test. Here the unexpected unipotent has been produced explicitly in (1052); the missing input is the ordered relative shadowing inequality, not the index of H_2.

The closest literature now falls into three complementary columns. McMullen's height is the direct Hodge-height precedent [17]. Arakelov slope and arithmetic-holonomy methods explain how a finite divisor can be fused with analytic transport [57, 59]. Relative/log-Anosov theory explains why ordered orbit tracking should force a parabolic boundary point [20, 41]. None of these inputs alone retains the single determinant divisor, the cyclic Galois weight, and the consecutive Rosen quotient. The comparison with the congruence hull says exactly where the new theorem must live: it is an archimedean and ordered failure-of-lifting inequality between two groups with the same finite closure.

II.20.18 The norm-Euclidean comparison: a successful reformulation

There is a more useful intermediate structure than an unlabelled overgroup. The field K has discriminant 49, and Cavallar–Lemmermeyer compute its exact Euclidean minimum

$$M(K) = \sup_{\xi \in K} \inf_{q \in \mathcal{O}_K} |N_{K/\mathbb{Q}}(\xi - q)| = \frac{1}{7}$$

[12]. Apply this to a Rosen transition

$$r_{j+1} = -r_{j-1} - n_j \lambda r_j, \qquad \xi_j = \frac{r_{j-1}}{\lambda r_j}.$$

There is a $q_j \in \mathcal{O}_K$ for which, on putting $\kappa_j = q_j + n_j$,

$$|N(r_{j+1} + \kappa_j \lambda r_j)| \le \frac{1}{7} |N(r_j)|. \tag{1054}$$

This is exact: it follows by multiplying $|N(\xi_j - q_j)| \le 1/7$ by $|N(\lambda r_j)| = |N(r_j)|$ and using the recurrence. If r_{j+1} is a new norm record, then $\kappa_j \neq 0$, since $\kappa_j = 0$ would make its norm at most one seventh of the preceding norm.

At group level the comparison has no error term. With

$$B(n) = \begin{pmatrix} -n\lambda & -1 \\ 1 & 0 \end{pmatrix}, \qquad A(q) = \begin{pmatrix} q\lambda & -1 \\ 1 & 0 \end{pmatrix},$$

one has

$$A(q_j) = T(\kappa_j \lambda) B(n_j). \tag{1055}$$

The ambient norm-reducing move is exactly the thin move followed by its labelled correction translation. This produces a relative cocycle for the pair $\Delta \subset \mathrm{SL}_2(\mathcal{O}_K)$. Replacing that cocycle by the subgroup generated by all its values is precisely the operation which forgets consecutiveness and leads back to arithmeticity.

This calculation isolates the rank two cleanly. The thin Rosen quotient is $-n_j \in \mathbb{Z}$; the ambient norm-Euclidean quotient is $q_j \in \mathcal{O}_K$; their transverse difference is

$$[\kappa_j] = [q_j] \in \mathcal{O}_K/\mathbb{Z},$$

a free rank-two group. A nonzero class is a genuinely unavailable quotient direction. A zero class means that another integer digit gives ambient norm descent but is rejected by the distinguished nearest-integer rule. Thus the factor $1/7$ does not prove the theorem, but it improves the target

substantially: the Hodge inequality has to control the consecutive correction process (κ_j), rather than arbitrary rational approximations.

The finite transport also has an exact consecutive law. Set

$$\rho_j = [\overline{r_j}] \in (\mathcal{O}_K/(r_{j-1}))^\times / \operatorname{im}(\mathcal{O}_K^\times).$$

Modulo r_j, the recurrence gives

$$\rho_{j+1} = [-\overline{r_{j-1}}] = [\overline{r_{j-1}}].$$

Because adjacent remainders are coprime, the CRT class

$$\overline{[r_j - r_{j-1}]} \in (\mathcal{O}_K/(r_j r_{j-1}))^\times / \operatorname{im}(\mathcal{O}_K^\times)$$

has these two successive ray classes as its reductions. It contains more information than the separately quotiented pair because it retains one common, rather than two independent, unit normalizations. The digit shear is invisible in the residue class and acts by choosing the next modulus. This is the first exact finite-place formula in the audit which uses consecutiveness without reducing it to a list of generators.

Allowing quotients in a rank-two module $\mathbb{Z} \subset M \subset \mathcal{O}_K$ gives an intermediate *algorithm*. For the proper subgroup $H_2 = \langle G_7, T(2)\rangle$, the corresponding module is $M_2 = \mathbb{Z} + 2\lambda^{-1}\mathbb{Z}$, since $\lambda M_2 = \lambda\mathbb{Z} + 2\mathbb{Z}$. Generating a group from all these moves forgets when the correction was used; allowing the full ring gives the norm-Euclidean algorithm and the full cusp lattice, which is arithmetic. The object worth exploiting is therefore the labelled path (κ_j), coupled to the Hodge transport, rather than H_2 by itself.

The correction does not change the finite ray state. Define

$$r_{j+1}^{\mathrm{E}} = q_j \lambda r_j - r_{j-1} = r_{j+1} + \kappa_j \lambda r_j.$$

Then

$$(r_{j+1}^{\mathrm{E}}, r_j) = \mathcal{O}_K, \qquad r_{j+1}^{\mathrm{E}} \equiv r_{j+1} \pmod{r_j}, \qquad |N r_{j+1}^{\mathrm{E}}| \le \frac{1}{7}|N r_j|.$$

At a record, unless $r_{j+1}^{\mathrm{E}} = 0$,

$$\frac{|N r_{j+1}|}{|N r_{j+1}^{\mathrm{E}}|} > 7.$$

Thus the Rosen and Euclidean remainders are two representatives of the same ray class with the same modulus and a norm gap greater than seven. For the corresponding splittings

$$s_j = -r_{j+1}/r_j, \qquad s_j^{\mathrm{E}} = -r_{j+1}^{\mathrm{E}}/r_j,$$

one has

$$s_j^{\mathrm{E}} = s_j - \kappa_j \lambda, \qquad (s_j, r_j^{-1}) = (s_j^{\mathrm{E}}, r_j^{-1}) = r_j^{-1}\mathcal{O}_K.$$

The finite content is literally identical. This succeeds in removing the moving-modulus ambiguity from the local comparison; the unresolved loss is entirely in the stopped archimedean norms.

There is also exact two-edge information:

$$(r_{j+1}^{\mathrm{E}}, r_{j+1}) = (\kappa_j, r_{j+1}), \qquad r_{j+1}^{\mathrm{E}} \equiv \kappa_j \lambda r_j \pmod{r_{j+1}}.$$

If κ_j is a unit, the Euclidean remainder represents both the current ray class modulo r_j and the reciprocal consecutive class modulo r_{j+1}. If it is not a unit, the one ideal (κ_j, r_{j+1}) is the exact obstruction to that two-edge fusion. This is global gcd information, not a prime-by-prime strategy.

II.20.18.1 The Eisenstein correction lattice

Choose τ so that

$$\tau(\lambda) = \lambda^2 - \lambda - 1, \qquad \tau(\lambda^2) = 2 - \lambda.$$

On $\mathcal{O}_K/\mathbb{Z}$, in the basis $(\bar{\lambda}, \overline{\lambda^2})$, it acts by

$$\begin{pmatrix} -1 & -1 \\ 1 & 0 \end{pmatrix}.$$

The invariant positive form is

$$Q(b,c) = b^2 + bc + c^2.$$

For $\kappa = a + b\lambda + c\lambda^2$, Newton sums give the intrinsic identity

$$\sum_{i=0}^{2} \left(\sigma_i(\kappa) - \frac{1}{3}\operatorname{Tr}(\kappa) \right)^2 = \frac{14}{3} Q(b,c).$$

Thus the two missing directions form one rank-one Eisenstein module, rather than two unrelated integer coordinates.

This also gives a new exact negative result about intermediate groups. If an additive module

$$\mathbb{Z} \subsetneq M \subseteq \mathcal{O}_K$$

is Galois stable, then its image in $\mathcal{O}_K/\mathbb{Z}$ has rank two: the matrix above has irreducible characteristic polynomial $X^2 + X + 1$. Hence M has rank three and finite index in $\mathcal{O}_K$. Adjoining the translations $T(\lambda M)$ therefore supplies a full Minkowski cusp lattice, and Benoist–Miquel forces finite index. Any rank-two intermediate quotient algorithm must break Galois symmetry. This is why H_2 can exist as a proper candidate but cannot itself encode the Galois-completed Hodge inequality.

Finally the quadratic form has a direct normal-distance meaning. On a compact return set $\mathcal{K} \Subset \mathbb{H}$, define

$$\mathcal{D}_z(\bar{\kappa}) = \inf_{m \in \mathbb{Z}} \sum_{i=0}^{2} d_{\mathbb{H}} \left(F_i(z), T(\sigma_i(\lambda(\kappa - m)))F_i(z) \right).$$

The horizontal-translation formula

$$d_{\mathbb{H}}(w, T(t)w) = 2 \operatorname{arsinh} \frac{|t|}{2 \operatorname{Im} w}$$

together with the trace identity proves, uniformly for $z \in \mathcal{K}$,

$$\mathcal{D}_z(\bar{\kappa}) \asymp_{\mathcal{K}} \log(1 + Q(b,c)).$$

It is exactly the distance from the ambient correction translate to the thin cusp orbit. Therefore a cumulative stopped estimate

$$\sum_{\text{record } j} \begin{cases} \log(1 + Q(\bar{\kappa}_j)), & \bar{\kappa}_j \neq 0, \\ \log(1 + |\kappa_j|), & \kappa_j \in \mathbb{Z} \setminus \{0\}, \end{cases} \leq 6h(P) + O_7(1)$$

would suffice for the classification. Every summand is at least $\log 2$, so only finitely many norm records could occur; bounded norms then close by the proved fixed-support S-unit theorem. What is not yet supplied is the telescoping Hodge potential which pays this correction charge. The

transverse term should come from the normal Hodge defect, while the integral term lies on the ordinary central cusp channel.

This also rules out a tempting bridge experiment. Inserting arbitrary elements of $\Gamma_{\mathrm{cong}}(G_7)$ between two G_7-blocks may destroy a graph correlation, but the resulting closing element belongs only to the hull, where closing is automatic. A valid bridge must be a consecutive G_7-word, or must cancel algebraically so that the final stabilizer lies in G_7. The natural intermediate object is therefore the relative pair $G_7 \subset \Gamma_{\mathrm{cong}}(G_7)$, used to define the defect, rather than an intermediate quotient surface used to erase it.

The suggested relative-Borel refinement was then checked in an adapted frame, and this produced a useful negative theorem rather than the missing gain. Choose a primitive cusp column C, put $\delta = \det(P, C)$, and use the determinant-one frame $(P, C/\delta)$. If

$$(P, C/\delta)^{-1} k (P, C/\delta) = \begin{pmatrix} a & c \\ b & d \end{pmatrix},$$

then

$$\alpha(k) := \delta a = \det(kP, C), \qquad \beta(k) := b = \det(P, kP).$$

The full Borel stabilizing P is cut out by the single equation $\beta(k) = 0$. Moreover

$$\delta kP = \alpha(k)P + \beta(k)C.$$

Localizing at a prime $\mathfrak{p}$, primitivity of kP gives the sharp fixed-gcd bound

$$\min\{v_{\mathfrak{p}}(\alpha(k)), v_{\mathfrak{p}}(\beta(k))\} \le v_{\mathfrak{p}}(\delta).$$

Thus, away from the fixed normalization ideal, the two boundary coordinates are coprime. The equation $\alpha = 0$ is not a second equation for the same Borel: it sends P to the other boundary line $[C]$, which is the cusp alternative. Targeting one specified unipotent would artificially impose further equations such as $a = d = 1$, but that would discard the desired special-hyperbolic alternative. Hence the correct ambient target has arithmetic codimension one, and the congruence hull cannot create a second finite divisor.

There is also an exact reason finite congruence information cannot make this codimension-one contact rigid. Choose a noncentral unipotent u in the hull Borel. Since G_7 and its hull have the same finite-adelic closure, for arbitrarily large ideals $\mathfrak{q}$ one can find $k_{\mathfrak{q}} \in G_7$ with $k_{\mathfrak{q}} \equiv u \pmod{\mathfrak{q}}$. In the hard case, where $G_7 \cap B_P$ is projectively trivial, these elements satisfy

$$0 \neq \beta(k_{\mathfrak{q}}) \in \mathfrak{q}.$$

So arbitrarily deep false Borel contact is allowed at the finite places. The attempted congruence-hull descent therefore failed for a precise reason: only the archimedean cost and the order of a consecutive Rosen word can distinguish false contact from actual incidence. This is not evidence for adding more primes; it is evidence that the remaining inequality must be an ordered relative-holonomy statement.

There is a second, independent obstruction to promoting a residue Borel to a genuine integral Borel. For

$$\begin{pmatrix} a & c \\ b & d \end{pmatrix} \in \mathrm{SL}_2(\mathcal{O}_K), \qquad b \neq 0,$$

the reduction modulo (b) is upper triangular. It is congruent to an integral upper triangular matrix with unit diagonal if and only if

$$[a] = 1 \quad \text{in} \quad \mathscr{R}_{(b)} = (\mathcal{O}_K/(b))^\times / \operatorname{im}(\mathcal{O}_K^\times).$$

Indeed a lift must have upper-left entry a global unit congruent to a, and that condition is also sufficient; the upper-right entry can be lifted arbitrarily and the lower-right congruence follows from the determinant. Changing the second vector in the P-adapted frame replaces a by $a - tb$, so this quotient class is intrinsic. It is the exact ray-class, or Levi, obstruction.

The experiment `levi_residue_tests.py` shows that this obstruction is present in the actual Rosen language. For

$$P = (6 + 2\lambda - \lambda^2, 1)^t$$

the first digits are $4, 1, 1$, and the corresponding adapted first column after two inversion steps is

$$a = 3\lambda^2 - 4\lambda - 2, \qquad b = 11\lambda^2 - 9\lambda - 16, \qquad |N(b)| = 769.$$

PARI gives $(\mathcal{O}_K/(b))^\times \simeq \mathbb{Z}/768\mathbb{Z}$. In one cyclic coordinate the logarithms of $a, \lambda, \lambda^2 - 1, -1$ are

$$100, \quad 519, \quad 432, \quad 384 \pmod{768}.$$

The three global-unit logs generate the multiples of 3, whereas $100 \notin 3\mathbb{Z}$. Thus this legitimate consecutive prefix has no integral Borel lift. This is not merely a short-word accident. The special expansion of 671 is detected at step 61, with period 23; among its 62 nontrivial initial-prefix residues, 32 fail the unit-lift test. The largest observed index of the unit image is 1176. Exact unit normalization is therefore a failed route.

The elementary closing threshold left after this failure is worth stating separately. If $0 \neq \mathfrak{q} \subset \mathcal{O}_K$ and $x \in \mathfrak{q}$, then

$$x \neq 0 \quad \Longrightarrow \quad \sum_{i=0}^{2} \log^+ |\sigma_i(x)| \geq \log N\mathfrak{q}.$$

This is just $|N(x)| \geq N\mathfrak{q}$ and the product formula. Hence a strict reverse inequality forces $x = 0$. It is the optimal final algebraic step, but not the missing estimate: taking $\mathfrak{q} = (x)$ lands exactly at equality. The stopped Hodge transport has to create strict slack while the determinant divisor is counted once. The Levi calculation explains why that slack cannot be created by first replacing every congruence Borel with a global-unit normalization.

This identifies a classical comparison which is more accurate than ordinary Roth. Since K has class number one, the ray-class exact sequence identifies

$$\mathscr{R}_{(b)} \simeq \mathrm{Cl}_{(b)}(K).$$

The rational shear $s = -a/b$, together with its finite content $(s, b^{-1}) = b^{-1}\mathcal{O}_K$ and its three Hodge norms, is therefore an Arakelov representative of this ray class. Schoof's Arakelov-class and infrastructure formalism is the closest established number-theoretic model: it combines an ideal with a logarithmic unit-shape coordinate and supplies reduced representatives [19]. It does not retain a modulus varying with a consecutive thin word, nor a metric transported by a Picard–Fuchs system. The remaining fusion inequality is precisely the additional non-drift statement for this Hodge-selected sequence of Arakelov ray classes. This makes the division between number theory and Hodge theory exact: the finite state is ray-class number theory, the strict surplus is Hodge curvature, and the consecutive thin orbit is the coupling between them.

II.21 Exact audits of the rank-two Arakelov and holonomy routes

Route II.21.1 (Two rational splittings: an exact global angle identity)**.** Let $q = Q_m$, $v = Q_n$, $\Delta = \det(q, v)$, and $V = v/\Delta$. The rational splitting vectors in the stopped rank-two extension are

$$W(s) = V + sq, \qquad \mathfrak{c}(s) = (s, \Delta^{-1}).$$

For two distinct splittings define

$$\mathfrak{J}(s, s') = \left(\frac{s - s'}{\Delta}\right) \mathfrak{c}(s)^{-1}\mathfrak{c}(s')^{-1}.$$

The coordinate determinant of $W(s), W(s')$ in the frame (q, v) is $(s - s')/\Delta$. It belongs to $\mathfrak{c}(s)\mathfrak{c}(s')$, so $\mathfrak{J}(s, s')$ is an integral ideal. The determinant-one Hodge metric gives, at every real place,

$$|\sigma_i(s - s')| = \|W(s)_i\| \, \|W(s')_i\| \sin\theta_i.$$

The product formula therefore yields the exact identity

$$\begin{aligned} &\widehat{\deg}(Ke(s)) + \widehat{\deg}(Ke(s')) - \widehat{\deg}(Ke_0) \\ &\quad = -\log|N\Delta| - \log N\mathfrak{J}(s, s') + \sum_{i=0}^{2} \log\sin\theta_i. \end{aligned} \tag{1056}$$

This is stronger than a generic Minkowski estimate: the finite intersection and all three archimedean angles occur with their exact multiplicities.

There is a canonical pair. The controlled splitting is $s_{\mathrm{ctl}} = 0$, while the true-past or dynamic splitting is $s_{\mathrm{dyn}} = -A/\Delta$. In a determinant-one integral P-frame,

$$\mathfrak{c}(s_{\mathrm{ctl}}) = \mathfrak{c}(s_{\mathrm{dyn}}) = \Delta^{-1}\mathcal{O}_K, \qquad \mathfrak{J}(s_{\mathrm{dyn}}, s_{\mathrm{ctl}}) = (A).$$

Thus

$$\begin{aligned} &\widehat{\deg}(Ke(s_{\mathrm{dyn}})) + \widehat{\deg}(Ke(s_{\mathrm{ctl}})) - \widehat{\deg}(Ke_0) \\ &\quad = -\log|N(A\Delta)| + \sum_i \log\sin\theta_i. \end{aligned} \tag{1057}$$

The two lines coincide exactly when $A = 0$, the cusp incidence; the extension loses rank exactly when $\Delta = 0$, the frame-recurrence incidence. Hence the two algebraic alternatives in the theorem are the two degeneracies of one elementary angle formula.

This calculation materially sharpens the description of the bottleneck. No auxiliary rank-two lattice and no second rational line remain to be invented. I tested the tempting terminal-degree inequality

$$\begin{aligned} &\widehat{\deg}_{t_n}(Ke(s_{\mathrm{dyn}})) + \widehat{\deg}_{t_n}(Ke(s_{\mathrm{ctl}})) - \widehat{\deg}_{t_n}(Ke_0) \\ &\quad \geq -\log|NA| + \eta\max\{X_0, \Delta\mathfrak{A}^{\mathrm{ctl}}\} - C. \end{aligned} \tag{1058}$$

Here

$$\Delta\mathfrak{A}^{\mathrm{ctl}} = \mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_n) - \mathfrak{A}^{\mathrm{ctl}}_{m,n}(t_m).$$

It is formally sufficient: comparing it with (1057) gives

$$\log|N\Delta| + \eta\max\{X_0, \Delta\mathfrak{A}^{\mathrm{ctl}}\} \leq C + \sum_i \log\sin\theta_i \leq C,$$

which is the rank-two stopped fusion estimate.

This candidate is rejected as a shortcut. The comparison shows that it is stopped fusion made strictly stronger by the extra nonpositive term $\sum_i \log \sin \theta_i$. Moreover, subtracting (1057) at the two endpoints says that the variation of the ordinary two-line degree is only

$$\sum_i \log \frac{\sin \theta_i(t_n)}{\sin \theta_i(t_m)};$$

it contains no interior Schwarz–Pick curvature. For arbitrary primitive pairs this angle can collapse without bound. Thus the elementary determinant calculation identifies the endpoint wall but does not cross it. A valid proof must insert the Hodge twist and its stopped Bott–Chern metric on the one-divisor cyclic pair; using ordinary terminal degrees simply restates the desired conclusion.

Experiment II.21.2 (Exact denominator cost of the heptagon periods)**.** For $k = 1, 2, 3$, the holomorphic Frobenius solution has parameters

$$(a_k, b_k) = \frac{1}{28}(19, 5), \quad \frac{1}{28}(17, 3), \quad \frac{1}{28}(15, 1).$$

Writing ${}_2F_1(a_k, b_k; 1; t) = \sum c_{k,n} t^n$, direct valuation gives

$$v_2(\operatorname{den} c_{k,n}) = 4n + 2v_2(n!), \qquad v_7(\operatorname{den} c_{k,n}) = 2n + 2v_7(n!),$$

and no other denominator prime occurs. Therefore

$$C_{\mathrm{hyp}} = 2^6 7^3 = 21952, \qquad {}_2F_1(a_k, b_k; 1; C_{\mathrm{hyp}} x) \in \mathbb{Z}[[x]]$$

for all three k, and the two exponents are minimal among integral rescalings.

The program `hypergeom_denominator_audit.py` checked the first 500 coefficients of each series exactly. In all three cases it found only the primes $2, 7$, verified integrality after multiplication of the n-th coefficient by 21952^n, found the first failure of a 2^5-scaling at $n = 4$, and the first failure of a 7^2-scaling at $n = 7$. The largest denominator checked had 1884 decimal digits. The valuation formula proves the result for all n; the computation is retained as a regression check.

Outcome: the differential system has one fixed, completely explicit denominator cost. It does not create the stopped surplus. Under $t = C_{\mathrm{hyp}} x$, the singularity $t = 1$ moves to $x = C_{\mathrm{hyp}}^{-1}$, so the better integral coefficients are paid for by a smaller analytic radius. Any useful holonomy gain must be relative overconvergence of the stopped combination, not a normalization trick.

Route II.21.3 (Arithmetic holonomy: exact hypothesis versus the needed stopped statement)**.** Theorem 2.6 of Calegari–Dimitrov–Tang does combine simultaneous approximation at several places, coefficient denominators, conformal size, and Bost–Charles metrics in one slope argument [57]. Its condition (1), however, requires an exact meromorphic continuation: $\varphi_u^* g$ must be meromorphic on the full disc, while the individual coefficient functions need converge only on the smaller disc of radius ρ_u. The theorem then controls rational approximation to the fixed local constants which occur in that overconvergent linear combination.

The stopped Hodge dichotomy supplies a different kind of datum. It says that one of two norms gains a definite fraction of the integrated Schwarz–Pick defect. It does not say that a period combination extends meromorphically across the first singularity. A long triangle-group word is monodromy of the same rank-two hypergeometric system; by itself it does not enlarge its conformal radius. If one chooses a word-dependent map $\varphi_{m,n}$ which crosses a singularity only after assuming

that the monodromy jump cancels, that cancellation is already the incidence $A\Delta = 0$ which the proof is meant to force. That use of the theorem would be circular.

There is nevertheless a precise possible bridge. The monodromy jump of the stopped dynamic–controlled pair is algebraic and its two finite incidence factors are exactly A and Δ, by (1057). A quantitative *stopped removability* lemma could show that the Hodge norm makes this jump smaller, simultaneously at all places, than its algebraic height allows. The product formula would then make the jump zero, after which the exact holonomy theorem would apply. Such a lemma must use the stopped cyclic Hodge metric, not the rejected ordinary-degree bound (1058). The published theorem supplies the global slope architecture after exact overconvergence; it does not presently supply this approximate-to-exact step.

Outcome: arithmetic holonomy remains structurally relevant, and its denominator side is now completely explicit. The unresolved input has not moved: it is the consecutive-thin-orbit noncollapse of the canonical two-line angle. Calling a long monodromy branch *overconvergence* does not prove it.

Route II.21.4 (Near-definable-set counting and thin-orbit amplification)**.** There is a genuinely different possible use of the whole thin group. Habegger's approximation form of the Pila–Wilkie theorem says, roughly, that algebraic points of fixed degree and height at most T, lying within $T^{-\lambda}$ of a closed set definable in a polynomially bounded o-minimal structure, number $O_\epsilon(T^\epsilon)$ away from the positive-dimensional semialgebraic locus, once λ is large enough [60]. The equality corridor has the right scale for such a theorem. If

$$h([g_nP]) = o(t_n)$$

and Hodge alignment gave distance e^{-ct_n} from a fixed definable set, then with $T_n = \exp(O(h([g_nP])))$ the approximation exponent $ct_n/\log T_n$ would tend to infinity.

The thin group also suggests an amplification device. From one word of length t_n, append words of length

$$L_n = \sqrt{t_n\big(h([g_nP]) + 1\big)}.$$

Then $h([g_nP]) \ll L_n = o(t_n)$. A free subsemigroup supplies $\exp(cL_n)$ distinct continuations, while the height of each continued state is at most $O(h([g_nP]) + L_n) = O(L_n)$. If all these continuations remained within $e^{-ct_n+O(L_n)}$ of one fixed transcendental definable curve, Habegger's T^ϵ bound, with ϵ smaller than the semigroup growth-to-height ratio, would give a contradiction. This would use exponential branching of the full thin group, not estimates for its generators.

The required hypotheses are not presently available, for two separate reasons.

1. A small reduced height $h([g_nP])$ does not itself put the Galois point $(\sigma_0(g_nP), \sigma_1(g_nP), \sigma_2(g_nP))$ near a fixed low-dimensional definable set. The holomorphic modular-embedding graph in $\mathbb{H}^3$ is real analytic on compact interior charts, but g_nP is a flat boundary line. The finite-time outer singular flags also depend on the full natural-extension past, not only on that boundary line. Replacing them by a global boundary map would require regularity and definability which have not been proved for the dense outer representations.
2. Left multiplication by a branching family preserves the elementary height bound, but it does not preserve the fact that the word is a consecutive prefix selected by the original endpoint. Thus the stopped Hodge alignment, even if known for g_n, need not hold for the exponentially many continued words. Without this persistence, one has only one algebraic point at each height, and a T^ϵ counting theorem gives no contradiction because those heights may grow arbitrarily fast.

Outcome: the Pila–Wilkie idea does not currently prove the theorem. It does isolate a possible alternative theorem: construct a polynomially bounded definable return correspondence for the complete future–past–Hodge state and prove that a positive-entropy family of thin continuations remains in its superpolynomial tube. That statement would subsume the missing stopped alignment, and is presently less direct than the one-divisor Arakelov formulation.

Route II.21.5 (Alternating unit corridors: an exact height lemma that works)**.** The one-step split-factor calculation initially appeared to create a new obstruction: a legal digit can take one of the two orientation minima from norm one to arbitrarily large norm. Following the orbit rather than freezing that one step reveals the structure. The state enters a long alternating $1, -1$ run, and that entire run is a diagonal translation by a unit. The correct question is whether its length creates new projective height. It does not.

Write

$$A = r + \alpha s, \quad B = r - \alpha^{-1} s, \quad C = r - \alpha s, \quad D = r + \alpha^{-1} s, \qquad \alpha = \lambda_7^2 - 1.$$

The exact branch transport is

$$\begin{array}{c|cc} a = 1 & A' = \alpha^{-1} C & B' = -\alpha D \\ a = -1 & C' = -\alpha^{-1} A & D' = \alpha B. \end{array}$$

Consequently, along a $1, -1$ pair,

$$C_{j+2} = -\alpha^{-2} C_j, \qquad D_{j+2} = -\alpha^2 D_j, \qquad z_{j+2} = \alpha^{-4} z_j, \quad z = C/D.$$

Along a $-1, 1$ pair the same formula holds for $w = A/B$.

Here is the all-place argument. Put $\kappa = \alpha^{-4}$ and $b_i = \log |\sigma_i \kappa|$. Then

$$b_0 < 0, \qquad b_1, b_2 > 0, \qquad b_0 + b_1 + b_2 = 0.$$

The last equality is the product formula, and the first three signs follow from $|\sigma_0 \alpha| > 1 > |\sigma_1 \alpha|, |\sigma_2 \alpha|$. The distinguished Rosen state lies in $I_7 = [-\lambda_7/2, \lambda_7/2)$. Since $\alpha - \alpha^{-1} = \lambda_7$, every $x \in I_7$ satisfies

$$\left| \frac{x - \alpha}{x + \alpha^{-1}} \right| \geq 1.$$

Thus the distinguished logarithm of z stays nonnegative at every even state for which the alternating block continues.

If $u_i = \log |\sigma_i(z_j)|$, multiplication by the unit κ^m changes no finite-place height contribution. The central positive part changes by exactly $m b_0$, while each outer positive part changes by at most $m b_i$:

$$\max(0, u_i + m b_i) - \max(0, u_i) \leq m b_i \quad (i = 1, 2).$$

Equation (II.21.5) therefore gives the exact monotonicity

$$h(z_{j+2m}) \leq h(z_j). \tag{1059}$$

The fixed fractional-linear changes between x and z, and the one fixed branch leading to an odd state, alter height by only $O_7(1)$. The same proof for w handles the other alternating orientation. Hence every state in an alternating block obeys

$$h(x_{j+k}) \leq h(x_j) + O_7(1), \tag{1060}$$

with an error independent of the block length.

This argument resolves an ambiguity in the earlier calculations. The failed common-invariant search found that the eight-dimensional Asai branch matrices preserve no nonzero symmetric quadratic form. Equation (1059) does not contradict that result: it is a nonlinear positive-part inequality, and it holds only on the Rosen cylinder where (II.21.5) fixes the central sign. This is the first place where the admissible thin orbit, rather than the ambient linear representation, supplies the missing information.

The family which exposed the one-step failure makes the point sharply. Let

$$\epsilon_N = (\lambda\alpha)^N, \qquad q_N = (\epsilon_N - 1)/\lambda, \qquad p_N = 1 - \alpha^{-1} q_N.$$

Its initial state has $A = \epsilon_N, D = 1$. It starts with alternating digits and leaves that block only when the central coordinate in (II.21.5) reaches a fixed branch boundary. Therefore the number m_N of complete digit pairs is

$$m_N = \frac{\log|\sigma_0(\lambda\alpha)|}{4\log\sigma_0(\alpha)} N + O(1) = 0.4318404528\ldots N + O(1). \tag{1061}$$

The exact program `rosen_neutral_corridor_height_audit.py` checks the branch convention and the height ledger. Its standard report is

N	alternating digits	pairs/N	$\max(h(x_k) - h(x_0))$
10	8	0.400	0
20	16	0.400	0.0069844
40	34	0.425	0.0184452
80	68	0.425	0.0071789

The computed increases of $h(z)$ at successive even states are below 10^{-170} in the largest run and are numerical roundoff around the proved nonpositive quantity. The table is only a check of conventions; the proof is the positive-part ledger above.

What works is now precise: arbitrarily long alternating corridors do not compound height. What does not follow is equally precise: their entry heights need not be bounded, and the two wall minima need not be monotone inside them. Cut the digit word before and after every digit with absolute value greater than one and between equal adjacent unit digits. The cut vertices are the junction states. Then

$$\max_{k\le n} h(x_k) \le \max_{\substack{b\le n \\ b\ ma\ junction}} h(x_b) + O_7(1).$$

Thus the junction estimate

$$h(x_b) \le 2h(x_0) + O_7(1)$$

would imply the all-time split-factor target, but the Pisot-unit family recorded later produces a growing tested excess over this bound. It is therefore not retained as a proposed estimate. The lemma says only that the duration of one alternating block is harmless; it does not control the transition to the next block.

II.22 The $q = 11$ comparison: scalar sign versus unit-rank fusion

The comparison with $q = 11$ clarifies which part of the heptagonal strategy is special to $q = 7$. Put $\lambda_{11} = 2\cos(\pi/11)$. The identity $\lambda_{11}^2 - 2 = 2\cos(2\pi/11)$, together with the equality of degrees, gives

$$\mathbb{Q}(\lambda_{11}) = \mathbb{Q}(\zeta_{11})^+.$$

Exact PARI computations give

$$\lambda_{11}^5-\lambda_{11}^4-4\lambda_{11}^3+3\lambda_{11}^2+3\lambda_{11}-1=0,\quad \operatorname{disc} K_{11}=11^4,\quad \mathcal{O}_{K_{11}}=\mathbb{Z}[\lambda_{11}],\quad h(K_{11})=1,\quad M(K_{11})=\frac{1}{11}.$$

The Euclidean-minimum equality is Cerri's exact computation, recorded in [13]. Thus the comparison below is not caused by a loss of norm-Euclidean division in degree five.

For odd q, index the projective real embeddings by

$$\mathcal{M}_q=\{m:1\le m<q/2,\ \gcd(m,q)=1\}.$$

The two elliptic rotation numbers are $1/2$ and m/q, so the relative Euler number and normalized Hodge weight are

$$e_m=\frac{q-2m}{2q},\qquad \chi_m=\frac{e_m}{e_1}=\frac{q-2m}{q-2}.$$

Define

$$\Delta_q=1-\sum_{m\in\mathcal{M}_q,\ m\neq 1}\chi_m.$$

The formal exponent of the full integral endpoint Asai tensor is $-\Delta_q$. The exact comparison is

q	Hodge weights	Δ_q	Asai exponent
7	$(1,3/5,1/5)$	$1/5$	$-1/5$
9	$(1,5/7,1/7)$	$1/7$	$-1/7$
11	$(1,7/9,5/9,3/9,1/9)$	$-7/9$	$7/9$.

Thus the same integral tensor which supplies a generic contraction at $q=7$ has the opposite sign at $q=11$.

The sign threshold is exact for the whole odd family, not merely for primes. If $q\ge 11$ and $3\nmid q$, the allowed Hodge indices $2,3$ already contribute

$$\frac{q-4}{q-2}+\frac{q-6}{q-2}>1.$$

If $3\mid q$, then $q\ge 15$, and the allowed indices $2,4$ contribute

$$\frac{q-4}{q-2}+\frac{q-8}{q-2}>1.$$

Thus $\Delta_q<0$ for every odd $q\ge 11$. Direct calculation gives

$$\Delta_3=1,\qquad \Delta_5=\frac{2}{3},\qquad \Delta_7=\frac{1}{5},\qquad \Delta_9=\frac{1}{7}.$$

Therefore 11 is exactly the first odd parameter above the scalar Hodge threshold. Its significance is not simply that the trace field has degree five.

For an odd prime $p=2g+1$, the sum is elementary:

$$\sum_{m=2}^{g}\chi_m=\frac{(g-1)^2}{2g-1},\qquad \Delta_p=\frac{-g^2+4g-2}{2g-1}.$$

Consequently $p=7$ is the last odd prime with a negative full-Asai endpoint exponent.

An asymmetric tensor does not repair the $q = 11$ sign after descent. If $\mathbf{a} = (a_\sigma)$ is any nonnegative multiplicity vector and

$$\ell(\mathbf{a}) = -a_1 + \sum_{\sigma \neq 1} a_\sigma \chi_\sigma, \qquad A = \sum_\sigma a_\sigma,$$

then transitivity of Galois on the embeddings gives the exact identity

$$\sum_{\tau \in \mathrm{Gal}(K_q/\mathbb{Q})} \ell(\tau \mathbf{a}) = -A\Delta_q.$$

When $\Delta_q < 0$, at least one conjugate component expands. Hence no flat endpoint height obtained only from tensor and symmetric powers can have negative formal weight in every descended component. A successful $q = 11$ construction would have to use a changing extension, consecutive incidence data, or another genuinely nonlinear structure. At $q = 7$, those structures are needed only to control algebraic endpoints which avoid the generic $1/5$ gap.

The neutral-corridor lemma gives a second, independent distinction. Let

$$E_q = K_q[t]/(t^2 - \lambda_q t - 1), \qquad u(X, Y) = \frac{tX + Y}{\iota(t)X + Y}.$$

The two rows are eigenvectors of the alternating two-step matrix, and

$$u(W_q P) = t^4 u(P), \qquad \iota(u) = u^{-1}.$$

For $q = 7$, $E_7 = K_7 \times K_7$. Selecting either split character leaves three signed logarithms, whose sum is zero; Rosen admissibility fixes the central sign and proves (1059). For $q = 11$, E_{11} is a totally real quadratic field over K_{11}. Above each base embedding, the two logarithms are a_σ and $-a_\sigma$, so their combined height after m pairs is

$$\max(0, a_\sigma + m b_\sigma) + \max(0, -a_\sigma - m b_\sigma) = |a_\sigma + m b_\sigma|.$$

There is no signed cancellation after the two extensions are assembled into one algebraic point. Central Rosen admissibility controls one of the five absolute values, while the other four may grow linearly. This shows why the useful $q = 7$ corridor compression is caused by the split torus, not merely by low degree.

The comparison suggests the following unified conjecture. After quotienting consecutive dynamics by its alternating norm-one-torus translations, the sign of the remaining Arakelov drift is the sign of the Hodge defect Δ_q. On the positive-defect side $q = 3, 5, 7, 9$, stopped drift should be nonpositive and the rational classification should follow. Starting at $q = 11$, negative defect permits an escaping trace-field orbit. The split factorization at seven is extra proof structure: it turns the neutral part of this conjecture into a one-line product-formula inequality. The exact $q = 11$ tests below probe the escaping side of the same statement.

Experiment II.22.1 (Exact $q = 11$ Rosen comparison)**.** The program `rosen_p11_comparison.py` propagates the homogeneous recurrence

$$(p_{j+1}, q_{j+1}) = (-q_j, p_j - a_j \lambda q_j)$$

exactly in $\mathbb{Z}[\lambda_{11}]^2$. Branch digits and all archimedean diagnostics are evaluated at five times the decimal coefficient length plus one hundred guard digits. Projective repetition is checked by exact cross multiplication in the quintic field.

For $q = 7$, all integers $1 \le x \le 100$ terminate within 61 steps. For $q = 11$, with cutoff 500, only $x = 1$ terminates and none of $x = 2, \dots, 100$ repeats. Starting at $x = 2$, 3000 steps give no exact repetition of period at most 500. The exact pair coefficients reach 1024 decimal digits. Winsor's logarithmic height of x_j/λ grows from 5.733065 at step zero to 2270.281759 at step 2999, with least-squares slope 0.726271 on the second half. The smallest observed distance to a half-integer branch boundary is $8.59 \cdot 10^{-5}$. There are 2989 blocks of length (12) in these 3000 digits, and all 2989 are distinct. The final digit suffix occurs earlier only through length (4). This block-complexity check is still finite evidence, but it distinguishes the run from a short period hidden by a long preperiod and is reproduced by the standard report of the same program.

The integer line is not special. A second exact census took all

$$x = a_0 + a_1\lambda_{11} + \cdots + a_4\lambda_{11}^4 \neq 0, \qquad a_i \in \{-1, 0, 1\}.$$

Of these 242 elements, exactly

$$\pm 1, \quad \pm\lambda_{11}, \quad \pm(1 + \lambda_{11}), \quad \pm(\lambda_{11} - 1)$$

terminate by step 300, none becomes projectively periodic with period at most 250, and the remaining 234 are still unresolved at the cutoff. Thus the long behavior of $x = 2$ is already the dominant behavior in the smallest full quintic coefficient box, not an isolated integer anomaly. This finite statement is reproduced by `rosen_p11_comparison.py -ring-box 1 -census-steps 300`.

The scalar height misses the main feature. Put

$$u_{m,j} = \log|\sigma_m(q_j)| - \log|\sigma_1(q_j)|, \qquad m = 2, 3, 4, 5,$$

where $m = 1$ is the uniformizing embedding. This is the full rank-four unit-shape vector of the exact denominator. At the midpoint, step 1499, it is

$$(2701.373, 2346.897, 2013.760, 1689.370),$$

and by step 2999 its maximum absolute coordinate is 5406.095. Its componentwise least-squares slopes on the second half are

$$(1.78720, 1.56499, 1.32739, 1.10638).$$

These are close to the Hodge prediction

$$(1 + \chi_2, 1 + \chi_3, 1 + \chi_4, 1 + \chi_5) = \left(\frac{16}{9}, \frac{14}{9}, \frac{12}{9}, \frac{10}{9}\right).$$

Indeed, the denominator is central-stable with exponent -1 and outer-unstable with exponents χ_m, so its shape ratios see $1 + \chi_m$. The abstract negative relative Hodge lines instead see $-1 + \chi_m$. The proposed stopped transport has to pass between these two different states; it is not a formal use of negative degree. The same exponent model predicts the norm exponent $-1 + \sum_{m=2}^{5} \chi_m = 7/9$. Solving for the five individual slopes from the four differences and the radial slope gives

$$(-1.01195, 0.77525, 0.55305, 0.31544, 0.09443).$$

In central-time units, the outer slopes are $(0.76610, 0.54652, 0.31172, 0.09331)$, while the radial slope is 0.71765. These are the finite-run versions of $(7/9, 5/9, 3/9, 1/9)$ and $7/9$, respectively. At

the same step $\log|\mathrm{N}(q_j)| = 1146.122$. Thus the denominator is not radially growing with bounded shape: all four unit directions escape linearly. The projective Hodge differences

$$\eta_{1,j} - \eta_{m,j}, \qquad \eta_{m,j} = \log \frac{\max\{|\sigma_m(p_j)|, |\sigma_m(q_j)|\}}{|\sigma_m(q_j)|},$$

remain at one-excursion scale: their maximum radius is 9.051, and their final radius is 2.093. The 300-step runs for $x = 3, 4, 5$ show the same separation: the largest shape coordinate lies between 542 and 547, while the logarithmic denominator norm lies between 219 and 255.

An initial run at insufficient precision falsely reported a period at step 196. Exact coefficients at that stage undergo severe cancellation in the distinguished embedding. Increasing precision from "coefficient length plus 80" to "five times coefficient length plus 100" removed the false branch and the false period. This failure is retained because it is a concrete warning: high-degree unit imbalance makes fixed-precision Rosen tests unreliable.

A second and subtler precision failure survived that correction. The high-precision roots returned by the branch routine were later evaluated under Python's default 15-digit `mpmath` context. The exact branch sequence was unchanged, but the height at step 1499 was falsely reported as 5745.916240, and the unit shape appeared bounded. The 512-digit coefficients undergo hundreds of digits of cancellation at the conjugate roots. Wrapping the height and shape evaluations in the same working precision gives the corrected values above. This is why the present program carries the precision context through every embedding calculation, rather than merely storing high-precision root objects.

There is an exact published control which predates these experiments. Hanson–Merberg–Towse–Yudovina compute the action modulo 2 at $q = 11$ and prove that λ_{11}^2 is not in the cusp orbit [62, Example 3]. Thus its Rosen expansion is known not to terminate. The same paper reports 10 000 nonterminating digits for λ_{11}^4, but does not prove that either point has trivial stabilizer. Their conjecture is the stronger statement that every $q > 9$ admits a trace-field point with nonperiodic expansion [62, Remark 7]; they also report that a search among monomials $a\lambda_q^n$ found no hyperbolic fixed point for odd $11 \le q \le 29$ [62, Remark 9].

The option `-hmty-controls` in `rosen_p11_comparison.py` follows the two published $q = 11$ test points with the same exact recurrence used above. Through 500 steps neither λ_{11}^2 nor λ_{11}^4 repeats with period at most 250. Their denominator unit-shape slopes are respectively

$$(2.038, 1.761, 1.511, 1.293), \qquad (2.074, 1.871, 1.578, 1.253),$$

and the corresponding radial norm slopes are 0.939 and 1.053. The cutoff is too short to interpret the discrepancy from the asymptotic Hodge slopes seen for $x = 2$, but it gives a reproducible exact-state check of the two points singled out in the earlier paper.

This changes the evidentiary status of the phase-transition proposal. The fact that $q = 11$ has noncuspidal field points is a theorem and does not distinguish periodic from aperiodic behavior. The conjectural threshold $q > 9$ in [62], however, was reached without Hodge theory and agrees exactly with the sign computation $\Delta_q < 0 \iff q \ge 11$ for odd q. That agreement is substantive evidence for the scalar-defect conjecture; it is not a proof of its negative half.

Experiment II.22.2 (Exact $q = 9$ control)**.** The program `rosen_q9_control.py` performs the same exact homogeneous recurrence in

$$\mathbb{Z}[\lambda_9], \qquad \lambda_9 = 2\cos(\pi/9), \qquad \lambda_9^3 - 3\lambda_9 - 1 = 0.$$

This is the nearest odd case on the contracting side of the exact threshold. Its Hodge weights and defect are

$$(1, 5/7, 1/7), \qquad \Delta_9 = 1/7.$$

Among the integers $1 \le x \le 100$, sixty-six expansions terminate and thirty-four become exactly projectively periodic; none is unresolved at 500 steps. The largest terminating time is 109, for $x = 55$, and the last periodic return in the sample is detected by step 107.

For $x = 2$, the first nine digits are

$$1, -4, 1, 1, 3, -4, 1, 1, 3.$$

Writing (p_j, q_j) for the exact state, direct multiplication in the cubic order gives

$$(p_9, q_9) = (9 + 14\lambda_9 - 10\lambda_9^2)(p_5, q_5), \qquad \mathrm{N}(9 + 14\lambda_9 - 10\lambda_9^2) = 1.$$

Thus this orbit has preperiod five and period four. No numerical recognition is used in the return identity.

Outcome. The test is evidence for, not a proof of, the defect criterion. Degree and defect still covary in these smallest cases, so it does not isolate the two variables. It does show that fast closing persists for the composite case immediately below the exact threshold, while the first case above the threshold exhibits the opposite numerical behavior.

The first interpretation of the *corrected* experiment is a Hodge-sign phase transition: $\Delta_q > 0$ would predict closing for every trace-field point, while $\Delta_q < 0$ would predict genuinely nonclosing orbits, with $x = 2$ the candidate at $q = 11$. It remains a hypothesis, but the linear escape of all four relative conjugate sizes now supports it rather than being invisible to the calculation.

The conditional calculation is exact. If every nonterminating, nonperiodic trace-field point is Bowen–Margulis generic for the Rosen natural extension, then its denominator exponent vector is bounded above by

$$(-1, \chi_2, \dots, \chi_g).$$

For $\Delta_q > 0$, the integral Asai exponent is negative and lattice discreteness forces closing. For $\Delta_q < 0$, the same argument allows growth. This suggests the uniform conjecture that positive defect gives termination or periodicity for every trace-field point, while negative defect permits at least one infinite aperiodic trace-field orbit. The positive half is the heptagonal theorem when $q = 7$; the negative half is tested first by $q = 11, x = 2$.

II.22.1 Reassessment: the natural rank is $[K : \mathbb{Q}] - 1$

Let $p = 2g + 1$ be prime, and let $\mathcal{L}_m$, $1 \le m \le g$, be the Hodge lines with normalized degrees $\chi_m = (p - 2m)/(p - 2)$. The natural relative bundle is

$$\mathcal{E}_p^{\mathrm{pair}} = \bigoplus_{m=2}^{g} \left(\mathcal{L}_m \mathcal{L}_1^{-1}\right).$$

It has rank $g - 1$, which is exactly the rank of the unit group of the totally real field. Every summand has strict Hodge contraction

$$1 - \chi_m = \frac{2(m-1)}{p-2} > 0,$$

and the total normalized gap is

$$\sum_{m=2}^{g}(1-\chi_m)=\frac{g(g-1)}{2g-1}.$$

Thus the full relative bundle has no sign transition. What changes with p is its rank.

This rank is forced by units. On the unit logarithm hyperplane $\sum_{m=1}^{g} t_m = 0$, the map

$$(t_1,\dots,t_g)\mapsto(t_2-t_1,\dots,t_g-t_1)$$

is an isomorphism. Its inverse is

$$t_1=-\frac{1}{g}\sum_{m=2}^{g}u_m,\qquad t_m=u_m-\frac{1}{g}\sum_{j=2}^{g}u_j.$$

Hence all $g-1$ pairwise coordinates are necessary and sufficient to control unit shape. Their determinant, a single scalar sum, has a kernel of dimension $g-2$ and is not proper.

There is a subtle finite-place correction to this statement. If one tensors the $g-1$ pairwise Hodge lines into their determinant, the homogeneous state is

$$\frac{X^{\otimes(g-1)}\otimes\bigotimes_{j=2}^{g}\sigma_j(X)}{v^{g-2}}.$$

Thus the moving incidence divisor occurs to the power $g-2$. At $g=3$ this happens to be one; at $g=5$ it is three. The determinant line is therefore not the desired one-divisor generalization.

For odd g, the full Asai exterior square repairs this without an ad hoc cancellation. If $W=\bigotimes_{j=1}^{g}V_j$, then

$$\bigwedge^{2}W=\bigoplus_{|S|\ \mathrm{odd}}\left(\bigotimes_{j\in S}\bigwedge^{2}V_j\right)\otimes\left(\bigotimes_{j\notin S}\operatorname{Sym}^{2}V_j\right).$$

For two primitive columns with determinant r, the coefficient ideal of the S-component is

$$\mathfrak{c}_S\prod_{j\in S}(\sigma_j r),$$

where the fixed factor $\mathfrak{c}_S$ is supported above 2 and comes from saturating the symmetric/alternating splitting. Away from 2 it is the unit ideal, and it contributes only a bounded Arakelov constant. The singleton components therefore have moving ideals $(\sigma_j r)$, up to those fixed factors, while the all-factor component, which occurs because g is odd, is the scalar $\mathrm{Nm}(r)$. Projectivizing the singleton components gives all $g-1$ unit-shape ratios, and the scalar component pays the radial norm once. This proves that the required finite carrier exists at both $p=7$ and $p=11$. It does not prove the stopped Hodge estimate on that carrier.

For $p=11$, the four pairwise gaps are

$$\frac{2}{9},\quad\frac{4}{9},\quad\frac{6}{9},\quad\frac{8}{9},$$

with total $20/9$. The Higgs ratios $\Xi_m=\theta_m/\theta_1$ vanish to orders $m-1$ at the order-eleven orbifold point, so the determinant of the rank-four bundle has the effective divisor

$$(1+2+3+4)P_{11}=10P_{11}.$$

This is the exact degree-five analogue of the heptagonal rank-two bundle and its divisor $3P_7$.

The new unifying candidate is therefore a rank-$(g-1)$ one-divisor fusion statement on this explicit Asai-wedge carrier: its stopped maximal-slope vector must retain a fixed proportion of every Hodge contraction until incidence vanishes or the affine frame repeats. For $p = 7$ this is the current rank-two stopped problem. For $p = 11$ it is rank four. Naively assigning one copy of the determinant ideal to each channel cancels the apparent gain and is not the conjecture.

This leaves two testable predictions. The scalar-sign hypothesis predicts that $x = 2$ at $p = 11$ never closes. The unit-rank fusion hypothesis predicts that it eventually terminates or repeats, perhaps after a very long transient. The 3000-step computation cannot prove which occurs, but it now distinguishes their mechanisms: the exact denominator escapes linearly in every unit direction, whereas the pointwise Hodge-proximity differences stay at excursion scale. Therefore a unit-rank proof would need a genuinely nonlocal stopped-transport theorem. Merely placing the denominator in the negative-degree relative bundle is false. The theoretical value of $p = 11$ is that it detects exactly this confusion in any proposed proof for $p = 7$.

II.22.2 The codifferent H_4 reduction: a sharper scalar target

The intermediate-group calculation contains one structure which was not used in the earlier numerical audit. Put

$$M_4 = \mathbb{Z} + 4\lambda^{-1}\mathbb{Z} = \mathbb{Z} + 4(\lambda - \lambda^2)\mathbb{Z}.$$

For $\kappa = a + b\lambda + c\lambda^2$, the free coordinate of $\mathcal{O}_K/M_4 \cong \mathbb{Z} \oplus \mathbb{Z}/4\mathbb{Z}$ is

$$d_4(\kappa) = b + c.$$

This coordinate is intrinsic. If

$$f(X) = X^3 - X^2 - 2X + 1, \qquad \vartheta = \frac{\lambda^2 + \lambda - 2}{7},$$

then exact reduction gives

$$(\lambda^2 + \lambda - 2)f'(\lambda) = 7\lambda, \qquad \vartheta = \frac{\lambda}{f'(\lambda)}.$$

Since λ is a unit and $\mathcal{O}_K = \mathbb{Z}[\lambda]$, one obtains

$$\mathfrak{D}_K^{-1} = \vartheta\mathcal{O}_K, \qquad \mathrm{Tr}(\vartheta) = 0, \qquad \mathrm{Nm}(\vartheta) = \frac{1}{49}.$$

The Newton sums give

$$\mathrm{Tr}(\vartheta\lambda) = \mathrm{Tr}(\vartheta\lambda^2) = 1,$$

and hence

$$d_4(\kappa) = \mathrm{Tr}(\vartheta\kappa).$$

Moreover the conjugates of ϑ have signs $+, -, -$. Thus d_4 is the primitive codifferent coordinate with the same central-versus-two-outer orientation as the heptagonal Hodge defect. It is not numerically equal to the logarithmic Hodge character, but it is the canonical integral linear functional pointing in that cone. Its kernel and the H_4 module are

$$\ker d_4 = \mathbb{Z} + \mathbb{Z}(\lambda - \lambda^2), \qquad M_4 = \mathbb{Z} + 4\mathbb{Z}(\lambda - \lambda^2).$$

The four-state label is exactly the index-four difference between the canonical trace hyperplane and the congruence-neutral module.

Route II.22.3 (Prime-family codifferent hyperplane)**.** Let $p = 2g+1 \geq 7$ be prime, let f_p be the minimal polynomial of $\lambda_p = 2\cos(\pi/p)$, and put $\vartheta_p = \lambda_p / f_p'(\lambda_p)$. The residue trace identities give, exactly,

$$\operatorname{Tr}\left(\vartheta_p \sum_{j=0}^{g-1} a_j \lambda_p^j\right) = a_{g-2} + a_{g-1}.$$

Consequently the primitive codifferent hyperplane has the basis

$$1, \lambda_p, \ldots, \lambda_p^{g-3}, \lambda_p^{g-2} - \lambda_p^{g-1}.$$

For every $N \geq 1$, normalization by

$$M_{p,N} = \mathbb{Z} + N\big(\mathbb{Z}\lambda_p + \cdots + \mathbb{Z}\lambda_p^{g-3} + \mathbb{Z}(\lambda_p^{g-2} - \lambda_p^{g-1})\big)$$

therefore has the exact quotient

$$\mathcal{O}_{K_p} / M_{p,N} \simeq \mathbb{Z} \oplus (\mathbb{Z}/N\mathbb{Z})^{g-2}.$$

This is a genuine uniform reduction: the correction lattice has rank $g-1$, but after a codifferent-hyperplane normalization it has only one unbounded coordinate and $g-2$ finite labels.

The real orientation is not uniform. Ordering the roots by $\lambda_{p,j} = 2\cos((2j-1)\pi/p)$, one has

$$\operatorname{sgn}\sigma_j(\vartheta_p) = \begin{cases} (-1)^{j-1}, & j \leq \lceil g/2 \rceil, \\ (-1)^{j}, & j > \lceil g/2 \rceil. \end{cases}$$

Hence $p = 7$ is the unique prime with $g \geq 3$ for which the normal has one positive sign and all outer signs negative. At $p = 11$,

$$d_{11}(a_0 + \cdots + a_4\lambda_{11}^4) = a_3 + a_4, \qquad \operatorname{sgn}(\vartheta_{11}) = (+,-,+,+,-),$$

and

$$M_{11,N} = \mathbb{Z} + N\mathbb{Z}\lambda_{11} + N\mathbb{Z}\lambda_{11}^2 + N\mathbb{Z}(\lambda_{11}^3 - \lambda_{11}^4).$$

Outcome. The algebraic rank reduction succeeds uniformly; the Hodge alignment succeeds only at $p = 7$. A natural $p = 11$ intermediate cusp candidate is generated by G_{11} and $T(\lambda_{11} M_{11,N})$, but no congruence-neutral level or exact cusp intersection has been established here. Even if that group step is proved, its remaining scalar quotient is not the uniformizing-versus-all- outer Hodge direction. Thus the $p = 11$ test favors a genuine phase change over a routine rank-four extension of the H_4 proof.

Route II.22.4 (The Euclidean extremals are a ramified root of the codifferent)**.** Let

$$f(X) = X^3 - X^2 - 2X + 1, \qquad \mathfrak{p} = (\lambda + 2), \qquad y = \frac{f'(\lambda)}{7}.$$

There is an exact chain of fractional ideals

$$(7) = \mathfrak{p}^3, \qquad \mathfrak{D}_K = \mathfrak{p}^2, \qquad (y) = \mathfrak{p}^{-1}, \qquad \mathfrak{D}_K^{-1} = \mathfrak{p}^{-2},$$

and the generators satisfy the stronger element identity

$$y^2 = (\lambda - 1)\frac{\lambda}{f'(\lambda)} = (\lambda - 1)\vartheta.$$

Indeed, modulo 7 one has $f(X) \equiv (X+2)^3$, while $|\mathrm{Nm}(\lambda+2)| = 7$; hence 7 is totally ramified. The monogenic different is $(f'(\lambda))$, of norm 49, and direct reduction by $f(\lambda) = 0$ gives

$$f'(\lambda)^3 = 49\lambda(\lambda - 1).$$

This ideal identity explains the six norm-1/7 Euclidean classes without a search. The quotient $\mathfrak{p}^{-1}/\mathcal{O}_K$ has seven elements. Because $\lambda \equiv -2 \pmod{\mathfrak{p}}$ has order six in $\mathbb{F}_7^\times$, its nonzero elements are

$$\lambda^j y + \mathcal{O}_K, \qquad 0 \le j \le 5.$$

Every lift of such a class has $\mathfrak{p}$-valuation -1, so its absolute norm is at least $1/7$; the displayed lift has norm exactly $1/7$. Together with the known global upper bound $M(K) = 1/7$, these are six global extremal classes.

For the prime family $p = 2g+1$, the corresponding ideal calculation is

$$(p) = \mathfrak{p}^g, \qquad \mathfrak{D} = \mathfrak{p}^{g-1}, \qquad \left(\frac{f_p'(\lambda_p)}{p}\right) = \mathfrak{p}^{-1}, \qquad \left(\frac{f_p'(\lambda_p)}{p}\right)^{g-1} = \mathfrak{D}^{-1}.$$

Thus the Euclidean fractional line is a square root of the codifferent for $p = 7$, but a fourth root for $p = 11$. This is the same rank change, two to four, predicted by the pairwise Hodge bundle.

In fact the exact Euclidean input also survives at 11. Put $y_{11} = f_{11}'(\lambda_{11})/11$. The residue of λ_{11} is -2, of order five in $\mathbb{F}_{11}^\times$, so

$$\pm\lambda_{11}^j y_{11} + \mathcal{O}_{K_{11}}, \qquad 0 \le j \le 4,$$

are the ten nonzero classes in $\mathfrak{p}^{-1}/\mathcal{O}_{K_{11}}$. Every lift has norm at least $1/11$, while the displayed lift has norm exactly $1/11$. Cerri's exact calculation $M(K_{11}) = 1/11$, recorded in [13], therefore makes all ten classes global extremals.

Outcome. The factor $1/7$ is not merely a favorable numerical constant. Its six equality classes form one canonical ramified fractional-ideal quotient, and the square of their carrier is the same codifferent line which defines the H_4 normal coordinate. This supplies an exact finite-place counterpart to the doubled-central Hodge construction. What is still missing is a metric or cocycle identity which makes this root relation telescope along consecutive Rosen returns. The $p = 11$ comparison now rules out a weaker Euclidean constant as the explanation for the different behavior: the first proved difference is the archimedean Hodge sign and orientation.

Route II.22.5 (The three $p = 7$ walls are the Eisenstein A_2 complex)**.** The single H_4 quotient is one member of a Galois orbit. For $\kappa = a + b\lambda + c\lambda^2$, define

$$\ell_0 = b + c, \qquad \ell_1 = b, \qquad \ell_2 = c$$

and

$$M_0 = \mathbb{Z} + 4\mathbb{Z}(\lambda - \lambda^2), \quad M_1 = \mathbb{Z} + 4\mathbb{Z}\lambda^2, \quad M_2 = \mathbb{Z} + 4\mathbb{Z}\lambda.$$

Each quotient is $\mathcal{O}_K/M_i \simeq \mathbb{Z} \oplus \mathbb{Z}/4\mathbb{Z}$, with free coordinate ℓ_i. The Galois action $(b, c) \mapsto (-b-c, b)$ cycles the three unoriented walls.

The corresponding overgroups

$$H_i = \langle G_7, T(4\lambda v_i)\rangle, \qquad (v_0, v_1, v_2) = (\lambda - \lambda^2, \lambda^2, \lambda),$$

are all strict congruence-neutral enlargements: their new translations belong to the principal level-four kernel, so they have the same finite-adelic closure as G_7. Distance from the H_i cusp plane at a compact return is comparable to $\log(1 + |\ell_i(\kappa)|)$.

The exact identity

$$b^2 + bc + c^2 = \frac{(b+c)^2 + b^2 + c^2}{2}$$

therefore gives

$$\log(1 + b^2 + bc + c^2) = 2 \max_i \log(1 + |\ell_i|) + O(1).$$

Outcome. The rank-two correction and the adaptive Hodge singular-gap problem have the same A_2 wall complex. The original H_4 normal coordinate $d = b + c$ is one wall normal; its unbounded k-fibre is motion along that wall and becomes transverse to one of the other two. This does not yet bound cumulative wall backtracking, but it replaces the apparently unrelated d- and k-payments by one Galois-symmetric Coxeter geometry. A successful stopped proof must control a path switching among these three walls, not force a fixed scalar sign.

Every correction has the exact decomposition

$$\kappa = n + 4k\lambda^{-1} + d\lambda + t(\lambda - \lambda^2), \qquad d, k, n \in \mathbb{Z}, \quad t \in \{0, 1, 2, 3\}.$$

For a correction written in the program's basis $\kappa = a + b\lambda + c(\lambda^2 - 1)$, these coordinates are

$$d = b + c, \qquad t \equiv -c \pmod 4, \qquad k = -\frac{c + t}{4}, \qquad n = (a - c) - 8k.$$

Define the stopped charge

$$\begin{aligned}\chi_4(\kappa) = {} & \log(1 + |d|) + \log(1 + |k|) + \mathbf{1}_{t \neq 0} \log 2 \\ & + \mathbf{1}_{d=k=t=0} \log(1 + |n|).\end{aligned}$$

The first term is the Hodge-transverse displacement from H_4; the second is the rational cusp fibre H_4/G_7. The finite label costs a fixed amount. Only when all three vanish does the ambient correction lie on the original integer Rosen line, where the ordinary central Green coordinate pays n.

Route II.22.6 (Canonical codifferent correction budget)**.** At each strict norm record choose a factor-1/7 Euclidean correction κ_j. Prove

$$\sum_{j \in \mathcal{R}} \chi_4(\kappa_j) \leq 6h([r_0 : r_{-1}]) + C_7.$$

This statement implies the full classification. Indeed a record correction is nonzero, so its charge is at least $\log 2$. The bound permits only finitely many records; all remainder norms are then bounded, and the proved fixed-support S-unit theorem closes the orbit.

This is presently the most concrete replacement for the rank-two fusion sentence. It asks for two geometrically different rank-one payments, not for a scalar inequality pretending that the unit plane has disappeared. Its unresolved part is telescoping: the transverse codifferent distance and the rational cusp distance must be charged along the original labelled Rosen path. Passing to arbitrary words in H_4 loses that order and creates dense real shortcuts.

Experiment II.22.7 (Codifferent correction stress test)**.** The program `rosen_euclidean_correction_tests.py` now evaluates χ_4 exactly. In three random batches with coefficient bounds $10^2, 10^4, 10^6$,

it followed 900 primitive states for 75,121 exact Rosen steps and found 1,149 norm-record corrections. Every selected correction satisfies the factor-1/7 norm inequality exactly, and the radius-five coefficient search missed no correction in these runs. The largest values were

coefficient bound	steps	records	$\max(\sum \chi_4 - 6h)$
10^2	13,664	393	22.987391
10^4	24,213	358	−17.675641
10^6	37,244	398	−36.673820

The negative values in the last two rows reflect that uniform sampling from a large coefficient box rarely selects low-height states; they are not evidence for a negative constant. A separate stress test of 781 states of the form $\lambda^a(\lambda^2-1)^b+s$, with $-6 \le a, b \le 6$ and $-2 \le s \le 2$, produced 652 record corrections and maximum excess 16.307205.

The experiment rejects the zero-constant form but has not found growth in the additive error. It is evidence only: the Euclidean search is finite, and terminating or periodic samples cannot test an a priori unbounded number of records. Unlike the former reservoir experiments, however, this charge remains bounded on the balanced unit families that made the proposed recharge estimates diverge.

Route II.22.8 (Rejected continuant/product-of-quotients proof)**.** One possible proof of Route II.22.6 would identify the successive pairs (d_j, k_j) with partial quotients of a reduced rank-one path. A classical continuant inequality would then bound the product of $1+|d_j|$ and $1+|k_j|$, giving the desired sum of logarithms. Exact record corrections do not have the required cone or orientation property.

In 500 exact test orbits, 196 had at least two record corrections. Among those, the nonzero d-coordinates changed sign in 121 orbits, the nonzero k-coordinates changed sign in 68, and the orientation

$$d_j k_{j+1} - k_j d_{j+1}$$

changed sign in 57. This is not a large-coefficient artifact. Take

$$x = 1 - 3\lambda - 3\lambda^2.$$

In the program basis $(1, \lambda, \lambda^2-1)$ this is $(-2,-3,-3)$. Its first strict records occur after digits -8 and 5, with exact record norms 7 and 13. In both cases an ambient Euclidean correction terminates immediately:

	κ	(d,k,t,n)
first record	$-7+2\lambda+\lambda^2$	$(3,-1,3,1)$
second record	$10-3\lambda^2$	$(-3,0,3,10)$.

Thus even two consecutive record corrections can reverse the canonical codifferent direction.

The conclusion is structural. The correction budget cannot telescope through a commutative net displacement or a product of ordinary partial quotients. A successful proof must retain the conjugations by the intervening G_7-blocks. Equivalently, it needs a weighted length in the relative correction cocycle, or a Bass–Serre/relative-Morse model in which sign reversals in the scalar projection do not cancel path length. This is exactly the ordered thin-orbit information absent from the static Arakelov class.

II.22.3 Fubini–Study/Bott–Chern audit of the one-divisor pencil

The two logarithmic-derivative sections Ψ_x, Ψ_y have the same zero divisor $D = 3P_7$. Put

$$\mathbf{\Psi}_{x,y} = (\Psi_x, \Psi_y), \qquad A = \log|\Psi_y/\Psi_x|.$$

With the orthogonal metric on the direct sum,

$$\log\|\mathbf{\Psi}_{x,y}\|^2 = \log\|\Psi_x\|^2 + \log(1+e^{2A}).$$

If $\mathcal{L}$ is their automorphy line and $\omega_{\mathrm{FS}} = dd^c\log(1+|z|^2)$, then

$$dd^c\log\|\mathbf{\Psi}_{x,y}\|^2 = -c_1(\mathcal{L},h) + \delta_{3P_7}(\Psi_y/\Psi_x)^*\omega_{\mathrm{FS}}.$$

This is the hoped-for clean ledger: the common order-seven divisor is counted once, and the projective ratio contributes a positive form.

Route II.22.9 (Outcome of the Fubini–Study audit)**.** The identity does not supply the required linear stopped gain. If $A_n - A_m = H$, its soft-maximum endpoint contribution is

$$\log(1+e^{2A_n}) - \log(1+e^{2A_m}).$$

Taking $A_m = -2H$, $A_n = -H$ makes this tend to zero while the controlled gain equals H. This is intrinsic: Fubini–Study geometry compactifies the multiplicative ratio, so a long excursion toward zero or infinity can have bounded projective cost. A singular toric maximum does not repair this directed example: relative to the first coordinate its endpoint term is $2\max(0,A)$, which is again zero at both ends.

The divisor diagnosis, however, changes. Locally write $\Psi_x = au_xe, \Psi_y = au_ye$, where $a = 0$ cuts out the common divisor. The sup metric satisfies

$$\log\max(\|\Psi_x\|^2, \|\Psi_y\|^2) = \log|a|^2 + \log\|e\|^2 + \max(\log|u_x|^2, \log|u_y|^2).$$

Its curvature is

$$-c_1(\mathcal{L},h) + \delta_{3P_7} + dd^c\max(\log|u_x|^2, \log|u_y|^2),$$

and the last term is the positive wall current on $|u_x| = |u_y|$. Thus the common Hodge divisor is counted once. The same factorization at a finite place pulls the common determinant Δ out of the sup norm once.

Corrected outcome. The earlier claim that a toric maximum necessarily duplicates the two divisors was false. The tropical metric has the correct one-divisor ledger, but it sees only the currently dominant side of a wall. What remains is a directed wall-backtracking estimate for the consecutive orbit. This matches the Eisenstein A_2 formulation above and is narrower than the original Fubini–Study objection.

II.22.4 Two-time period determinant and the exact holonomy audit

The one-time Wronskian compares two period evaluations and their first derivatives. I tested a different carrier which compares the same two period evaluations at the two ends of a stopped interval. This removes the moving-affine-centre problem completely on the real side.

Let $s < t$ lie on the lifted central geodesic. Parallel transport identifies the flat fibres along the segment. At the i-th real place choose polarization-normalized Hodge vectors $\widehat{\omega}_i(s), \widehat{\omega}_i(t)$, and put

$$\mathcal{E}_i^{m,n}(s,t) = \det\begin{pmatrix} \ell_{m,i}(\widehat{\omega}_i(s)) & \ell_{n,i}(\widehat{\omega}_i(s)) \\ \ell_{m,i}(\widehat{\omega}_i(t)) & \ell_{n,i}(\widehat{\omega}_i(t)) \end{pmatrix}.$$

Alternating algebra gives the exact factorization

$$\mathcal{E}_i^{m,n}(s,t) = \sigma_i(\Delta_{m,n})\big(\widehat{\omega}_i(s) \wedge \widehat{\omega}_i(t)\big).$$

If $d_i = d_{\mathbb{H}}(Z_i(s), Z_i(t))$, the determinant-one Hodge normalization gives

$$\frac{|\mathcal{E}_i^{m,n}(s,t)|}{|\sigma_i(\Delta_{m,n})|} = \frac{|Z_i(s) - Z_i(t)|}{\sqrt{\operatorname{Im} Z_i(s) \operatorname{Im} Z_i(t)}} = 2\sinh(d_i/2).$$

This identity is independent of the two rational vectors after their minor $\Delta_{m,n}$ has been removed.

The cyclic separation is

$$\mathfrak{T}(s,t) = 2\log\left(2\sinh\frac{d_0}{2}\right) - \log\left(2\sinh\frac{d_1}{2}\right) - \log\left(2\sinh\frac{d_2}{2}\right).$$

Write $L = t - s$ and

$$L_i = \int_s^t \|DF_i(\gamma(u))\|_{\mathrm{hyp}}\, du.$$

Then $d_0 = L$ and $d_i \le L_i$. Since

$$\frac{d}{dx}\log(2\sinh(x/2)) = \frac{1}{2}\coth(x/2) \ge \frac{1}{2},$$

one obtains the proved inequality

$$\mathfrak{T}(s,t) \ge \frac{1}{2}(2L - L_1 - L_2) = \frac{1}{2}\mathfrak{C}(s,t).$$

If an outer endpoint distance is zero, the left side is $+\infty$. Thus this kernel retains one half of the whole Schwarz–Pick defect on the actual stopped segment. No shadowing of a rational geodesic and no dynamic-to-controlled affine shear is involved.

Route II.22.10 (Finite-place outcome of the two-time determinant)**.** The construction does not automatically carry the record divisor. Let τ generate the cubic Galois group. The three algebraic coefficient factors are

$$\Delta, \qquad \tau(\Delta), \qquad \tau^2(\Delta).$$

The contracting cyclic quotient has coefficient

$$\frac{\Delta^2}{\tau(\Delta)\tau^2(\Delta)} = \frac{\Delta^3}{\mathrm{Nm}(\Delta)},$$

whose field norm is one. It records only the unit shape of Δ. On the other hand, the product of the three determinants has coefficient $\mathrm{Nm}(\Delta)$, and therefore counts the radial divisor once, but its real factor is

$$\prod_{i=0}^{2} 2\sinh(d_i/2),$$

which expands instead of measuring the cyclic defect.

This is a useful positive and negative result. Positively, the two-time kernel solves the real transport problem exactly. Negatively, its contracting scalar loses the norm divisor, while its radial scalar loses the contraction. The remaining assertion can be written in the particularly short sufficient form

$$\log|\operatorname{Nm}\Delta_{m,n}| + \eta\,\mathfrak{T}(t_m, t_n) \le C$$

for consecutive stopped returns before recurrence. Because $\mathfrak{T} \geq \mathfrak{C}/2$, this would prove the classification. The calculation does not prove this inequality: it shows that the inequality is again a one-divisor fusion statement, now for an explicit two-time kernel whose analytic transport is already understood.

Route II.22.11 (The tempting cube root fails on the coincidence divisor)**.** There is an apparent way around the preceding finite-place failure. The cyclic coefficient is

$$\frac{\Delta^3}{\mathrm{Nm}(\Delta)}.$$

After multiplication by $\mathrm{Nm}(\Delta)$, it is a cube. Moreover, the fundamental Hodge relation $\mathcal{N} \simeq \mathcal{N}_2^{\otimes 3}$ says that the corresponding two-point Hodge line is also a cube. If the two-time cyclic *section* were a cube, its cube root would have coefficient Δ, exactly the desired one-divisor normalization.

It is not a cube. Near the pair (P_7, P_7), use orbifold-cover coordinates u, v. The three period maps have local degrees $1, 3, 2$, so their two-point determinants are, up to units,

$$K_0 = u - v, \qquad K_1 = u^3 - v^3, \qquad K_2 = u^2 - v^2.$$

Consequently

$$\frac{K_0^2}{K_1 K_2} = \frac{\text{unit}}{(u+v)(u^2+uv+v^2)}.$$

It has simple poles on the three off-diagonal coincidence branches

$$u = -v, \qquad u = \zeta_3 v, \qquad u = \zeta_3^2 v.$$

Its divisor is not divisible by three, so it has no local meromorphic cube root. The line-bundle cube root exists, but the particular section selected by the two-time evaluation does not lie in the image of the cube map.

The failed cube has the canonical factorization

$$\frac{K_0^2}{K_1 K_2} = \left(\frac{K_0}{K_2}\right)\left(\frac{K_0}{K_1}\right),$$

whose local denominators are respectively $u+v$ and u^2+uv+v^2. Along the diagonal these become the inverse fundamental Higgs section and its square, in agreement with $\Xi_1 = \Xi_2^2$. Away from the diagonal the quadratic factor is not the square of the linear factor. The rank-two pairwise carrier is therefore forced by the off-diagonal coincidence geometry, even though the infinitesimal Hodge divisor has a single fundamental generator.

Along one real segment avoiding the coincidence branches, one may still take a positive real cube root of the norm. That operation is not holomorphic and has no integral transition law around the branches. It therefore cannot be used as the missing Arakelov section. Retaining the three branches separately brings back a vector-valued carrier. This is a geometric, rather than merely notational, reason that the scalar approaches repeatedly return to rank-two fusion.

Route II.22.12 (Local resultant fusion and its exact global obstruction)**.** The two off-diagonal factors look more promising if they are treated as a Koszul pair. Put

$$L = u + v, \qquad Q = u^2 + uv + v^2.$$

Then

$$Q - uL = v^2, \qquad Q - vL = u^2,$$

and hence

$$(L, Q) = (u + v, u^2) = (u + v, v^2) \quad \text{in } \mathbb{C}\{u, v\}.$$

The common zero scheme has length two and is supported at $(u, v) = (0, 0)$. In projective tangent coordinate this is the numerical identity

$$\mathrm{Res}_r(r + 1, r^2 + r + 1) = 1.$$

Thus the linear and quadratic coincidence branches have no common off-diagonal tangent direction. This is a real local fusion statement: there is no second hidden divisor which would force the determinant ideal to be counted twice.

The attempted globalization fails for a group-theoretic reason which can be proved exactly. Let

$$A_3 = \begin{pmatrix} \alpha & \alpha \\ \lambda & \alpha \end{pmatrix}, \qquad \alpha = \lambda^2 - 1.$$

This element is hyperbolic in the distinguished factor and elliptic of infinite order in both outer factors. Its joint outer powers lie in a compact product of two rotation circles, so there are $n_j \to \infty$ for which

$$\rho_1(A_3)^{n_j}, \rho_2(A_3)^{n_j} \longrightarrow 1.$$

For any fixed $z \in \mathbb{H}$, equivariance now gives

$$F_i(A_3^{n_j} z) \longrightarrow F_i(z) \quad (i = 1, 2),$$

while

$$d_{\mathbb{H}}(z, A_3^{n_j} z) = n_j \ell_0(A_3) + O_z(1) \longrightarrow \infty.$$

In invariant norms the two divided differences are

$$\mathscr{D}_i(z, w) = \frac{2 \sinh(d(F_i(z), F_i(w))/2)}{2 \sinh(d(z, w)/2)}.$$

Both $\mathscr{D}_i(z, A_3^{n_j} z)$ tend to zero. Therefore a global Bézout identity between the two divided differences cannot have bounded coefficient norms: along this sequence its coefficients must diverge.

Outcome. The local resultant correctly fuses the ramification branches, but the thin group has arbitrarily deep simultaneous outer near-returns. Transporting the local Koszul identity around those returns creates exactly the unbounded coefficient height which the stopped Arakelov inequality is meant to control. Moreover u, v are complex period coordinates, not algebraic integers, so resultant one is not a coprimality theorem for the finite ideal $(\Delta_{m,n})$. This route does not prove the one-divisor inequality, but it identifies the obstruction more sharply than the earlier statement that the section is merely "not a cube."

Route II.22.13 (Chebyshev–Baker audit of the coherent cyclic obstruction)**.** The global obstruction above uses powers of one fixed element. This allows an exact calculation which was missing from the compactness argument. For $h \in \mathrm{SL}_2(K)$, put

$$t = \mathrm{tr}(h), \qquad C_0 = 0, \quad C_1 = 1, \quad C_{n+1} = tC_n - C_{n-1}.$$

Cayley–Hamilton gives

$$h^n = C_n(t)h - C_{n-1}(t)I, \qquad \det(Q, h^nQ) = C_n(t)\det(Q, hQ).$$

If the central image is hyperbolic of translation length ℓ_0, then

$$|C_n(\sigma_0 t)| = \frac{\sinh(n\ell_0/2)}{\sinh(\ell_0/2)}.$$

If the i-th outer image is elliptic, write $\sigma_i(t) = 2\cos\theta_i$. Then

$$C_n(\sigma_i t) = \frac{\sin(n\theta_i)}{\sin\theta_i}.$$

For the element A_3, both outer rotations have infinite order. Therefore $e^{2i\theta_i}$ is algebraic and is not a root of unity. Baker–Wüstholz [55] gives

$$|\sin(n\theta_i)| \geq c_i n^{-C_i}.$$

It follows, exactly, that

$$\log|\operatorname{N} C_n(t)| = \frac{n\ell_0}{2} + O(\log n), \qquad \log|\operatorname{N}\det(Q, h^nQ)| = \frac{n\ell_0}{2} + O_{h,Q}(\log n).$$

The same estimate quantifies the metric obstruction. For a Hodge point which is not an outer elliptic center,

$$\mathfrak{T}(z, h^n z) = n\ell_0 + O(\log n), \qquad \log \mathscr{D}_i(z, h^n z) = -\frac{n\ell_0}{2} + O(\log n).$$

Thus the Bézout coefficients needed to globalize the local resultant grow like $\exp(n\ell_0/2)$, up to a power of n. The global cost is not an uncontrolled artifact of the recurrence subsequence: it is precisely half the central height.

There is also a fixed-target Roth consequence. The attracting fixed point ξ^+ of h has degree two over K: its discriminant is $t^2 - 4$, which is positive centrally and negative at each elliptic outer embedding, and hence cannot be a square in the totally real field K. If z_n remains in a compact set away from the repelling fixed point and $y_n = h^n z_n \in K$, then

$$|y_n - \xi^+| \ll e^{-n\ell_0}.$$

Roth over K therefore gives, for distinct y_n,

$$n\ell_0 \leq (2+\epsilon)h_K(y_n) + O(1), \qquad h_K = [K:\mathbb{Q}]h_{\mathrm{abs}}.$$

So a repeated fixed hyperbolic block cannot be supercritical relative to endpoint height. This does not improve the constant at the critical H^{-2} scale.

Outcome. The coherent cyclic channel is now understood. Baker shows that simultaneous outer returns of one word are only polynomially close in the power, and Roth excludes supercritical rational shadowing of that word's fixed quadratic axis. The obstruction left for the Rosen algorithm is more specific: successive blocks vary, so the quadratic target and its height vary with the approximant, and critical H^{-2}-scale blocks remain allowed. A useful moving-target theorem would have to exploit that these targets are consecutive fixed points of the thin-group coding, not arbitrary algebraic numbers.

This calculation is uniform in the odd triangle parameter. In particular, it behaves the same way for $p = 7$ and $p = 11$, with four outer rotation factors instead of two in the latter case. Hence fixed cyclic powers cannot explain the $p = 11$ phase change. The distinction between 7 and 11 occurs for genuinely varying words: the scalar defect is $1/5$ at 7 and $-7/9$ at 11, and the exact $p = 11$ orbit displays the four corresponding unit-shape drifts.

Route II.22.14 (Global fixed-form resultant fusion)**.** The local period-coordinate resultant has no integral finite-place meaning. The fixed-point forms do. For

$$g = \begin{pmatrix} u & v \\ w & z \end{pmatrix} \in \mathrm{SL}_2(\mathcal{O}_K), \qquad Q_g(X, Y) = wX^2 + (z - u)XY - vY^2,$$

direct evaluation of the Sylvester determinant gives the exact identity

$$\mathrm{Res}(Q_g, Q_h) = \mathrm{tr}([g, h]) - 2. \tag{1062}$$

If P is primitive, every common ideal divisor of $Q_g(P)$ and $Q_h(P)$ divides the right side of (1062). Indeed, in a local unimodular frame with $P = (1, 0)$, the two evaluations are the leading coefficients. Every term in the resultant contains at least one of them. Thus the global Bézout coefficient is not mysterious: it is the single commutator divisor.

For two consecutive Rosen branches this simplifies sharply. Put

$$B_a = \begin{pmatrix} 0 & -1 \\ 1 & -a\lambda \end{pmatrix}, \qquad v_{j+1} = B_{a_j} v_j, \qquad D_j = \det(v_j, v_{j+1}).$$

Pulling the next fixed-point form back to the v_j-frame gives

$$\begin{aligned} Q_{B_a}(X, Y) &= X^2 - a\lambda XY + Y^2, \\ Q_{B_a^{-1} B_b B_a}(X, Y) &= X^2 - (2a - b)\lambda XY + \{1 + a(a - b)\lambda^2\}Y^2. \end{aligned}$$

Consequently

$$D_{j+1} - D_j = (a_{j+1} - a_j)\lambda (v_j)_2 (v_{j+1})_2, \tag{1063}$$

$$\mathrm{Res}(Q_{B_{a_j}}, Q_{B_{a_j}^{-1} B_{a_{j+1}} B_{a_j}}) = \lambda^2 (a_{j+1} - a_j)^2. \tag{1064}$$

Primitivity gives $(D_j, (v_j)_2) = (D_j, (v_{j+1})_2) = \mathcal{O}_K$. Therefore

$$\min\{v_{\mathfrak{p}}(D_j), v_{\mathfrak{p}}(D_{j+1})\} \le v_{\mathfrak{p}}(a_{j+1} - a_j)$$

at every finite prime. This is the desired finite fusion at one step: the common part of two incidence divisors is paid globally by the one principal divisor of the digit difference. No prime-by-prime choice is involved.

The exterior-square packaging is

$$Q_g \wedge Q_h \in \bigwedge^2 \mathrm{Sym}^2(K^2)^*.$$

Under the three-dimensional cross-product identification, its discriminant is the resultant above. This looked like a possible new integral Hodge carrier. The exact $p = 11, x = 2$ test rejects the naive version. Over 300 exact steps, in embedding order $1, 3, 5, 7, 9$, the fixed-form slopes are

$$(1.08971, 0.64142, 0.13420, 0.39555, 0.88854),$$

whereas the consecutive-wedge slopes are

$$(2.17880, 1.28160, 0.26684, 0.79034, 1.77476).$$

After subtracting the two form sizes, the five wedge-angle slopes are all within 0.0011 of zero. Thus consecutive fixed-form lines do not converge in the required projective metric; the wedge simply carries the sum of their growth. The calculation is exact until the final archimedean logarithms and is reproduced by `fixed_form_resultant_tests.py`.

Outcome. The route succeeds at the finite places and fails at a precisely identified archimedean step. For raw consecutive branches, the divisor overlap is completely controlled by the cusp digit. For long Hodge-stopped blocks, the corresponding overlap is controlled by $\operatorname{tr}([g,h])-2$. What remains is a thin-orbit bound for the height of this commutator divisor after parabolic acceleration. A static exterior square has no extra negative Hodge character, so that bound must use the consecutive stopped geometry, not merely the representation $\bigwedge^2 \operatorname{Sym}^2$.

Route II.22.15 (Normalized commutator as the fixed-axis cross ratio)**.** The commutator divisor admits a second exact factorization which explains what its height is measuring. If the fixed points of g are α_g, β_g, those of h are α_h, β_h, and

$$t = \frac{(\alpha_g - \alpha_h)(\beta_g - \beta_h)}{(\alpha_g - \beta_h)(\beta_g - \alpha_h)},$$

then

$$\frac{\operatorname{tr}([g,h])-2}{\{\operatorname{tr}(g)^2-4\}\{\operatorname{tr}(h)^2-4\}} = \frac{t}{(1-t)^2}. \tag{1065}$$

This follows by writing each fixed-point form as a product of its two linear factors. The resultant is the product of the four cross differences; division by the two discriminants leaves the right side of (1065). The identity is unchanged if either fixed-point pair is reversed.

At the uniformizing real place, if the two axes meet at angle θ,

$$\operatorname{tr}([g,h])-2 = -4\sinh^2(\ell_g/2)\sinh^2(\ell_h/2)\sin^2\theta.$$

For disjoint axes one replaces θ by i times their distance. Thus the global commutator height has two conceptually different parts: the two spectral discriminants, controlled by matrix/Hodge growth, and the normalized separation of the two axes, controlled only if the words are known to be consecutive blocks of one coded geodesic.

There is a useful whole-group theorem at this point. If M belongs to the Hecke reflection group and its distinguished trace is hyperbolic, then

$$|\sigma(\operatorname{tr} M)| \le |\sigma_0(\operatorname{tr} M)|$$

for every field embedding. Bugeaud–Hubert–Schmidt use this trace dominance to bound the heights of Rosen convergents [3, Lemma 3.1 and the preceding trace theorem]. Elliptic and parabolic elements have all conjugate traces bounded by two. Applying this to $M = [g,h]$ gives the unconditional global estimate

$$\log|\operatorname{N}(\operatorname{tr}([g,h])-2)| \le [K_q:\mathbb{Q}]\,\log\bigl(|\sigma_0(\operatorname{tr}([g,h])-2)|+4\bigr). \tag{1066}$$

Thus the commutator route does not require separate alignment estimates at all embeddings. It requires one sharp two-sided axis-alignment estimate in the uniformizing plane.

The exact program `rosen_commutator_trace_audit.py` tests this remaining quantity on consecutive norm-record interval words. For the rational input

$$x = \frac{4971}{407} - \frac{2084}{407}\lambda - \frac{96}{11}(\lambda^2 - 1),$$

which terminates after 235 steps, the record triples

$$(0, 9, 12),\ (9, 12, 14),\ (12, 14, 29)$$

give

$$\frac{\log|\mathrm{N}\, R|}{\ell(g) + \ell(h)} = 1.95024, \quad 1.30757, \quad 1.96238.$$

The first resultant is exactly

$$42730861497044 + 76750922714184\lambda + 95878061444880(\lambda^2 - 1),$$

with absolute norm 247046007883811069131693977664. The largest outer-to-central commutator-trace ratio in the three tests is only 0.0754. Thus trace dominance works extremely well, but the central commutator itself can grow at essentially the full two-block scale even on a closing orbit.

Outcome. The all-place part succeeds: the global resultant height is reduced to the uniformizing commutator by (1066). The naive central bounded-angle assertion fails badly at finite record blocks. A usable theorem must be stopped and asymptotic, allowing a constant depending on the initial algebraic endpoint, and must show that persistent full-scale central commutators force recurrence of the complete frame. This is a more precise obstruction than the former phrase "all-embedding alignment." It also leaves a genuinely non-scalar uniform possibility at $q = 11$: trace dominance is valid there despite the wrong scalar Hodge sign. The symbolic resultant and cross-ratio identities are checked by `fixed_form_resultant_tests.py`.

Route II.22.16 (Parabolic rays turn trace dominance into fixed-form coefficient dominance)**.** There is a stronger use of the same Hecke-group trace theorem. It controls not just a commutator but every coefficient of the fixed-point equation for every element of the full thin group.

For a cusp column $v = (p, q)^t$, put

$$N_v = \begin{pmatrix} -pq & p^2 \\ -q^2 & pq \end{pmatrix}, \qquad U_v = I + \lambda N_v.$$

If $v = A(1, 0)^t$ with $A \in G_7$, then $U_v = ATA^{-1} \in G_7$. For an arbitrary $M \in G_7$, the elementary rank-one calculation

$$\mathrm{tr}(MN_v) = \det(v, Mv) = Q_M(v)$$

gives

$$\mathrm{tr}(MU_v^n) = \mathrm{tr}(M) + n\lambda Q_M(v). \tag{1067}$$

For large $|n|$, the central trace is hyperbolic unless $Q_M(v) = 0$. Apply trace dominance to MU_v^n, divide by $|n|$, and pass to the limit. One obtains

$$|\sigma_i(\lambda Q_M(v))| \le |\sigma_0(\lambda Q_M(v))| \tag{1068}$$

at every outer embedding.

This is useful because three fixed group cusps determine a binary quadratic. Write

$$Q_M(X, Y) = cX^2 + eXY - bY^2.$$

The columns $(1,0)^t, (0,1)^t, (\lambda,1)^t$ are the first columns of I, S, TS, and the corresponding values are

$$c, \qquad -b, \qquad \lambda^2 c + \lambda e - b.$$

Solving for the coefficients after applying (1068) shows that

$$\|\sigma_i Q_M\|_{\mathrm{coeff}} \le C_{q,i} \|\sigma_0 Q_M\|_{\mathrm{coeff}}. \tag{1069}$$

No word decomposition occurs in this proof. The only deep input is the Hecke-group trace-dominance theorem; the parabolic-ray identity and the three-cusp interpolation are elementary.

Let the trace field have degree d, let $H_0(Q_M) = \|\sigma_0 Q_M\|_{\mathrm{coeff}}$, and fix a primitive integral target P. If $Q_M(P) \neq 0$, the product formula and (1069) give

$$1 \le |\operatorname{Nm} Q_M(P)| \le C_q(P) |\sigma_0 Q_M(P)| H_0(Q_M)^{d-1}. \tag{1070}$$

Hence little-oh of H_0^{-2} forces exact closing for $q = 7$, while little-oh of H_0^{-4} is required for $q = 11$. This makes the earlier objection about constants precise. A bare $O(H_0^{-2})$ estimate at $q = 7$, with no control of its constant, is exactly critical and does not force vanishing.

The exact program `rosen_parabolic_ray_audit.py` tests the most immediate candidate for the missing little-oh: interval matrices between successive norm-record prefixes. For the same terminating rational input used in the commutator audit, the record steps are

$$0,\ 9,\ 12,\ 14,\ 29.$$

For the four successive interval matrices, the observed exponents

$$-\frac{\log |\sigma_0 Q_M(P)|}{\log H_0(Q_M)}$$

are

$$0.78224, \qquad 0.88352, \qquad 0.92767, \qquad 0.95354.$$

They approach the ordinary one-power stable-axis scale, not the required two-power threshold. The corresponding values of $\log |\sigma_0 Q_M(P)| + 2 \log H_0(Q_M)$ are

$$21.84886, \qquad 39.83891, \qquad 43.28910, \qquad 70.66512,$$

so the naive one-form completion is decisively not present in these blocks.

Conjugating to the moving reduced frame gives a more useful conclusion. If $X_m = g_m P$, $X_n = g_n P$, and $B = g_n g_m^{-1}$, then

$$Q_B(X_m) = \det(X_m, X_n).$$

Equation (1069) removes all outer coefficient norms of Q_B; the only outer quantities left are the two sizes of the transported state X_m. Explicitly,

$$|\operatorname{Nm} Q_B(X_m)| \le C_7 |\sigma_0 Q_B(X_m)| H_0(Q_B)^2 \prod_{i=1}^{2} \|\sigma_i X_m\|^2.$$

This is the same rank-two Arakelov wall reached from the Wronskian, but with a real simplification: the varying quadratic form itself is now controlled at every outer place by a theorem about the full thin group. What remains is only the outer height of the transported rational state.

There is a natural intermediate group attached to this exact obstruction. For the target P, let

$$\mathcal{U}_P = \{I + tN_P : t \in K\}, \qquad \Gamma(P,t) = \langle G_7, I + tN_P \rangle$$

with t chosen so that the second generator is integral. For every group cusp column v,

$$\operatorname{tr}\big((I + tN_P)U_v^{-1}\big) - 2 = t\lambda \det(P,v)^2. \tag{1071}$$

If P itself is a group cusp, both parabolics lie in G_7, and trace dominance applies. Otherwise the mixed product belongs only to $\Gamma(P,t)$. The intermediate-group dichotomy shows that its central projection is dense as soon as the adjoined unipotent is not already in G_7, so applying the discrete Hecke-group trace theorem to the mixed trace would be a logical error.

This target-dependent group nevertheless identifies the correct equality case. The intersection $G_7 \cap \mathcal{U}_P$ is nontrivial exactly in the cusp case. A nonparabolic element $h \in G_7$ normalizes $\mathcal{U}_P$ exactly when $hP = \varepsilon P$; then

$$h(I + sN_P)h^{-1} = I + \varepsilon^2 sN_P,$$

which is precisely the special hyperbolic alternative. Thus a relative trace theorem would need the following dichotomy: failure of dominance for the mixed traces in (1071) is either paid by the Hodge gap, forcing the integral incidence to vanish, or is coherent under this normalizer. This is the stopped Arakelov equality problem in target-unipotent language, not a new consequence of the elementary trace identity.

Outcome. The route proves a new all-place reduction and explains the $q = 7$ versus $q = 11$ difference by field degree: one needs one central power for every outer embedding. It does not prove the needed central surplus. Raw interval recurrence gives approximately one power. At $q = 7$ one further power must come from the stopped Hodge–thin-orbit constraint; at $q = 11$ three further powers would be needed. This rules out a degree-independent *same Roth estimate for every q* conjecture and points instead to a rank-$(d-1)$ fusion or relative height statement.

Route II.22.17 (Adjacent cusp parabolics: no fixed height, but a split torus special to $q = 7$)**.** The parabolic formula suggests trying to remove the moving fixed-point form altogether. This can be done algebraically. Let v, w be cusp columns, put

$$r_v = \det(P,v), \qquad r_w = \det(P,w), \qquad D = \det(v,w),$$

and set $M = U_v U_w^{-1}$. Since $N_v x = v \det(v,x)$, direct expansion gives

$$Q_M(P) = \lambda(r_w^2 - r_v^2) + \lambda^2 D\, r_v r_w. \tag{1072}$$

If (v,w) is an oriented group frame, then $D = 1$ and

$$M = A \begin{pmatrix} 1+\lambda^2 & \lambda \\ \lambda & 1 \end{pmatrix} A^{-1}, \qquad A = (v\ w).$$

Moreover

$$A^{-1}P = (r_w, -r_v)^t,$$

so (1072) is simply

$$Q_M(P) = \lambda F(r_w, r_v), \qquad F(X,Y) = X^2 + \lambda XY - Y^2.$$

The opposite orientation changes the sign of the middle term. Thus the entire conjugacy class is represented by one fixed quadratic form in the two moving remainder coordinates. This identity is

elementary but structurally attractive: exact vanishing gives either a cusp degeneration or one of the special hyperbolic lines.

It does not create a decreasing arithmetic height. Under one Rosen transition

$$(X, Y) \longmapsto (-a\lambda X - Y, X),$$

one has

$$F(-a\lambda X - Y, X) = \{a^2\lambda^2 - a\lambda^2 - 1\}X^2 + (2a - 1)\lambda XY + Y^2,$$

which is not a scalar multiple of $F(X, Y)$. Conjugating the matrix to a fixed matrix has merely transferred the moving-frame cost into the coordinates (r_w, r_v).

The exact program `rosen_adjacent_parabolic_form_audit.py` tests both orientation forms

$$F_\pm(X, Y) = X^2 \pm \lambda XY - Y^2$$

at strict remainder-norm records. On 200 deterministic random primitive inputs, there are 236 record rows. The largest value of

$$\log |\operatorname{Nm} F_\pm(r_j, r_{j-1})| - 6h(P)$$

is

$$18.5859291950.$$

It occurs at step 58 for

$$P = \left[-9411 + 2023\lambda + 2572(\lambda^2 - 1) : 20\right].$$

At the same transition the next remainder satisfies

$$\log |\operatorname{Nm} r_{j+1}| - 6h(P) = -12.5185734142.$$

Thus the fixed quadratic can grow far beyond the initial height budget even while the record remainder is well below that budget.

There is, however, a second fact which is special to the heptagonal field and was hidden by taking the norm of the whole quadratic. Put

$$\alpha = \lambda^2 - 1, \qquad \alpha^{-1} = \alpha - \lambda, \qquad s = \alpha + \alpha^{-1} = 2\lambda^2 - \lambda - 2.$$

The cubic equation for λ gives

$$s^2 = \lambda^2 + 4, \qquad \operatorname{Nm}_{K_7/\mathbf{Q}}(s) = -13,$$

and hence

$$F_+(X, Y) = (X + \alpha Y)(X - \alpha^{-1}Y). \tag{1073}$$

The two factors are the incidences with the two K_7-rational eigenlines of the fixed adjacent-parabolic word. Their coefficient determinant is s, so for a primitive pair the two factor ideals are coprime away from the single fixed divisor above 13. This is not a height estimate; it is an exact split-torus structure.

The amended exact audit records the two factor norms separately. On the same 236 strict norm-record rows, the largest value of

$$\max_{\epsilon=\pm1} \min_{L|F_\epsilon} \log |\operatorname{Nm} L(r_j, r_{j-1})| - 6h(P)$$

is -12.3593608694. Thus, in this sample, for either orientation at least one special-eigenline incidence stays inside a fixed budget even though the product F_ϵ can be much larger. A larger 2500-input stress test, with coefficients up to 10^6, again found no violation for the + orientation. This observation must not be overstated. Wider coefficient-box tests found a largest excess of 1.65 in the smallest box, followed by $-2.82, -3.43, -7.49, -15.36$ in boxes of radii $10^2, 10^3, 10^4, 10^5$. Thus the literal estimate with additive constant zero is false. The data are compatible with an estimate of the form $6h(P) + O_7(1)$, but only for the *absolute norms* of the factors. Such a norm records the finite divisor and loses the unit distribution among the three real embeddings. It is therefore not yet the Arakelov inequality required by the proof.

The reason that the factors nevertheless interact naturally with the actual Rosen dynamics is exact. For the recurrence

$$r_{j+1} = -a_j \lambda r_j - r_{j-1},$$

put

$$\begin{aligned} A_j &= r_j + \alpha r_{j-1}, \quad B_j = r_j - \alpha^{-1} r_{j-1}, \\ C_j &= r_j - \alpha r_{j-1}, \quad D_j = r_j + \alpha^{-1} r_{j-1}. \end{aligned}$$

Then $F_+ = A_j B_j$ and $F_- = C_j D_j$. If $a_j = 1$, direct substitution gives

$$A_{j+1} = \alpha^{-1} C_j, \qquad B_{j+1} = -\alpha D_j, \tag{1074}$$

and if $a_j = -1$,

$$C_{j+1} = -\alpha^{-1} A_j, \qquad D_{j+1} = \alpha B_j. \tag{1075}$$

For a general digit the four errors are exactly

$$\begin{aligned} A_{j+1} - \alpha^{-1} C_j &= -(a_j - 1)\lambda r_j, & B_{j+1} + \alpha D_j &= -(a_j - 1)\lambda r_j, \\ C_{j+1} + \alpha^{-1} A_j &= -(a_j + 1)\lambda r_j, & D_{j+1} - \alpha B_j &= -(a_j + 1)\lambda r_j. \end{aligned}$$

Consequently an alternating $1, -1$ block is precisely the channel in which the two split characters are transported by units. Its two-step matrix is

$$-W_7 = -\begin{pmatrix} 1 + \lambda^2 & \lambda \\ \lambda & 1 \end{pmatrix},$$

and it scales the two eigenline incidences by $-\alpha^{-2}$ and $-\alpha^2$. This is a genuine use of the thin group: the exceptional digit language is the normalizer channel of the special K_7-split torus. The elementary calculation identifies that channel; it does not prove that a general stopped word is forced into it.

There is a decisive negative control on using the absolute factor norm alone. The program `rosen_split_torus_double_coset_audit.py` uses the exact free-product presentation

$$G_7 = \langle S, R \mid S^2 = R^7 = 1 \rangle, \qquad W_7 = SRSR^{-1}, \qquad e = (\alpha, 1),$$

and enumerates reduced words g. After quotienting exactly by the left and right actions of $\langle W_7 \rangle$, the cumulative numbers of double cosets satisfying $|\operatorname{Nm} \det(e, ge)| \le B$ are

maximum word length	0	2	4	6	8	10
$B = 1$	1	5	11	17	33	87
$B = 13$	1	18	54	112	262	510
$B = 100$	1	18	87	298	684	1528.

This finite enumeration does not prove that a bounded-norm set contains infinitely many double cosets. It does show that bounded norm is violently non-proper at the tested scales even after the obvious normalizer drift has been removed. The source is visible: units can have norm one while their three real sizes diverge. Any successful stopping theorem must therefore bound a *metrized* incidence at all real places, not merely its principal ideal norm.

The $q = 11$ comparison now gives a genuinely new answer. Let

$$K_{11} = \mathbf{Q}(\lambda_{11}), \qquad F_{11}(X, Y) = X^2 + \lambda_{11}XY - Y^2.$$

Taking the resultant with the quintic minimal polynomial gives

$$\mathrm{Nm}_{K_{11}/\mathbf{Q}}(\lambda_{11}^2 + 4) = 5741, \tag{1076}$$

and 5741 is prime. A square in K_{11} has square absolute norm, so $\lambda_{11}^2 + 4$ is not a square and F_{11} is irreducible over K_{11}. In particular, the fixed product of two adjacent parabolics has no K_{11}-rational eigenline. The exact program `rosen_adjacent_parabolic_p11_audit.py` found 1724 strict record rows on 30 deterministic integral inputs; the largest value of

$$\log|\,\mathrm{Nm}\, F_{11,\epsilon}(r_j, r_{j-1})| - 10h(P)$$

was 610.9565903907. There is no quintic linear factor whose norm could play the heptagonal stopped-factor role.

The uniform object is instead the quadratic K_p-algebra

$$E_p = K_p[t]/(t^2 - \lambda_p t - 1), \qquad F_p(X, Y) = \mathrm{Nm}_{E_p/K_p}(X + tY). \tag{1077}$$

The eigenvalues of the adjacent-parabolic word are t^2, t^{-2}. For $p = 7$, $E_7 \simeq K_7 \times K_7$, and (1073) is the resulting pair of characters of a split torus. For $p = 11$, E_{11} is a quadratic field extension of K_{11}; the two characters are not separately defined over the trace field.

Passing to E_p does not reproduce the heptagonal minimum of two character norms. In the nonsplit case the two conjugate factors satisfy

$$\left|\mathrm{N}_{E_p/\mathbf{Q}}(X + tY)\right| = \left|\mathrm{N}_{K_p/\mathbf{Q}}\, F_p(X, Y)\right| = \left|\mathrm{N}_{E_p/\mathbf{Q}}(X + \iota(t)Y)\right|. \tag{1078}$$

Thus they carry the same finite divisor size. Their distinction lies only in the distribution of their real sizes, which is an Arakelov metric rather than another algebraic norm. This explains structurally, and not just numerically, why the split S-unit reduction is special to seven.

In fact this splitting is isolated in the whole odd-prime family. If P_n is the Pell sequence $P_0 = 0, P_1 = 1$, $P_{n+2} = 2P_{n+1} + P_n$, a cyclotomic product gives

$$\mathrm{Nm}_{K_p/\mathbf{Q}}(\lambda_p^2 + 4) = P_p. \tag{1079}$$

Indeed, with $v = 1 + \sqrt{2}$, the right side is $(v^p + v^{-p})/(2\sqrt{2})$ for odd p. Ljunggren's classification of square Pell numbers says that only P_0, P_1, P_7 are squares [1, 2]. Since a square in K_p has square absolute norm, E_p can split for an odd prime only at $p = 7$, where the explicit factorization shows that it does. Thus the heptagonal split torus is not a low-degree feature which should be expected again at a larger prime.

The quadratic algebra also gives a genuinely uniform all-place coordinate. Let ι be its involution and define

$$u_p(X, Y) = \frac{X + tY}{X + \iota(t)Y}, \qquad \mathrm{Nm}_{E_p/K_p}\, u_p(X, Y) = 1. \tag{1080}$$

Thus $u_p(X,Y)$ is a K_p-point of the norm-one torus

$$\mathcal{T}_p = \mathrm{Res}^1_{E_p/K_p}\,\mathbf{G}_m.$$

Its ordinary Weil height is an all-place quantity: it includes both the finite divisor and the logarithmic imbalance at every real embedding. When $p=7$, formula (1080) is the ratio of the two linear factors in (1073). When $p=11$, it remains defined over K_{11} as a point of $\mathcal{T}_{11}$, although neither numerator nor denominator is a K_{11}-character. This is the correct uniform replacement for the split-factor language.

It would be wrong to conjecture that this torus height is bounded for every prime: the exact $p=11, x=2$ orbit has four escaping unit-shape coordinates. The family-level conjecture suggested by the comparison is instead a phase transition controlled by the scalar Hodge defect. For $p=2g+1$, put

$$\Delta_p = \frac{-g^2+4g-2}{2g-1}.$$

The trace-field orbit should be forced to terminate or repeat when $\Delta_p > 0$; when $\Delta_p < 0$, at least one trace-field orbit should escape in the metrized norm torus. The last prime on the contracting side is $p=7$, and the first prime on the expanding side is $p=11$. Independently, the adjacent algebra is split only at $p=7$. These are not identical criteria: $p=5$ has positive Hodge defect but a nonsplit adjacent torus. More strongly, the Rosen–Leutbecher theorem says that every point of $\mathbf{P}^1(K_5)$ is a cusp [18, 17]. The comparison among $5,7,11$ therefore separates the two mechanisms: positive Hodge defect occurs at both 5 and 7, while adjacent splitting occurs only at 7; at 11 both advantages disappear. The Hodge sign is the plausible family-level criterion, while splitting is an additional structure available only in the $p=7$ proof.

Outcome. The whole fixed quadratic still gives no descent: its growth is the moving unit-shape obstruction in another frame. Its splitting gives a concrete $q=7$-specific structure, and (1074)–(1075) identify its precise thin-group normalizer channel. But the double-coset audit shows why the absolute norm of a factor cannot supply the missing theorem. The heptagonal target is a stopped *metrized* split-character inequality; its uniform formulation is a Hodge-selected height statement on $\mathcal{T}_p$, not a degree-independent Roth estimate. At $p=11$ the torus is nonsplit over K_{11} and the Hodge defect has the opposite sign. The comparison therefore suggests that the heptagonal theorem is the last contracting prime case, rather than the first easy case of a classification valid for every prime.

Route II.22.18 (The split-factor estimate reduces to four sparse unit orbits)**.** The negative double-coset audit does not make the split factors useless. It identifies the exact theorem that a factor-norm estimate would leave. Keep

$$A = r+\alpha s, \quad B = r-\alpha^{-1}s, \quad C = r-\alpha s, \quad D = r+\alpha^{-1}s,$$

so $F_+ = AB$ and $F_- = CD$.

Proposition II.22.19 (Two bounded split incidences give a unit parametrization)**.** *Fix $B_0 \geq 1$. There is a finite set $\mathcal{C}(B_0) \subset K^\times$ such that every $[r:s] \in \mathbf{P}^1(K)$, with $r,s \in \mathcal{O}_K$, satisfying*

$$\min\{|\,\mathrm{N}\,A|, |\,\mathrm{N}\,B|\} \leq B_0, \qquad \min\{|\,\mathrm{N}\,C|, |\,\mathrm{N}\,D|\} \leq B_0 \tag{1081}$$

has the form

$$[r:s] = \Phi_\mu(c\epsilon) \quad \textit{for some} \quad c \in \mathcal{C}(B_0), \quad \epsilon \in \mathcal{O}_K^\times,$$

where $\mu \in \{AC, AD, BC, BD\}$ *and*

$$\Phi_{AC}(z) = \alpha\,\frac{z+1}{z-1}, \qquad \Phi_{AD}(z) = \frac{\alpha - \alpha^{-1}z}{z-1}, \tag{1082}$$

$$\Phi_{BC}(z) = \frac{\alpha z - \alpha^{-1}}{z-1}, \qquad \Phi_{BD}(z) = \alpha^{-1}\frac{z+1}{1-z}. \tag{1083}$$

The formulas are interpreted projectively at their poles. Moreover,

$$|h([r:s]) - h(\epsilon)| \le C(B_0)$$

for a constant depending only on the displayed finite data.

Proof. Choose one of A, B and one of C, D which satisfy (1081). The four coefficient determinants are

$$\det(A,C) = -2\alpha, \quad \det(A,D) = -\lambda, \quad \det(B,C) = -\lambda, \quad \det(B,D) = 2\alpha^{-1}.$$

Thus every chosen pair determines (r,s); the only nonunit divisor in the inverse matrices lies above 2.

There are only finitely many integral principal ideals of absolute norm at most B_0. Choose one generator for each. If the selected incidences are L, M, write

$$L = \beta u, \qquad M = \gamma v,$$

where β, γ belong to the resulting finite set and $u, v \in \mathcal{O}_K^\times$. Then

$$z = \frac{L}{M} = \frac{\beta}{\gamma}\,\frac{u}{v}$$

belongs to one of finitely many cosets of $\mathcal{O}_K^\times$. Solving the four possible two-by-two systems gives precisely (1082)–(1083). Finally, a fixed degree-one Möbius transformation changes absolute logarithmic height by $O(1)$. There are only finitely many transformations here, including the constants β/γ, which proves the height assertion. □

Proposition II.22.20 (The two split minima are one quadratic divisor)**.** *Put*

$$Q_{AC} = r^2 - \alpha^2 s^2, \qquad Q_{AD} = r^2 + (\alpha + \alpha^{-1})rs + s^2,$$
$$Q_{BC} = r^2 - (\alpha + \alpha^{-1})rs + s^2, \qquad Q_{BD} = r^2 - \alpha^{-2}s^2.$$

Away from the four special eigenlines $ABCD = 0$*, one has the exact identity*

$$\min_{\mu \in \{AC,AD,BC,BD\}} |\,\mathrm{N}\,Q_\mu(r,s)| = \min\{|\,\mathrm{N}\,A|, |\,\mathrm{N}\,B|\}\,\min\{|\,\mathrm{N}\,C|, |\,\mathrm{N}\,D|\}. \tag{1084}$$

Consequently the two-orientation bound (1081) *is equivalent, after replacing* B_0 *by* B_0^2 *in one direction, to*

$$\min_\mu |\,\mathrm{N}\,Q_\mu(r,s)| \le B_0.$$

For a height bound the comparison is

$$\log \max\{m_+, m_-\} \le \log \min_\mu |\,\mathrm{N}\,Q_\mu(r,s)| \tag{1085}$$

$$\le 2 \log \max\{m_+, m_-\}, \tag{1086}$$

where $m_+ = \min(|\,\mathrm{N}\,A|, |\,\mathrm{N}\,B|)$ *and* $m_- = \min(|\,\mathrm{N}\,C|, |\,\mathrm{N}\,D|)$.

Proof. The four displayed forms are respectively AC, AD, BC, BD. Hence the left side of (1084) is

$$\min\{|\,\mathrm{N}\,A|, |\,\mathrm{N}\,B|\}\min\{|\,\mathrm{N}\,C|, |\,\mathrm{N}\,D|\}.$$

Both factors are positive integers away from the four zero loci, so $\max(u,v) \le uv \le \max(u,v)^2$ proves (1086). On a zero locus the orbit has already reached a special fixed line. □

Route II.22.21 (Outcome of the one-divisor rewrite)**.** The phrase "reuse one divisor in two pairwise estimates" can now be made literal in the heptagonal case. There is no need to reuse a divisor: the product of the two required minima is the norm of one of four fixed quadratic incidences. What is still missing is not an algebraic packaging lemma but an Arakelov estimate selecting one of these four forms along every remainder state. The rewrite is useful because a symmetric-square or height argument can now be tested against one scalar inequality, and it is harmless because the subsequent four-term S-unit closing argument still remembers which of the four forms was selected. Proposition II.20.91 further shows that these are the four hyperbolic walls in the nontrivial level-four coset of H_2; the scalar target is therefore an ordered thin-orbit nearest-wall inequality, not an arbitrary minimum over quadratic polynomials.

Lemma II.22.22 (Exact classification of degenerate split-coordinate transitions)**.** *Let*

$$P_{AC} = \begin{pmatrix} 1 & \alpha \\ 1 & -\alpha \end{pmatrix}, \quad P_{AD} = \begin{pmatrix} 1 & \alpha \\ 1 & \alpha^{-1} \end{pmatrix},$$

$$P_{BC} = \begin{pmatrix} 1 & -\alpha^{-1} \\ 1 & -\alpha \end{pmatrix}, \quad P_{BD} = \begin{pmatrix} 1 & -\alpha^{-1} \\ 1 & \alpha^{-1} \end{pmatrix},$$

and $B(n) = \left(\begin{smallmatrix} -n\lambda & -1 \\ 1 & 0 \end{smallmatrix}\right)$. *For* $\mu, \nu \in \{AC, AD, BC, BD\}$, *consider the four-term identity obtained by expressing the shared coordinate in*

$$X = P_\mu(r_j, r_{j-1})^t, \qquad Y = P_\nu(r_{j+1}, r_j)^t = P_\nu B(n) P_\mu^{-1} X.$$

At a nonterminal state, a degenerate identity has one of the two cross pairings. Each pairing determines $z = X_1/X_2$ *as a rational function* $z_{\mu,\nu,\pi}(n) \in K(n)$. *For all* 32 *choices, either*

$$\mathrm{N}_{K/\mathbf{Q}}\, z_{\mu,\nu,\pi}(n)$$

is a nonconstant quotient of polynomials of degree at most three, or $z_{\mu,\nu,\pi}(n)$ *itself is constant. The constant cases are*

μ	ν	z
AC	BD	-1
AD	AD	$\lambda^2+\lambda$
AD	BC	$\lambda^2+\lambda$
BC	AD	$2-\lambda$
BC	BC	$2-\lambda$
BD	AC	-1.

Proof. Write R_μ for the first row of P_μ^{-1} and S_ν for the second row of P_ν^{-1}. The shared-coordinate relation is

$$S_\nu Y - R_\mu X = 0.$$

All four scalar coefficients are nonzero. A cross pairing equates the first term of $S_\nu Y$ with one of the two terms of $R_\mu X$. Substituting $Y = P_\nu B(n) P_\mu^{-1} X$ is one linear equation for $z = X_1/X_2$,

and hence gives a degree-one rational function over $K(n)$. Taking the cubic field norm produces a quotient of cubic polynomials unless cancellation makes it constant.

The remaining assertion is the finite calculation in `rosen_split_sunit_degeneracy_certificate.py`. It performs all arithmetic exactly in

$$\mathbf{Q}(n)[\lambda]/(\lambda^3 - \lambda^2 - 2\lambda + 1),$$

checks both pairings for all sixteen ordered type pairs, asserts that constant norm always implies constant z, and prints the six rows above. No numerical root selection is used. Direct substitution in the six displayed rows supplies a short independent check of the only constant cases. □

Theorem II.22.23 (An all-time split-factor bound proves the classification)**.** *Let $x \in K$ have primitive exact remainders (r_j, r_{j-1}). Suppose there is a constant B_0 such that, at every nonterminal state,*

$$\max\{\min(|\,\mathrm{N}\,A_j|, |\,\mathrm{N}\,B_j|), \min(|\,\mathrm{N}\,C_j|, |\,\mathrm{N}\,D_j|)\} \le B_0. \tag{1087}$$

Then the Rosen orbit terminates or repeats. Consequently x is a cusp or a G_7-translate of a K-rational eigenline of a hyperbolic element.

Proof. If a selected factor vanishes, the state is one of the K-rational eigenlines already supplied by the adjacent special words. Otherwise choose at time j one bounded factor from each orientation and call the resulting coordinate vector $X_j = (L_j, M_j)$, of type μ_j. There are only finitely many principal ideals of norm at most B_0. After choosing generators and adjoining them to the unit group, every coordinate of every X_j belongs to one fixed finitely generated subgroup $\Gamma \subset K^\times$.

For consecutive states, the second coordinate of (r_{j+1}, r_j) equals the first coordinate of (r_j, r_{j-1}). Hence

$$S_{\mu_{j+1}} X_{j+1} - R_{\mu_j} X_j = 0, \tag{1088}$$

with the notation of Lemma II.22.22. This is a fixed homogeneous four-term equation in Γ once the two types and the finite ideal-generator choices are fixed. The projective S-unit theorem leaves only finitely many nondegenerate projective solutions. Each such solution determines the projective state $[r_j : r_{j-1}]$.

A proper three-term vanishing subsum in (1088) would leave a nonzero singleton, and a pairing of the two terms belonging to one time would make the shared remainder r_j vanish. Every nonterminal degeneracy is therefore one of the two cross pairings in Lemma II.22.22. The norm

$$\mathrm{N}(L_j/M_j) = \mathrm{N}(L_j)/\,\mathrm{N}(M_j)$$

belongs to a finite set of rational numbers determined by B_0. In a nonconstant case of the lemma, equating the displayed rational function of the integer digit n_j to one of those finitely many values leaves only finitely many digits. In a constant case, L_j/M_j, and hence the projective state, is already fixed. Thus the degenerate transitions also contain only finitely many projective current states.

The whole nonterminal orbit consequently ranges over a finite set of states. Determinism forces a repetition. The standard closing criterion identifies the intervening word as the required special hyperbolic element. □

This is the first place where the split norm and the thin recurrence combine to give an unconditional closing mechanism. It explains why the proliferation in the double-coset audit is not fatal: arbitrary bounded-norm group points proliferate, but consecutive Rosen states must satisfy the fixed four-term equation (1088).

The exact program `rosen_split_factor_all_states_audit.py` tests (1087) at every state. On 200 deterministic random primitive inputs, comprising all 200 closing orbits and thousands of exact states, the largest value of

$$\log \max\{\min(|\,\mathrm{N}\,A|, |\,\mathrm{N}\,B|), \min(|\,\mathrm{N}\,C|, |\,\mathrm{N}\,D|)\} - 6h(P)$$

was -12.1368907270. On the adversarial unit rays with exponents

$$8, 16, 24, 32, 40, 48, 64, 80, 100, 120,$$

the program checked $180, 426, 594, 1024, 1025, 1404, 1629, 2335, 2449, 3340$ states respectively. Their largest excesses lay between -38.58 and -31.42. These calculations do not prove (1087). They also missed a different and simpler unit direction. Put

$$U = \lambda(\lambda^2 - 1), \qquad x_N = U^{-N}.$$

The same exact program now includes this family. In its first one hundred states it finds

N	10	20	40	60
$\max_j\{E_j - 6h(x_N)\}$	1.2821496	3.5023095	12.4091459	18.2149036,

where E_j is the logarithmic split-factor minimax. Thus the earlier sentence that the $6h$ form survived the adverse unit tests was false: the test family was too narrow. The calculation does not prove that the displayed excess is unbounded, so it does not disprove the existence of some field constant. It is enough to withdraw this quantitative route.

Proposition II.22.19 is unconditional. It says that a stopped estimate for the two orientation minima would place every norm record in four explicit torus orbits. It does *not* put all intermediate remainders in a fixed S, so Theorem II.9.14 cannot yet be applied. The remaining question is sparse return rigidity.

Proposition II.22.24 (Conditional sparse-return rigidity)**.** *Assume the four-term number-field n-conjecture II.9.23. Let j_ν be an infinite sequence of strict Rosen norm records satisfying (1081) with one fixed B_0. After passage to a subsequence, write*

$$x_{j_\nu} = \Phi(\epsilon_\nu), \qquad \epsilon_\nu \in \Gamma,$$

where $\Phi \in \mathrm{PGL}_2(K)$ is fixed and $\Gamma \subset K^\times$ is a fixed finitely generated group. Let $G_\nu \in G_7$ be the intervening Rosen word and write

$$H_\nu = \Phi^{-1} G_\nu \Phi = \begin{bmatrix} a_\nu & b_\nu \\ c_\nu & d_\nu \end{bmatrix}.$$

Set

$$A_\nu = 1 + h([a_\nu : b_\nu : c_\nu : d_\nu]), \qquad Q_\nu = h([\epsilon_\nu\epsilon_{\nu+1} : \epsilon_{\nu+1} : \epsilon_\nu : 1]).$$

If

$$Q_\nu \longrightarrow \infty, \qquad A_\nu = o(Q_\nu), \tag{1089}$$

then the original endpoint is a cusp or a G_7-translate of a K-rational eigenline of a hyperbolic element.

Proof. The transition equation is

$$c_\nu\epsilon_\nu\epsilon_{\nu+1} + d_\nu\epsilon_{\nu+1} - a_\nu\epsilon_\nu - b_\nu = 0. \tag{1090}$$

Its four multiplicative variables belong to the fixed group Γ. Proposition II.9.25 and (1089) imply that every sufficiently late relation (1090) is degenerate.

We classify that degeneracy rather than treating it as a closing conclusion. If all four matrix entries are nonzero, the partitions

$$\{c\epsilon\epsilon', d\epsilon'\} \mid \{-a\epsilon, -b\}, \qquad \{c\epsilon\epsilon', -a\epsilon\} \mid \{d\epsilon', -b\}$$

would each force $ad = bc$, contrary to invertibility. The only possible pairing is

$$c\epsilon\epsilon' = b, \qquad d\epsilon' = a\epsilon.$$

It gives

$$\epsilon^2 = \frac{bd}{ac}.$$

The second equality shows that $h(\epsilon') = h(\epsilon) + O(A_\nu)$; the first then gives $Q_\nu = O(A_\nu)$, contradicting (1089). If exactly three coefficients are nonzero, a degenerate proper subsum leaves a nonzero singleton, which is impossible. Therefore the matrix is eventually diagonal or antidiagonal. Equivalently,

$$G_\nu \in G_7 \cap N_{\mathrm{PGL}_2(K)}\big(\Phi\mathbf{G}_m\Phi^{-1}\big).$$

The group in (II.22.18) is elementary: it is finite or virtually cyclic because G_7 is a discrete Fuchsian group. All selected Rosen prefixes consequently lie in one coset of this elementary subgroup. If the set of prefixes is finite, a state repeats. If it is infinite, their cusp approximants converge both to the original endpoint and to one of the two endpoints of the elementary axis. Those endpoints lie in K, since Φ is defined over K. An infinite intersection contains a hyperbolic element, so the limit is its K-rational eigenline. This gives the stated classification. □

What has and has not been gained. The conditional theorem uses the thin group at the final, essential step: an infinite arithmetic-torus intersection inside a discrete Fuchsian group is elementary and hence special. It also identifies one precise replacement for the stopped Arakelov inequality: (1089). The latter is not supplied by the ordinary factor norm. It compares the height of the entire intervening Rosen word with the unit height of its two record endpoints. Standard fixed-coefficient S-unit theory proves the proposition when the H_ν range over a finite set. Ru–Vojta moving-target theory does not by itself prove the moving-coefficient case, because (1090) makes the points linearly degenerate over the coefficient field. The generalized *abc* input supplies that missing height-uniformity. The remaining geometric task is now concrete: deduce (1089), or a fixed positive-margin variant sufficient for the same argument, from the 1/5 Hodge defect and the strict-record condition.

The slow-word condition is a genuine branch, not a numerical consequence of the split norm. The exact diagnostic `rosen_split_sunit_return_audit.py` computes $\Phi^{-1}G_\nu\Phi$ and verifies its action on the two unit coordinates. On the adversarial unit ray at exponent 80, the three repeated-type record pairs had ratios

$$\frac{h(\Phi^{-1}G_\nu\Phi)}{h([\epsilon_\nu\epsilon_{\nu+1} : \epsilon_{\nu+1} : \epsilon_\nu : 1])} = 0.9059, \qquad 0.9794, \qquad 0.9852.$$

Thus the sharp family approaches the critical ratio one rather than the little-oh regime. Generalized *abc* proves the slow branch but does not finish this critical branch. Theorem II.22.23 avoids that moving-coefficient criticality by using the fixed consecutive-state equation (1088); this is why the all-time version of the factor inequality is presently the better target.

There is an exact geometric reformulation of that target. Put

$$e_A = [-\alpha : 1], \quad e_B = [\alpha^{-1} : 1], \quad e_C = [\alpha : 1], \quad e_D = [-\alpha^{-1} : 1],$$

and let $\Lambda_e(X)$ be the sum over the three real places of the usual logarithmic projective proximity of X to e. The global incidence identity gives

$$\begin{aligned} &\log \max\{\min(|\,\mathrm{N}\,A|, |\,\mathrm{N}\,B|), \min(|\,\mathrm{N}\,C|, |\,\mathrm{N}\,D|)\} \\ &\quad = 3h(X) + 3h(\alpha) - \min\{\max(\Lambda_{e_A}(X), \Lambda_{e_B}(X)), \max(\Lambda_{e_C}(X), \Lambda_{e_D}(X))\}. \end{aligned} \tag{1091}$$

Thus the remaining all-time bound says exactly that a high Rosen state has large total proximity to one endpoint of each of two adjacent special axes. The four choices of endpoints are the four unit loci above. This is the clean interface: Hodge contraction must prove the two-axis proximity covering, and the thin consecutive recurrence turns that covering into the fixed S-unit equation.

The hierarchy with the earlier height conjecture is elementary but important. For every fixed linear factor $L \in \{A, B, C, D\}$,

$$\log |\,\mathrm{N}\,L(X)| \le 3h(X) + O_7(1).$$

Hence the conjectural factor-two excursion bound $h(X_j) \le 2h(P) + O_7(1)$ implies the split minimax bound. It proves far more than is needed. Conversely, the split bound allows $h(X_j) \to \infty$ inside one of the four unit parametrizations $X_j = \Phi(c\epsilon_j)$. The fixed consecutive-state S-unit equation, not Northcott boundedness of each state, is what rules out an indefinite nonclosing journey through those corridors. The Pisot-unit family above also gives a growing tested excess over the constant-two height budget while its sampled orbits still terminate or become periodic. Neither the $6h$ split target nor the factor-two height target is retained here as a plausible global estimate.

Route II.22.25 (Why the standard arithmetic-holonomy theorem does not finish the two-time route)**.** For a fixed real place, the functions

$$f_{m,i} = \ell_{m,i}(\omega_i), \qquad f_{n,i} = \ell_{n,i}(\omega_i)$$

belong to one fixed rank-two Picard–Fuchs module. They are independent exactly when $\Delta_{m,n} \neq 0$. The dimension theorem of Calegari–Dimitrov–Tang bounds the number of independent functions from two inputs: their Taylor-coefficient denominator types and the boundary energy of a common continuation map [58, Theorem 2.5.1]. Neither input depends on the integral coefficient matrix which changes one basis of this fixed rank-two module into another.

Consequently, whenever $\Delta_{m,n} \neq 0$, the standard theorem must return a bound at least two:

$$\frac{\mathscr{G}(\varphi)}{\log |\varphi'(0)| - \tau(\mathbf{b})} \ge 2.$$

The same inequality already holds for any fixed basis of the Picard–Fuchs module. A long thin-group word changes the basis by monodromy; it does not enlarge the continuation domain or improve the denominator type. Therefore the unmodified holonomy theorem cannot see that the coefficient vectors were selected by consecutive Rosen returns, and cannot force their determinant to vanish.

There are three ways a genuine holonomy completion could escape this coefficient-blindness: a word-selected linear combination could overconverge through a singularity; the one-divisor normalization could have a strictly better denominator type than the full module; or one could prove a relative, coefficient-sensitive slope theorem. None of these properties follows from the existing dimension theorem. The first is essentially the desired incidence cancellation, and the third is the stopped Arakelov inequality in another form. Thus arithmetic holonomy remains a possible proof architecture, but citing its standard rank bound is not progress on the missing implication.

II.22.5 Relative-Anosov/Morse audit after the $p = 11$ test

The doubled-central estimate suggests trying a relative Morse lemma for the cusped triangle group. The natural representation for that attempt is

$$\rho_\star = \operatorname{Sym}^2 \rho_0 \oplus (\rho_1 \otimes \rho_2).$$

If $a_i = \log s_1(\rho_i(g))$, its three largest singular logarithms are

$$2a_0, \qquad a_1 + a_2, \qquad |a_1 - a_2|,$$

up to the bounded endpoint change between adapted Hodge norms. Hence the first two simple-root gaps are exactly

$$D(g) = 2a_0 - a_1 - a_2, \qquad M(g) = 2\min(a_1, a_2).$$

This recovers the two analytic channels already present in the stopped argument.

The fixed-type relative-Anosov route fails for a precise reason. On a parabolic power T^N, all three a_i equal $\log|N| + O(1)$, so

$$D(T^N) = O(1), \qquad M(T^N) = 2\log|N| + O(1).$$

On the escaping sequence supplied by outer adelic/topological density, both outer images remain compact, so

$$D(g_n) \longrightarrow \infty, \qquad M(g_n) = O(1).$$

Thus $\rho_\star$ is neither relatively P_1- nor relatively P_2-Anosov. Relative P_k-Anosov theory requires the same k-th singular gap to diverge on every escaping sequence, including peripheral ones. The quantity which survives is instead

$$\max\{D, M\},$$

or the top-to-third gap $D + M = 2a_0 - |a_1 - a_2|$. This is an adaptive change of flag type and is not the hypothesis of the existing relative-Anosov orbit-map theorems.

Route II.22.26 (Outcome of the relative-Anosov audit)**.** The theory of Canary–Zhang–Zimmer and Zhu–Zimmer supplies an exact analogy for cusp-adapted norms and explains why a fixed contracting flow would give a proper orbit map. It does not prove the required assertion: the heptagonal representation switches between the first and second singular gaps. The already constructed rank-two affine extension is the object which would have to make the adaptive maximum into one stopped Arakelov slope while counting the determinant ideal once. Calling this a relative Morse lemma merely renames the stopped fusion problem.

II.23 Primary references used for the Roth comparison

References

[1] W. Ljunggren, *Zur Theorie der Gleichung* $x^2 + 1 = Dy^4$, Avh. Norske Vid. Akad. Oslo I, no. 5 (1942).

[2] J. H. E. Cohn, *Perfect Pell powers*, Glasgow Math. J. **38** (1996), 19–20, https://doi.org/10.1017/S0017089500031207.

[3] Y. Bugeaud, P. Hubert, and T. A. Schmidt, *Transcendence with Rosen continued fractions*, J. Eur. Math. Soc. **15** (2013), 39–51, https://arxiv.org/abs/1007.2050.

[4] B. Adamczewski and Y. Bugeaud, *On the complexity of algebraic numbers II. Continued fractions*, Acta Math. **195** (2005), 1–20, https://arxiv.org/abs/math/0511677.

[5] P. Arnoux and T. A. Schmidt, *Cross sections for geodesic flows and* α*-continued fractions*, Nonlinearity **26** (2013), 711–726, https://arxiv.org/abs/1207.7299.

[6] K. Calta and J. Smillie, *Algebraically periodic translation surfaces*, J. Mod. Dyn. **2** (2008), 209–248, https://arxiv.org/abs/math/0703567.

[7] P. Arnoux and T. A. Schmidt, *Veech surfaces with non-periodic directions in the trace field*, J. Mod. Dyn. **3** (2009), 611–629, https://arxiv.org/abs/0906.4571.

[8] J. Boulanger, *Central points of the double heptagon translation surface are not connection points*, Bull. Soc. Math. France **150** (2022), 459–472, https://arxiv.org/abs/2009.01748.

[9] Y. Benoist and S. Miquel, *Arithmeticity of discrete subgroups containing horospherical lattices*, Duke Math. J. **169** (2020), 1485–1539, https://arxiv.org/abs/1805.00045.

[10] R. Canary, T. Zhang, and A. Zimmer, *Cusped Hitchin representations and Anosov representations of geometrically finite Fuchsian groups*, Adv. Math. **404** (2022), Paper No. 108439, https://arxiv.org/abs/2103.06588.

[11] F. Zhu and A. Zimmer, *Relatively Anosov representations via flows I: theory*, https://arxiv.org/abs/2207.14737.

[12] S. Cavallar and F. Lemmermeyer, *The Euclidean algorithm in cubic number fields*, Proceedings of the International Conference, Eger, 1996, 123–146 (1998), https://arxiv.org/abs/1202.6019.

[13] P. Lezowski and K. J. McGown, *The Euclidean algorithm in quintic and septic cyclic fields*, Math. Comp. **86** (2017), 2535–2549, https://arxiv.org/abs/1601.03433.

[14] P. Vojta, *A quantitative proof of Roth's theorem with moving targets*, source and theorem statement available at https://math.berkeley.edu/~vojta/ftp/mtroth2.tex.

[15] M. Maculan, *An explicit version of the theorem of Roth with moving targets*, https://arxiv.org/abs/1305.0926.

[16] M. Fujimori, *The algebraic groups leading to the Roth inequalities*, J. Théor. Nombres Bordeaux **24** (2012), 257–292, https://doi.org/10.5802/jtnb.796.

[17] C. McMullen, *Billiards, heights, and the arithmetic of non-arithmetic groups*, Invent. Math. **228** (2022), 1309–1351, https://doi.org/10.1007/s00222-022-01101-4.

[18] D. Rosen, *An arithmetic characterization of the parabolic points of* $G(2\cos(\pi/5))$, Glasgow Math. J. **6** (1963), 88–96, https://doi.org/10.1017/S204061850003478X.

[19] R. Schoof, *Computing Arakelov class groups*, in *Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography*, Math. Sci. Res. Inst. Publ. **44**, Cambridge Univ. Press, 2008, 447–495, https://arxiv.org/abs/0801.3835.

[20] S. Filip, *Uniformization of some weight* 3 *variations of Hodge structure, Anosov representations, and Lyapunov exponents*, arXiv:2110.07533v2, 2022, https://arxiv.org/abs/2110.07533.

[21] M. Ru and P. Vojta, *Schmidt's subspace theorem with moving targets*, Invent. Math. 127 (1997), 51–65, https://doi.org/10.1007/s002220050114.

[22] G. Zhao and Y. Li, *On a theorem of Schmidt's subspace theory with moving targets*, preprint (2026), https://arxiv.org/abs/2603.01627.

[23] H. Nakada, *On the Lenstra constant associated to the Rosen continued fractions*, J. Eur. Math. Soc. 12 (2010), 55–70, https://ems.press/content/serial-article-files/31700.

[24] C. Kraaikamp, H. Nakada, and T. A. Schmidt, *Metric and arithmetic properties of mediant–Rosen maps*, Acta Arith. 137 (2009), 295–324, https://doi.org/10.4064/aa137-4-1.

[25] C. Kraaikamp, T. A. Schmidt, and I. Smeets, *Natural extensions for* α*-Rosen continued fractions*, J. Math. Soc. Japan **62** (2010), 649–671, https://arxiv.org/abs/0905.4588.

[26] D. Rosen and T. A. Schmidt, *Hecke groups and continued fractions*, Bull. Austral. Math. Soc. 46 (1992), 459–474, https://doi.org/10.1017/S0004972700012120.

[27] A. D. Pohl, *Symbolic dynamics for the geodesic flow on Hecke surfaces*, preprint, https://arxiv.org/abs/0801.3951.

[28] G. Panti, *Decreasing height along continued fractions*, Ergodic Theory Dynam. Systems 40 (2020), 763–788, https://doi.org/10.1017/etds.2018.55.

[29] J.-H. Evertse, H. P. Schlickewei, and W. M. Schmidt, *Linear equations in variables which lie in a multiplicative group*, Ann. of Math. (2) 155 (2002), 807–836, https://doi.org/10.2307/3062133.

[30] J.-H. Evertse and K. Győry, *Effective results for unit equations over finitely generated domains*, Math. Proc. Cambridge Philos. Soc. **154** (2013), 351–380, https://arxiv.org/abs/1107.5756.

[31] P. Vojta, *A more general abc conjecture*, Int. Math. Res. Not. 1998, no. 21, 1103–1116, https://arxiv.org/abs/math/9806171.

[32] P. Corvaja and U. Zannier, *On the rational approximations to the powers of an algebraic number: solution of two problems of Mahler and Mendès France*, Acta Math. **193** (2004), 175–191, https://arxiv.org/abs/math/0403522.

[33] A. Kulkarni, N. M. Mavraki, and K. D. Nguyen, *Algebraic approximations to linear combinations of powers: an extension of results by Mahler and Corvaja–Zannier*, Trans. Amer. Math. Soc. **371** (2019), 3787–3804, https://arxiv.org/abs/1511.08525.

[34] V. Kumar and R. Thangadurai, *On simultaneous approximation of algebraic numbers*, Mathematika **68** (2022), 1153–1175, https://arxiv.org/abs/2001.00386.

[35] P. S. Nair, V. Kumar, and S. S. Rout, *Algebraic approximations to linear combinations of S-units*, preprint (2025), https://arxiv.org/abs/2506.02898.

[36] N. Grieve and J. T.-Y. Wang, *Greatest common divisors with moving targets and consequences for linear recurrence sequences*, Trans. Amer. Math. Soc. **374** (2021), 7673–7705, https://arxiv.org/abs/1902.09109.

[37] K. Golubev and A. Kamber, *On Sarnak's density conjecture and its applications*, Forum Math. Sigma **11** (2023), e48, https://arxiv.org/abs/2004.00373.

[38] A. Kamber and P. P. Varjú, *Lifting all elements in* $\mathrm{SL}_n(\mathbb{Z}/q\mathbb{Z})$, preprint, https://arxiv.org/abs/2310.10269.

[39] M. Kapovich, B. Leeb, and J. Porti, *Morse actions of discrete groups on symmetric spaces*, arXiv:1403.7671, https://arxiv.org/abs/1403.7671.

[40] T. Weisman, *An extended definition of Anosov representation for relatively hyperbolic groups*, arXiv:2205.07183, https://arxiv.org/abs/2205.07183.

[41] F. Zhu, *Relatively dominated representations*, Ann. Inst. Fourier (Grenoble) **71** (2021), 2169–2235, https://arxiv.org/abs/1912.13152.

[42] R. Aoun, *Transience of algebraic varieties in linear groups and application to generic Zariski density*, Ann. Inst. Fourier (Grenoble) **63** (2013), 2049–2080, https://arxiv.org/abs/1103.0944.

[43] S. Edwards, M. Lee, and H. Oh, *Anosov groups: local mixing, counting, and equidistribution*, Geom. Topol. **27** (2023), 513–573, https://arxiv.org/abs/2003.14277.

[44] M. Hochman and B. Solomyak, *On the dimension of Furstenberg measure for* $\mathrm{SL}_2(\mathbb{R})$ *random matrix products*, https://arxiv.org/abs/1610.02641.

[45] J. Li, *Fourier decay, renewal theorem and spectral gaps for random walks on split semisimple Lie groups*, Ann. Sci. Éc. Norm. Supér. (4) **55** (2022), 1613–1686, https://arxiv.org/abs/1811.06484.

[46] A. Rapaport, *Exact dimensionality and Ledrappier–Young formula for the Furstenberg measure*, https://arxiv.org/abs/2007.06326.

[47] T. Rush, *Frostman dimension of Furstenberg measure for* $\mathrm{SL}_2(\mathbb{R})$ *random matrix products*, https://arxiv.org/abs/2601.14061.

[48] W. He and N. de Saxcé, *Sum-product for real Lie groups*, https://arxiv.org/abs/1806.06375.

[49] W. He and N. de Saxcé, *Semisimple random walks on the torus*, Ergodic Theory Dynam. Systems **45** (2025), 3090–3147, https://arxiv.org/abs/2204.11453.

[50] D. Mayer, T. Mühlenbruch, and F. Strömberg, *The transfer operator for the Hecke triangle groups*, Discrete Contin. Dyn. Syst. **32** (2012), 2453–2484, https://arxiv.org/abs/0912.2236.

[51] A. Adam and A. Pohl, *A transfer-operator-based relation between Laplace eigenfunctions and zeros of Selberg zeta functions*, https://arxiv.org/abs/1606.09109.

[52] H. Oh and D. Winter, *Uniform exponential mixing and resonance free regions for convex cocompact congruence subgroups of* $\mathrm{SL}_2(\mathbb{Z})$, https://arxiv.org/abs/1410.4401.

[53] M. Magee, H. Oh, and D. Winter, *Uniform congruence counting for Schottky semigroups in* $\mathrm{SL}_2(\mathbb{Z})$, https://arxiv.org/abs/1601.03705.

[54] G. Tomanov, *Locally divergent orbits on Hilbert modular spaces and Margulis conjectures*, J. Reine Angew. Math. **689** (2014), 1–42, https://arxiv.org/abs/1012.6006.

[55] A. Baker and G. Wüstholz, *Logarithmic forms and group varieties*, J. Reine Angew. Math. **442** (1993), 19–62, https://doi.org/10.1515/crll.1993.442.19.

[56] F. Calegari, *The adelic closure of triangle groups*, preprint, Theorems 1.1 and 1.3, https://www.math.uchicago.edu/~fcale/papers/Chaos.pdf.

[57] F. Calegari, V. Dimitrov, and Y. Tang, *Arithmetic holonomy bounds and effective Diophantine approximation*, 2026, https://arxiv.org/abs/2510.04156.

[58] F. Calegari, V. Dimitrov, and Y. Tang, *The linear independence of* 1, $\zeta(2)$, *and* $L(2,\chi_{-3})$, preprint (2024), https://arxiv.org/abs/2408.15403.

[59] F. Calegari, V. Dimitrov, and Y. Tang, *The unbounded denominators conjecture*, J. Amer. Math. Soc. **38** (2025), 627–702, https://arxiv.org/abs/2109.09040.

[60] P. Habegger, *Diophantine approximations on definable sets*, Selecta Math. (N.S.) **24** (2018), 1633–1675, https://arxiv.org/abs/1608.04547.

[61] L. N. Vaserstein, *On the group* SL_2 *over Dedekind rings of arithmetic type*, Math. USSR-Sb. **18** (1972), 321–332.

[62] E. Hanson, A. Merberg, C. Towse, and E. Yudovina, *Generalized continued fractions and orbits under the action of Hecke triangle groups*, Acta Arith. 134 (2008), 337–348, https://doi.org/10.4064/aa134-4-4.

[63] Y. Hao, *Groups generated by root unipotents: higher-rank and rank-one*, preprint (2026), https://arxiv.org/abs/2607.04580.

[64] K. Winsor, *Hecke triangle groups and special hyperbolic elements*, preprint (2026), https://arxiv.org/abs/2605.30064.

[65] J. I. Burgos Gil, D. Holmes, and R. de Jong, *Positivity of the height jump divisor*, Int. Math. Res. Not. IMRN 2019, no. 7, 2044–2068, https://arxiv.org/abs/1701.00370.

[66] P. Brosnan and G. Pearlstein, *Jumps in the Archimedean height*, Duke Math. J. **168** (2019), 1737–1842, https://arxiv.org/abs/1701.05527.

[67] J.-B. Bost, *Algebraic leaves of algebraic foliations over number fields*, Publ. Math. IHÉS **93** (2001), 161–221, https://www.numdam.org/item/PMIHES_2001__93__161_0.pdf.

[68] J. Bourgain and A. Kontorovich, *On representations of integers in thin subgroups of* $\mathrm{SL}_2(\mathbb{Z})$, Geom. Funct. Anal. **20** (2010), 1144–1174, https://arxiv.org/abs/1001.4534.

[69] J. Bourgain and A. Kontorovich, *The affine sieve beyond expansion I: thin hypotenuses*, Int. Math. Res. Not. IMRN 2015, no. 19, 9175–9205, https://arxiv.org/abs/1307.3535.

[70] X. Zhang, *On representation of integers from thin subgroups of* $\mathrm{SL}_2(\mathbb{Z})$ *with parabolics*, https://arxiv.org/abs/1610.00770.

[71] A. Warren and J. Wheeler, *Incidences of Möbius transformations in* $\mathbb{F}_p$, Finite Fields Appl. **82** (2022), 102056, https://arxiv.org/abs/2107.12286.