Persistence and Regularity in Unstable Model Theory

by

Maryanthe Elizabeth Malliaris

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Logic and the Methodology of Science

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Professor Thomas Scanlon, Chair
Professor Leo Harrington
Professor George Bergman
Professor W. Hugh Woodin
Professor Theodore Slaman

Spring 2009

The dissertation of Maryanthe Elizabeth Malliaris is approved:

Chair Date

Date

Date

Date

Date

University of California, Berkeley

Spring 2009

Persistence and Regularity in Unstable Model Theory

Abstract

Persistence and Regularity in Unstable Model Theory

by

Maryanthe Elizabeth Malliaris

Doctor of Philosophy in Logic and the Methodology of Science

University of California, Berkeley

Professor Thomas Scanlon, Chair

The starting point is a question about the structure of Keisler's order, a preorder on theories which compares the difficulty of producing saturated regular ultrapowers. In Chapter 1 we show that Keisler's order reduces to the analysis of types in a finite language, i.e. that the combinatorial barriers to saturation are contained in the parameter spaces of the formulas of T. In Chapter 2 we define the characteristic sequence of hypergraphs $\langle P_n : n < \omega \rangle$ associated to a formula which describe the relevant incidence relations, and develop a general framework for analyzing the complexity of a formula in terms of the complexity of its characteristic sequence.

Specifically, we are interested in analyzing consistent partial types, which correspond to sets A such that $A^n \subset P_n$ for all n. The key issues studied in Chapter 2 are localization and persistence, which describe the difficulty of separating some

fixed complex configuration from a complete graph under analysis by progressive restrictions of the base set. We characterize stability and simplicity of φ in terms of persistence in the characteristic sequence.

Chapter 3 restricts attention to the behavior of the graph P_2 in the characteristic sequence of a given formula. We ask how subsets of the parameter space can generically interrelate by asking what densities can occur between sufficiently large ϵ-regular pairs $A, B \subset P_1$, in the sense of Szemerédi. When the formula is stable, after localization the density must always be 1. In a class including simple theories, after localization the density must approach either 0 or 1. In the absence of strict order, we characterize the property that P_1 contains large disjoint ϵ-regular sets of any reasonable density δ in terms of instability of P_2.

Chapter 4 observes and explicates a discrepancy between the model-theoretic notion of an infinite random k-partite graph and the finitary version given by Szemerédi regularity, showing that a class of infinite k-partite random graphs which do not admit reasonable finite approximations must have the strong order property SOP_3.

Chapters 2-4 take place in a general setting. Chapter 5 describes how the formalism of characteristic sequences may be applied to the analysis of types in ultrapowers.

Professor Thomas Scanlon
Dissertation Committee Chair

To my parents

Contents

Acknowledgments

I would like to express my great gratitude to my thesis advisor Thomas Scanlon for his deep wisdom and constant encouragement. He has taught me a tremendous amount about the many levels on which it is possible to read a mathematical idea.

It has been an honor to learn from Leo Harrington, whose understanding is immense and who can with a light touch reveal the surface below the surface of things.

Because of them, these past few years have been an illumination and a pleasure.

Thanks to George Bergman for many helpful remarks which significantly improved the exposition of this thesis. Thanks also to the many people who helped with administrative matters, in particular Barb Waller and Thomas Brown.

I thank my parents for their neverending encouragement and for their example of the examined life well and fully lived.

Finally, I would like to acknowledge the support of an Anglo-California fellowship to Pembroke College, a Phoebe Apperson Hearst scholarship, Mathlogaps and the student support provision of Scanlon's NSF grant.

Chapter 0

Introduction

0.1 Thesis work

Historically one of the great successes of model theory has been Shelah's stability theory: a program, described in [26], of showing that the arrangement of first-order theories into complexity classes according to a priori set-theoretic criteria (e.g. counting types over sets) in fact pushes down to reveal a very rich and entirely model-theoretic structure theory for the classes involved: what we now call stability, superstability, and ω-stability, as well as the dichotomy between independence and strict order in unstable theories. The success of the program may be measured by the fact that the original set-theoretic criteria are now largely passed over in favor of definitions which mention ranks or combinatorial properties of a particular formula.

Because of this shift, Keisler's 1967 order (defined below) may strike the modern reader as an anachronism. It too seeks to coarsely classify first-order theories in terms of a more set-theoretic criterion, the difficulty of producing saturated regular ultrapowers, but its structure has remained largely open. Partial results from the 70s suggest a mine of perhaps comparable richness, one which has remained largely inaccessible to current tools.

Keisler's criterion of choice, saturation of regular ultrapowers, is natural for two reasons. First, when the ultrapower is regular, the degree of its saturation depends only on the theory and not on the saturation of the index models. Second, ultrapowers are a natural context for studying compactness, and Keisler's order can be thought of as studying the fine structure of compactness by asking: what families of consistent types are realized or omitted together in regular ultrapowers? Thus the relative difficulty of realizing the types of T_1 versus those of some T_2 in regular ultrapowers gives a measure of the combinatorial complexity of the types each T_i is able to describe.

Definition 0.1. (Keisler's order [12]) $T_1 \leq T_2$ *if for all infinite* λ*,* $\mathcal{D}$ *regular on* λ*,* $M_1 \models T_1, M_2 \models T_2$*, we have: if* $(M_2)^\lambda/\mathcal{D}$ *is* λ^+*-saturated then* $(M_1)^\lambda/\mathcal{D}$ *is* λ^+*-saturated.*

Shelah in the 1970s gave a beautiful and surprising series of results showing deep links between Keisler's order and the underlying structure of first-order theories. His dividing lines will be familiar to model theorists who have not worked on ultrapowers:

Theorem A. (Shelah [26]; see Chapter 1, Theorem H below for the full statement) *In the Keisler order we have:* $\mathcal{T}_1 < \mathcal{T}_2 < ...?... \leq \mathcal{T}_s$*, where:*

1. *$\mathcal{T}_1$ is the set of countable theories without the finite cover property, which form the minimum Keisler equivalence class.*

2. *$\mathcal{T}_2$ is the set of countable theories which are stable but have fcp, which form the second Keisler equivalence class.*

3. *$\mathcal{T}_s$ is the maximum class, which is known to exist and to include theories with the strict order property.*

4. *and the intermediate structure of the unstable ...?..., as well as the question of determining the boundary of the maximum class, remains open.*

Notice the coarseness of the order. Stability is a classic model-theoretic frontier, but the finite cover property crosscuts all of its usual refinements. Recent work of Shelah [27] and Shelah and Usvyatsov [28] has shown that SOP_3, a weakening of strict order, is sufficient for maximality; however, the identity of the maximal class, as well as the structure of the order on unstable theories without SOP_3, has remained open.

Notice also that stability, fcp and strict order are all properties of formulas. In the first chapter of this thesis we show that this is paradigmatic: the Keisler order reduces to the study of types in a single formula (Theorem 1.33 below). In other words, the combinatorial barriers to saturation are contained in the parameter spaces of the formulas of T. This mirrors the crucial move of stability theory in reducing questions of a priori infinitary combinatorics to properties of formulas. But proof itself suggests the importance of a new kind of combinatorial structure.

Thus in Chapter 2 we associate to each formula φ a countable sequence of hypergraphs, called the "characteristic sequence," which describe incidence relations on the parameter space of φ. We then begin the investigation of the model-theoretic complexity of φ in terms of the graph-theoretic complexity of its characteristic sequence, that is, the distribution and recurrence of complex configurations around the base set of a φ-type under analysis.

Definition 0.2. *The* characteristic sequence $\langle P_n : n < \omega \rangle$ *associated to a formula* φ *of* T *is given by: for* $n < \omega$, $P_n(z_1, \dots z_n) := \exists x \bigwedge_{i \leq n} \varphi(x; z_i)$. *Write* $(T, \varphi) \mapsto \langle P_n \rangle$.

This move is a natural consequence of the proof of Theorem 1.33. Classification theory typically isolates particular configurations which signal complexity (the order property, the independence property...); an interest in saturation of ultrapowers shifts the emphasis onto understanding how the many fragments of configurations are distributed in the parameter space of the formula and how they cluster into larger constellations, into constellations of constellations, etc. Once observed and made precise, this relation of questions of "presence" as seen in the formula φ to questions of "persistence" as seen in the hypergraphs is an interesting structural issue beyond the context of ultrapowers.

Chapter 2 applies the characteristic sequence to the analysis of consistent partial φ-types, which correspond to complete P_∞-graphs, i.e. sets $A \subseteq M$ such that $A^n \subseteq P_n$ for all n. A first goal is to definably restrict the predicate P_1 around A

so that the localized graph is as uncomplicated as possible, Definition 2.27 below. A combinatorial configuration will be called *persistent* around A if it appears in every finite localization around the complete graph A under analysis, Definition 2.36. The main results of the chapter are characterizations of stability (Theorem 2.55) and simplicity (Theorem 2.60) in terms of persistence.

Chapter 3 restricts attention to some fixed localization and considers what the complexity of configurations there imply for T. This provides a second motivation for characteristic sequences: linking classification theory for φ to structural issues of distributions of edges in the characteristic sequence of hypergraphs is potentially quite powerful, because as properties like edge density, randomness, and regularity of the graphs are shown to give meaningful model-theoretic information about φ, this opens up the possibility of using a deep collection of structure theorems for graphs, for instance Szemerédi-type regularity lemmas [29], to give model-theoretic information. In the notation of Chapter 3, Section 3.3,

Definition 0.3. ([29], [15]) *Fix $0 < \epsilon < 1$, and write $\delta(X,Y)$ for the edge density $e(X,Y)/|X||Y|$. The finite bipartite graph (X,Y) is ϵ-regular if for every $X' \subseteq X$, $Y' \subseteq Y$ with $|X'| \geq \epsilon|X|$, $|Y'| \geq \epsilon|Y|$, we have: $|\delta(X,Y) - \delta(X',Y')| < \epsilon$.*

Theorem B. (Szemerédi [29]) *For every $0 < \epsilon < 1, m_0 \in \mathbb{N}$ there exist $N = N(\epsilon, m_0)$, $m = m(\epsilon, m_0)$ such that: for any graph X, $|X| \geq N$, for some $m_0 \leq k \leq m$ there exists a partition $X = X_1 \cup \cdots \cup X_k$ satisfying:*

- $||X_i| - |X_j|| \leq 1$ *for $i, j \leq k$*

- *All but at most ϵk^2 of the pairs (X_i, X_j) are ϵ-regular.*

Analogous lemmas for hypergraphs exist, e.g. [10], though the issue of how to extend regularity to hypergraphs is a subtle one [11].

The organizing principle of Chapter 3 is the question of how subsets of the parameter space can generically interrelate, i.e., what densities can occur between sufficiently large ϵ-regular pairs $A, B \subseteq P_1$, in the sense of Szemerédi. We obtain an interesting picture. When the formula is stable, after localization the density must always be 1. In a class including simple theories, after localization the density must approach either 0 or 1. We may assume NSOP as strict order is already Keisler-maximal; with this hypothesis, we characterize the property that P_1 contains large disjoint ϵ-regular sets of any reasonable density δ in terms of instability of P_2, in the sense of model theory, and obtain several corollaries.

Chapter 4, a slightly more technical interlude, takes as its starting point the observation of a gap between the kind of bipartite randomness given by model theory (i.e. the independence property) and that given by Szemerédi regularity. This gap has to do with the way in which the finite subgraphs approximate the infinite. We formalize this gap and use it to describe a general principle: what might be called "the depth of independence" of an infinite k-partite graph. We show that graphs which are partially, but not fully, independent in this sense give rise to SOP_3. This gives a new motivation for the property, which is known to imply maximality in the Keisler order.

Chapter 1 began with the Keisler order, showing that an analysis of its structure depends on an analysis of φ-types, that is, types in a finite language. Chapters 2-4 developed a framework for analyzing the combinatorial complexity of φ-types in a general context. Chapter 5 gives arguments necessary to apply the constructions of Chapters 2-4 to the analysis of types in ultrapowers.

0.2 Notational conventions

This section records some conventions which will be in place throughout this thesis. Further conventions which require preliminary definitions are given in §0.4. Local conventions are laid out at the beginning of the chapters.

The letter λ will denote an infinite cardinal, typically identified with the base set of some regular ultrafilter. The letters φ, ψ, θ will be formulas of some ambient theory T. Unless otherwise specified, other lower-case Greek letters will be cardinals or ordinals (as indicated by context), usually infinite.

Typically, lower-case Roman letters satisfy: a, b, c, d, e are elements of a model, f, g, h are functions, i, j, k, l, m, n, s, t are integers or indices, p, q, r types, u, v, w, x, y, z are variables ranging over elements, and the remaining letter is $o$mitted for clarity.

$\mathcal{D}$ is a regular ultrafilter on the index set λ, Definition 1.1.

T is a first-order theory, almost always countable, and M, N are models of T. In the sections on regular ultrapowers, by convention $N := M^\lambda/\mathcal{D}$.

$\mathcal{P}_{\aleph_0}(\lambda) = \{\sigma \subseteq \lambda \ : \ |\sigma| < \aleph_0\}$.

The proof in Chapter 1 of the reduction to φ-types (Chapter 1, Theorem 1.33) makes use of predicates P_i and Q_j defined in the course of that proof.

However, beginning with Chapter 2, $\langle P_n \rangle$ is always the characteristic sequence of some formula $\varphi \in T$, and subscripted predicates P_n refer to elements of the characteristic sequence. Superscripted P^f_n are *localized* P_n (Definition 2.27, page 71). Superscripted φ^f are the analogously localized formulas (Definition 2.27).

In Chapter 2, $\alpha, \alpha_n, \beta_n$ refer to specific counting functions defined on the characteristic sequence.

In Chapters 2-5, $\varphi_n = \varphi_n(x; y_1, \ldots y_n)$ refers to the conjunction $\bigwedge_{i \leq n} \varphi(x; y_i)$.

In Chapter 5, m^*, k^* are distinguished nonstandard integers which record the size of specified profinite sets.

Throughout the thesis, if a variable or a tuple is written x or a rather than $\overline{x}, \overline{a}$, this does not necessarily imply that $\ell(x), \ell(a) = 1$.

0.3 Background from classification theory

This section is intended to organize the definitions given in Section 0.4; the key terms appear in italics below.

The project of model-theoretic classification theory, as articulated and largely developed by Shelah, is to find "good dividing lines" among first-order theories. The test of a good dividing line, as opposed to simply an interesting property, is whether its presence and absence both have strong structural implications. Let us review some of the major discoveries.

The key dividing line of stability arose from the question of counting types (see [22], [21]). Say that T is *unstable* if for every cardinal λ there exists a model $M \models T$ and $A \subseteq M$, such that $|S(A)| > |A| = \lambda$; otherwise T is *stable*. Shelah gave a deep series of characterizations of this dividing line, relating this count to very strong structural properties of the theory and its models. For instance, T is stable just in case all types are definable; just in case no formula has the *order property*; just in case all *indiscernible sequences* are indiscernible sets (see [26].II.2). In the 40 years since [22], stable theories have served as the crucible in which most of the techniques of contemporary model theory have been developed.

On the other side of this dividing line, some paths have been marked through the large territory of *unstable theories*, though much wilderness remains. A theory is unstable just in case it contains a formula with the *order property*. [26].II.4 reveals a polarization of the order property: T unstable implies that either T contains a definable linear order (SOP), or a bipartite random graph (IP), or both. (The definitions in the next section will explain the "polarization" remark.)

$$\text{Stable} = (\text{not } IP) \cap (\text{not } SOP) \tag{0.1}$$

Our main focus in this thesis will be unstable theories and, in particular, theories which have the *independence property* (IP) but not the *strict order property* (SOP).

Shelah developed, and Kim characterized, a generalization of stability known as *simplicity* [23], [13], [14]. Simple theories are the largest class of theories on which nonforking satisfies a series of desirable properties (symmetry, transitivity...). The boundary between simple/non-simple marks a point where, for at least one formula φ, *dividing* becomes pervasive. Indeed the local characterization of simplicity, T is simple iff no formula φ of T has the *tree property*, says precisely that, for each φ and each $k < \omega$, there is a uniform finite bound on the number of times φ can sequentially k-divide. When T is not simple, arbitrarily long φ-dividing chains called trees abound, and another surprising polarization appears:

$$\text{Stable} \subsetneq \text{Simple} = ((\text{not } TP_1) \cap (\text{not } TP_2)) \subsetneq \text{not } SOP \tag{0.2}$$

These two possible kinds of *tree property*, TP_1 and TP_2, are clearly combinatorially interesting but it is not yet clear what their strength is as dividing lines. Let us close by mentioning a proposed stratification of theories with the tree property but without SOP: the hierarchy of *n-stronger order properties* SOP_n ($n \geq 3$) [27]. SOP_n says that T contains a binary relation with infinite chains and no loops of length n (Definition 4.7, 130). SOP_2 is defined to be TP_1.

$$\text{Tree property} \Leftarrow SOP_2 \Leftarrow SOP_3 \Leftarrow \cdots \Leftarrow SOP_n \Leftarrow SOP_{n+1} \Leftarrow \ldots ... \Leftarrow SOP \tag{0.3}$$

Little is known about the strength or structure theory of this very interesting sequence of properties (for SOP_3, see Chapter 4). The implications are known to be strict with one exception: it remains open whether $SOP_2 \implies SOP_3$.

0.4 Background properties

This section collects a series of model-theoretic properties referred to throughout the thesis. See the previous section for context. References are interspersed. For basic properties of first-order theories, formulas, models, and types, see [7].

Convention 0.4. (Throughout the thesis)

- *T is a first-order theory, usually countable.*
- *$\mathcal{L}$ is the language of T.*
- *$\varphi(x; y)$ will be an $\mathcal{L}$-formula consistent with T. Write "φ is a formula of T."*
- *Lower-case variables and tuples $x, y, a \ldots$ written without an overline need not be singletons, though they are always finite. We will generally omit overlines, except to emphasize the construction of some $\overline{z}$ from given tuples $z_1, \ldots$.*

Definition 0.5. (Consistency) *For any formula $\varphi(x; y)$ of T, $M \models T$ and $A \subseteq M$,*

1. *A set of instances of φ, $\Sigma = \{\varphi(x; a) : a \in A\}$ is* consistent *if for every finite $A_0 \subseteq A$, $M \models \exists x(\bigwedge_{a \in A_0} \varphi(x; a))$. Equivalently, Σ is consistent if there is an elementary extension M' of M and an element $c \in M'$ such that, for each $a \in A$, $M' \models \varphi(c; a)$.*
2. *A set of instances of φ is* inconsistent *if it is not consistent, i.e. if some finite subset is inconsistent.*

3. *A set of instances of φ is k-*consistent *if every subset of size k is consistent.*

4. *A set of instances of φ is k-*inconsistent *if every subset of size k is inconsistent.*

Definition 0.6. (Types) *Let T be a theory in the language $\mathcal{L}$, $M \models T$.*

1. *A* type *$p(x)$ over $A \subseteq M$ ($A = \emptyset$ possible) is a consistent set of instances of formulas of T, in the free variable(s) x and with parameters from A. A* complete type *over A is a maximal consistent such set. We will follow these notational conventions:*

   - *$S(A)$ is the set of all types with parameters in A.*
   - *$S_n(A)$ is the set of all types in $S(A)$ in n free variables.*
   - *Write $S_\varphi(A)$ to indicate that the formulas in the type are restricted to positive and negative instances of some fixed formula φ of T.*
   - *The type of an element c in M over $A \subseteq M$ is $\{\varphi(x;a) : \varphi(x;y) \in \mathcal{L}, a \in A^{\ell(y)}, M \models \varphi(c;a)\}$.*
   - *p is called a* partial type *if it is not necessarily complete.*

2. *A φ-*type *is a type generated by positive and negative instances of a single formula.*

Definition 0.7. (Indiscernibles, dividing, forking) *Fix $M \models T$.*

1. *Let $A, B \subseteq M$, with B possibly $\emptyset$, and fix an enumeration $\langle a_i : i < \omega \rangle$ of A. The sequence A is a B-*indiscernible sequence (*or is* indiscernible over B) *if for*

any $n < \omega$ and any n-tuple $a_{i_1}, \dots a_{i_n}$ from A, the type of $a_{i_1}, \dots a_{i_n}$ over B depends only on the order-type of $i_1, \dots i_n$.

*If the enumeration of A does not matter (i.e., for any $\sigma : n \to n$, $tp(a_{i_1}, \dots a_{i_n})$ $= tp(a_{\sigma(i_1)}, \dots a_{\sigma(i_n)})$), A is an B-*indiscernible set*.*

Indiscernible means $\emptyset$-indiscernible.

2. *The formula $\varphi(x;a)$ k-*divides *(over B) if there is an infinite (B-)indiscernible sequence $A = \langle a_i : i < \omega \rangle$, with $a = a_0$, such that $\{\varphi(x;a_i) : i < \omega\}$ is 1-consistent but k-inconsistent. The formula $\varphi(x;a)$* divides *if it k-divides for some finite k.*

Definition 0.8. (Saturation) *Let λ be an infinite cardinal.*

1. *The model M* realizes *the type $p(x) \in S(A)$, $A \subseteq M$, if there exists $c \in M^{\ell(x)}$ such that for all formulas $\psi(x;a) \in p$, $M \models \psi(c;a)$. If no such c exists, then M* omits *p.*

2. *The model M is λ-*saturated *if for every $A \subseteq M$, $|A| < \lambda$, and every type $p \in S(A)$, there exists an element $a \in M$ realizing $p(x)$.*

   Below, we focus on λ^+-saturation, that is, on whether a given model realizes all types over sets of size $\leq \lambda$.

Below, the finite cover property is due to Keisler [12]; the remainder of the properties below are due to Shelah [26].

Definition 0.9. (FCP, OP, IP, SOP) *[26].II.4 Let φ be a formula of T and $M \models T$ be any sufficiently saturated model. The formula $\varphi(x; y)$ has:*

1. not the finite cover property, *written* nfcp, *if there exists $k < \omega$ such that: for any $A \subseteq M$ and any set $X = \{\varphi(x; a) : a \in A\}$ of instances of φ, k-consistency implies consistency. (This does not depend on the model chosen.)*

2. *the* finite cover property, *written* fcp, *if it does not have nfcp: i.e. if for cofinally many $k < \omega$ there is a set X_k of instances of φ which is k-consistent but $(k+1)$-inconsistent.*

3. *the* order property *if there exist elements a_i, b_i $(i < \omega)$ such that $\varphi(b_j; a_i)$ iff $j \leq i$.*

   Formulas with the order property are called unstable.

4. *the* independence property *if there exist elements a_i $(i < \omega)$ such that:*

$$\bigwedge_{\sigma,\tau \in \omega^{<\omega},\ \sigma \cap \tau = \emptyset} \exists x \left((i \in \sigma \implies \varphi(x; a_i)) \wedge (j \in \tau \implies \neg\varphi(x; a_j)) \right)$$

5. *the* strict order property *if there exist elements a_i $(i < \omega)$ such that:*

$$\forall i, j \left(\exists x (\varphi(x; a_i) \wedge \neg\varphi(x; a_j)) \iff j < i \right)$$

Convention 0.10. *Say that a theory T has one of the properties of Definition 0.9 iff one of its formulas does.*

The next set of tree properties appeared in various forms in different papers, thus the multiplicity of names. See also [26].III.7.

Definition 0.11. (Tree properties) *Let $\subseteq$ indicate initial segment. To simplify notation, say that the nodes $\rho_1, \rho_2 \in \omega^{<\omega}$ are* *incomparable *if*

$$\neg(\rho_1 \subseteq \rho_2) \wedge \neg(\rho_2 \subseteq \rho_1) \wedge \neg(\exists \nu \in \omega^{<\omega}, i, j \in \omega)(\rho_1 = \nu^\frown i, \rho_2 = \nu^\frown j)$$

i.e., if they do not lie along the same branch and are not immediate successors of the same node.

Then the formula φ has:

- *the k-*tree property*, where $k < \omega$, if there is an $\omega^{<\omega}$-tree of instances of φ where paths are consistent and the immediate successors of any given node are k-inconsistent, i.e. $X = \{\varphi(x; a_\eta) : \eta \in \omega^{<\omega}\}$, and:*

  1. *for all $\nu \in \omega^\omega$, $\{\varphi(x; a_\eta) : \eta \subseteq \nu\}$ is a consistent partial type;*
  2. *for all $\rho \in \omega^{<\omega}$, $\{\varphi(x; a_{\rho^\frown i}) : i < \omega\}$ is k-inconsistent.*

  *Call any such X a φ-*tree*, or if necessary a φ-k-tree.*

- *the* tree property *if it has the k-tree property for some $2 \leq k < \omega$.*

- *the* non-strict tree property *TP_2 if there exists a φ-tree with $k = 2$ and for which, moreover:*

  $(3)_2$ *for any two *incomparable $\rho_1, \rho_2 \in \omega^{<\omega}$, $\exists x(\varphi(x; a_{\rho_1}) \wedge \varphi(x; a_{\rho_2}))$.*

- *the* strict tree property*, also known as TP_1 or SOP_2, if there exists a φ-tree with $k = 2$ and for which, moreover:*

$(3)_1$ *for any two* **incomparable* $\rho_1, \rho_2 \in \omega^{<\omega}$, $\neg\exists x(\varphi(x; a_{\rho_1}) \wedge \varphi(x; a_{\rho_2}))$.

Definition 0.12. *(Stable and simple theories) We will say that the theory* T *is:*

1. unstable *if it satisfies either of the equivalent conditions:*

   (a) *T contains an unstable formula, i.e. a formula having the order property.*

   (b) *For every $\lambda \geq \aleph_0$ there is $M \models T$, $A \subseteq M$, $|A| = \lambda$ with $|S(A)| > \lambda$, where $S(A)$ is the set of types with parameters in A.*

   Otherwise T is stable ([26].II.2).

2. not simple *if it satisfies either of the equivalent conditions:*

   (a) *There is a formula $\varphi \in T$ with the tree property.*

   (b) *For each formula $\varphi \in T$ and each $k < \omega$, there is a uniform finite bound n_k on the length of a k-φ-dividing chain.*

   Otherwise T is simple ([13]). *Stable implies simple.*

Theorem C. (Implications among these properties (Shelah))

- *[26].II.4. T nfcp, meaning that no formula in T has fcp, $\implies$ T stable. That is, every unstable theory will contain a formula with the finite cover property.*

- *In particular, if φ is unstable then the formula*

  $(\varphi(x; y_1) \iff \neg\varphi(x; y_2)) \wedge (\varphi(x; y_3) \iff \varphi(x; y_4))$

  will have the finite cover property, though φ need not have fcp.

- *[26].II.4. Refinement of instability: φ has the order property iff φ has the independence property or some boolean combination of instances of φ has the strict order property.*

- *If T is simple unstable, then T has the independence property but not the strict order property.*

- *[26].III.7. Refinement of non-simplicity: φ has the tree property iff φ has TP_1 (a tree which is as inconsistent as possible) or TP_2 (a tree which is as consistent as possible).*

- *[27] SOP implies TP_1 but the reverse is not true.*

To conclude this chapter, we mention a possible source of confusion: there will be three distinct uses of the word *regular*, which nonetheless will be clear from context and scope. Namely, there will be regular cardinals in the sense of set theory, regular ultrafilters used in Keisler's order, and regular graphs in the sense of Szemerédi.

Chapter 1

Regular ultrapowers

The Keisler order is a preorder on countable first order theories which compares the relative difficulty of producing saturated regular ultrapowers. That is, $T_1 \leq T_2$ if for any infinite cardinal λ, any $M_1 \models T_1, M_2 \models T_2$ and any regular ultrafilter $\mathcal{D}$ on λ, we have $M_2^\lambda/\mathcal{D}$ is λ^+-saturated implies $M_1^\lambda/\mathcal{D}$ is λ^+-saturated (Definition 1.22 below). This order exposes certain basic, though often surprising, tensions between finite combinatorial problems posed by fragments of a theory in each index model and infinitary combinatorial problems posed by the aggregate in the ultrapower. The dividing lines which work on the order has exhibited are of deep, and independent, model-theoretic interest.

It is known (Shelah 1978, [26].VI.5) that countable stable theories fall into precisely two equivalence classes, those with and those without the finite cover property, and that theories with SOP_3 (and thus with the strict order property) are maximal

[27], though the identity of the maximal equivalence class is not known. The classification for unstable theories, in particular unstable theories with the independence property but without SOP_3, has remained open.

Section 1.1 is an extended introduction to the kinds of combinatorial issues visible in this context. Section 1.2 gives a summary of the known results. Section 1.3 contains the main result of the chapter, Theorem 1.33, which says that any failure of saturation in an unstable theory must come from the omission of a type in a finite language.

1.1 Introduction to regular ultrapowers

This section defines regular ultrapowers and reviews some of their basic properties. Most of the material is not new. An ultrapower is a reduced product where equivalence is computed modulo an ultrafilter $\mathcal{F}$ and the index models are taken to be the same, see for instance [26].VI, [8].

Definition 1.1. *(Regular ultrapowers) Let $\lambda \geq \kappa$ be infinite cardinals.*

1. *A* κ-regularizing set *is any $X \subset \mathcal{P}(\lambda)$, $X = \langle X_i : i < \kappa \rangle$, satisfying:*
   - *X has the finite intersection property, i.e. for any $\sigma \in \mathcal{P}_{\aleph_0}(\kappa)$, $\bigcap_{i \in \sigma} X_i \neq \emptyset$*
   - *for any $t \in \lambda$, $|\{i < \kappa : t \in X_i\}| < \aleph_0$.*
2. *An ultrafilter $\mathcal{D}$ on λ is* κ-regular *if it contains a κ-regularizing set. $\mathcal{D}$ is* regular *if it is λ-regular.*

3. *A set* $A \subset N := M^\lambda/\mathcal{D}$ *is called* small *if* $|A| \leq \lambda$. *Any* $p \in S(A)$ *is a* small type *if* A *is small.*

For the remainder of this chapter $\mathcal{D}$ *will denote a regular ultrafilter on* $\lambda \geq \aleph_0$*.*

Regular ultrapowers are "flat" in the sense that any small set in the ultrapower is contained in a product of finite subsets of the index models (see Definition 1.4). As a consequence, the amount of saturation in the ultrapower does not depend on the level of saturation of the index model, but only on its theory T. We omit the proof of Theorem D, which relies on a back-and-forth game of length λ:

Theorem D. ([26].VI.1) *Suppose that* $M_0 \equiv M_1$*, the ambient language is countable, and* $\mathcal{D}$ *is a regular ultrafilter on* λ*. Then* ${M_0}^\lambda/\mathcal{D}$ *is* λ^+*-saturated iff* ${M_1}^\lambda/\mathcal{D}$ *is* λ^+*-saturated.*

This ensures that the Keisler order, which will be the focus of this chapter, is well defined (Definition 1.22).

Fact 1.2. *Let* $\lambda \geq \aleph_0$ *be an infinite cardinal. Regular ultrafilters on* λ *always exist.*

Proof. Uniquely in this proof, let us write I for the index set of the filter to avoid confusion. Let $f : \mathcal{P}_{\aleph_0}(\lambda) \rightarrow I$ be a bijection. For $\eta \in \lambda$, define $X_\eta = \{i \in I : \eta \in f^{-1}(i)\}$. Then $\{X_\eta : \eta \in \lambda\}$ is a regularizing set of size λ. X has the finite intersection property and does not contain $\emptyset$, so it can be extended to a nonprincipal ultrafilter by Zorn's lemma. $\square$

1.1.1 Distributions

Let us spell out the "distribution" of a type $p \in S(A)$, $A \subset N$ small, across the index models M as a way of illustrating how types are realized or omitted in regular ultrapowers.

Convention 1.3. *For the purposes of this chapter,*

- *$\mathcal{D}$ is a regular ultrafilter on λ, and $t \in \lambda$ is an element of the index set.*
- *T is a countable theory, M is a model of T and $N := M^\lambda/\mathcal{D}$.*
- *"Small" means of cardinality $\leq \lambda$.*
- *Write $M[t]$ for the model M considered as the index model at index t.*
- *For each ultrapower $M^\lambda/\mathcal{D}$, fix a lifting: $[a]_\mathcal{D} \in M^\lambda/\mathcal{D} \mapsto a \in M^\lambda$. The parameter $a \in N$ is thus identified with $(\prod_{t<\lambda} a[t])/\mathcal{D}$, and the projections of elements and sets $a[t] \in M[t]$, $X[t] \subset M[t]$ are well defined.*

Definition 1.4. *(Distributions) Fix $T, M \models T, \lambda \geq \aleph_0$, $\mathcal{D}$ regular on λ, $N := M^\lambda/\mathcal{D}$, and a small type $p \in S(A)$, $A \subset N$. A* distribution *$d : \mathcal{P}_{\aleph_0}(p) \rightarrow \mathcal{D}$ of the type p is a monotonic assignment of each finite subset of p to an element of $\mathcal{D}$, such that d refines the Łoś map and the image of d is a regularizing set. More precisely:*

1. *For each $u \in \mathcal{P}_{\aleph_0}(p)$, $d(u) \in \mathcal{D}$.*
2. *d is monotonic, that is, for all $\varphi_i, \varphi_j \in p$, $d(\{\varphi_i, \varphi_j\}) \subset d(\{\varphi_i\})$.*

3. For each $u \in \mathcal{P}_{\aleph_0}(p)$, d is a refinement of the Łoś map. That is,

$$d(u) \subset \left\{ t < \lambda : M[t] \models \exists x \left(\bigwedge_{\varphi(x;a) \in u} \varphi(x; a[t]) \right) \right\}$$

4. For each $t \in \lambda$, $|\{u : t \in d(u)\}| < \aleph_0$.

Observation 1.5. *For any small type p in a regular ultrapower, a distribution exists.*

Proof. Let us sketch a possible construction:

- Write p as $\{\varphi_i(x; a_i) : i < \lambda\}$, where each φ_i is a formula of $\mathcal{L}$ and the parameters a_i are from A.
- Let $d_0 : p \to \mathcal{D}$ be the Łoś map, i.e. $\varphi_i(x; a_i) \mapsto \{t < \lambda : M[t] \models \exists x(\varphi_i(x; a_i[t]))\}$.
- Let $X = \langle X_i : i < \lambda \rangle$ be a regularizing set in $\mathcal{D}$. Define $d_1 : p \to \mathcal{D}$ by $d_1(\{\varphi_i\}) = d_0(\{\varphi_i\}) \cap X_i$.
- To finish, we extend the definition to $d : \mathcal{P}_{\aleph_0}(p) \to \mathcal{D}$ by:

$$d(\{\varphi_{i_1}, \ldots \varphi_{i_n}\}) := \left\{ t : \bigwedge_{k \leq n} t \in d_1\left(\{\varphi_{i_k}\}\right) \right\} \cap \left\{ t \ : \ M[t] \models \exists x \bigwedge_{k \leq n} \varphi_{i_k}(x; a_{i_k}[t]) \right\}$$

The first set is equal to $\bigcap_{k \leq n} d_1(\{\varphi_{i_k}\})$, which is large because $\mathcal{D}$ has the finite intersection property; the second is large by Łoś' theorem. □

Remark 1.6. *Paired with Observation 1.10 below, this basic construction shows the combinatorial issues at stake in realizing, or omitting, a small type in a regular ultrapower. Namely, let $\{\varphi_{i_1}\}, \ldots \{\varphi_{i_n}\}$ be singleton elements of $\mathcal{P}_{\aleph_0}(p)$ whose images* *under d all contain t. Then for all $k \in \{1, \ldots n\}$,*

$$M[t] \models \exists x \left(\varphi_{i_k}(x; a_{i_k}[t])\right) \tag{1.1}$$

But unless $t \in d(\{\varphi_{j_1}, \ldots \varphi_{j_r}\})$, for $\{j_1, \ldots j_r\} \subset \{i_1, \ldots i_n\}$, it need not be the case that:

$$M[t] \models \exists x \left(\bigwedge_{\ell \leq r} \varphi_{j_\ell}(x; a_{j_\ell}[t])\right) \tag{1.2}$$

The important class of distributions which satisfy $d(u) \cap d(v) = d(u \cup v)$ are called multiplicative*, Definition 1.7. A multiplicative distribution of p exists just in case p is realized; see Observation 1.10.*

1.1.2 Some examples

Remarks on some typical cases:

Example 1: Algebraically closed fields. Let M be an algebraically closed (hence infinite) field, $N := M^\lambda/\mathcal{D}$. For some small $A \subset N$ small let $p(x) \in S(A)$ be the type describing an element which does not satisfy any nontrivial polynomial with coefficients in A. So we can write $p(x) = \{\neg f_i(x; \overline{a}_i) : i < \lambda\}$, where each f is a finite conjunction of polynomial equations with coefficients in the finite set $\overline{a}_i$. A distribution d assigns finitely many of the f_i to each index model $M[t]$. We look in $M[t]$ for an element $c[t]$ satisfying the finitely many relevant $\neg f_i(x; \overline{a_i}[t])$, which will always exist. Then $c := \prod_{t<\lambda} c[t]/\mathcal{D}$ will satisfy the type p, because it avoids each f_i on the large set $d(f_i)$, by construction. This gives an easy proof that in any regular ultrapower of M, the transcendence degree over the prime field will be at least λ^+. Indeed, $M^\lambda/\mathcal{D}$ will always be λ^+-saturated for any infinite λ and regular ultrafilter $\mathcal{D}$ on λ.

Example 2: The random graph. The language contains equality and a binary edge relation R. The axioms say that the graph is infinite, and for each set of $2n$ distinct elements $y_1, \ldots y_n$, $z_1, \ldots z_n$,

$$\exists x \left(\bigwedge_{i \leq n} xRy_i \wedge \bigwedge_{j \leq n} \neg xRz_j \right)$$

Again, let $M \models T$ and $N := M^\lambda/\mathcal{D}$. By quantifier elimination, a small type p in N can be written as $p = \{xRa_i \wedge \neg xRb_i : i < \lambda\}$. Let d be a distribution, so:

$$t \in d(\{xRa_i \wedge \neg xRb_i\}) \implies M[t] \models \exists x \left(xRa_i[t] \wedge \neg xRb_i[t]\right)$$

The distribution may fail to be multiplicative because of "collisions" between parameters in the index models. That is:

$$M[t] \models \quad \exists x \left(\bigwedge\nolimits_{i \leq n} xRa_i[t] \wedge \neg xRb_i[t] \right)$$

$$\iff \qquad M[t] \models \left\{ \bigcup\nolimits_{i \leq n} a_i[t] \right\} \cap \left\{ \bigcup\nolimits_{j \leq n} b_j[t] \right\} = \emptyset$$

Let us write $A[t]$ for the set $\{a_i[t] : t \in d(u), u \in \mathcal{P}_{\aleph_0}(p), xRa_i \wedge \neg xRb_i \in u\} \subset M[t]$, and likewise for $B[t]$. The type p will be realized just in case there exists a distribution d in which, for almost every index model $M[t]$, $A[t] \cap B[t] = \emptyset$.

Equivalently, an ultrafilter will be able to realize all small types in models of the random graph iff for any pair of disjoint small sets $A, B \subset N$ it is possible to expand each index model by a new monadic predicate $X[t]$ so that $X := \prod_t X[t]/\mathcal{D}$ separates A and B.

We will see that there are ultrafilters which fail to have this property; in fact there is a Keisler equivalence class strictly between algebraically closed fields and the random graph.

Example 3: The finite cover property. Let M be the standard model of the finite cover property (see Definition 0.9): the language contains equality and a binary equivalence relation E, and the theory says that E has a single class of size n for each $n < \omega$. Let $N = M^\lambda/\mathcal{D}$.

Let $A \subset N$ be an infinite set contained in a single E-equivalence class, and $p(x, A) := \{xEa \wedge x \neq a : a \in A\}$. Let $B \subset N$ be a set of representatives of distinct E-equivalence classes, and $q(x, B) := \{\neg xEb : b \in B\}$. Then it is easy to see that the type q is always realized, as any finite fragment assigned to $M[t]$ by a distribution is satisfiable. For d a distribution, write $A[t]$ for the set $\{a_i[t] : (\exists u \in \mathcal{P}_{\aleph_0}(p))(t \in d(u) \;\; \wedge \;\; xEa_i \wedge x \neq a_i \;\; \in u)\}$. For p (see first line of paragraph) the following are equivalent:

1. There exists $c \in N$ such that $c \models p$.

2. Fixing some $a_0 \in A$, $\{x \in N : xEa_0 \wedge x \neq a_0\} \supsetneq A$.

3. For any $a \in A$, there exists a distribution $d' : p \to \mathcal{D}$ whose associated $A[t]$ satisfy, almost everywhere,

   $|A[t]| < |\{x \in M[t] : xEa[t] \wedge x \neq a[t]\}|$.

In fact, extrapolating from condition (2) one can see that N realizes all such types over small sets A' iff the size of every nonstandard equivalence class is large, i.e. $\geq \lambda^+$. We shall see that the size of pseudofinite sets is sensitive to the ultrafilter $\mathcal{D}$ in Theorem E below. Thus regular ultrapowers of theories with the finite cover property will not always be λ^+-saturated.

1.1.3 Multiplicative refinements

Let us describe a class of ultrafilters, the *good* ultrafilters, which are subtle enough to untangle any type in a countable language. Because there exists a first order theory whose $\mathcal{D}$-ultrapowers are λ^+-saturated iff $\mathcal{D}$ is good (Observation 1.12), we see that there must be a *maximum*, rather than simply maximal, class in the Keisler order.

Definition 1.7. *(Multiplicativity)*

1. *A function* $f : \mathcal{P}_{\aleph_0}(\lambda) \to \mathcal{D}$ *is* multiplicative *if* $f(u) \cap f(v) = f(u \cup v)$*, and* monotonic *if* $f(u \cup v) \subset f(u)$.

2. *If every monotonic* $f : \mathcal{P}_{\aleph_0}(\lambda) \to \mathcal{D}$ *has a multiplicative refinement, then* $\mathcal{D}$ *is called* λ^+-good.

The existence of λ^+-good ultrafilters on λ is a theorem of Kunen [16].

Fact 1.8. *Every ultrafilter is $\aleph_1$-good. When T is countable, this means that every ultrapower of $M \models T$ is $\aleph_1$-saturated.*

See for instance [26].VI.2.

Definition 1.9. *A distribution $d : p \to \mathcal{D}$ is* accurate *if for each index $t < \lambda$ and each finite subset $\{\varphi_{i_1}, \dots \varphi_{i_n}\} \subset D(t) := \{\varphi_j : t \in d(\varphi_j)\}$, we have that $t \in d(\{\varphi_{i_1}, \dots \varphi_{i_n}\})$ iff $M[t] \models \exists x \bigwedge_{k \leq n} \varphi_{i_k}$.*

Observation 1.10. *Choose $T, M, \lambda, \mathcal{D}$, $N := M^\lambda/\mathcal{D}$, $A \subset N$ small, $p \in S(A)$. Then the following are equivalent.*

1. *Some distribution d of p has a multiplicative refinement.*
2. *Every accurate distribution d of p has a multiplicative refinement.*
3. *The type p is realized in N.*

Proof. (2) $\implies$ (1) The construction of Observation 1.5 above shows that accurate distributions always exist.

(1) $\implies$ (3) Let d' be the multiplicative refinement. Then the formulas $\varphi_1, \dots \varphi_n$ assigned to index model $M[t]$ have a common realization in that model, because multiplicativity implies that if $\bigwedge_{i \leq n} (t \in d'(\{\varphi_i\}))$ then $t \in d'(\{\varphi_1, \dots \varphi_n\})$. Let $\alpha[t]$ be some such common realization in $M[t]$, and set $\alpha := \prod_{t<\lambda} \alpha[t]/\mathcal{D}$. Now for each formula $\varphi(x; c) \in p$, we have that $\varphi(\alpha; c)$ by Łoś' theorem, so $\alpha \models p$.

(3) $\implies$ (2) Let d be some accurate distribution. Suppose that p is realized by the element α. For $v \in \mathcal{P}_{\aleph_0}(p)$, $v = \{\varphi_{i_1}, \ldots \varphi_{i_k}\}$ set

$$d'(v) := \left\{ t : M[t] \models \bigwedge_{\ell \leq k} \varphi_{i_\ell}(\alpha[t]; a_{i_\ell}[t]) \right\} \cap d(v)$$

Now d' refines d by construction. Suppose $u, v \in \mathcal{P}_{\aleph_0}(p)$. Then $t \in d'(u) \cap d'(v)$ implies $\alpha[t]$ is a common witness, in $M[t]$, to both sets of formulas. A fortiori $t \in d(u) \cap d(v)$, and because there is a common witness and we assumed d was accurate, $t \in d(u \cup v)$. Thus $t \in d'(u \cup v)$ so d' is multiplicative. $\square$

Corollary 1.11. *If $\mathcal{D}$ is λ^+-good and $Th(M)$ is countable then $M^\lambda/\mathcal{D}$ is λ^+-saturated.*

That is to say, we have a way of assigning finitely many of the formulas of a small type to each index model in such a way that the finitely many formulas assigned to $M[t]$ have a common realization in $M[t]$.

Observation 1.12. *Let M be the model whose elements are the finite subsets of ω. The language is $\{=, \subseteq\}$, interpreted in the natural way. Let $T = T(\mathcal{P}_{\aleph_0}, \subseteq) := Th(M)$. Let $\varphi(x; y) = x \subset y$. Suppose that the ultrafilter $\mathcal{D}$ on λ is not λ^+-good. Then there is a small φ-type omitted in $N = M^\lambda/\mathcal{D}$.*

Proof. Let $f : \mathcal{P}_{\aleph_0}(\lambda) \to \mathcal{D}$ be a monotonic function with no multiplicative refinement. We would like to find elements $\{a_i : i \in \lambda\} \subset N$ such that f is an accurate distribution of a consistent partial φ-type $p = \{x \subseteq a_i : i < \lambda\}$. It would be enough to define $a_i[t] \in M[t]$ when $\{i\} \in f^{-1}(t)$ so that $M[t] \models \exists x(\bigcap_{j \leq k} x \subset a_{i_j}[t])$ just in case

$t \in f(\{i_1, \dots i_k\})$. In other words, in the index model $M[t]$, we choose finitely many sets $a_i[t]$ so that the pattern of incidence is precisely that described by f, and set all other $a_j[t] = \emptyset$. The existence of such a_i is clearly consistent with the theory, by monotonicity of f. Set $a_i := \prod_{i<\lambda} a_i[t]$ to finish; the distribution is accurate by construction, so we are done. □

Corollary 1.13. *A necessary and sufficient condition for maximality in the Keisler order is: for all λ, $M \models T$, $M^\lambda/\mathcal{D}$ is λ^+-saturated iff $\mathcal{D}$ is λ^+-good.*

Proof. Sufficiency is Corollary 1.11. Necessity follows from Observation 1.12: if the ultrafilter is not λ^+-good, then there is a theory whose $\mathcal{D}$-ultrapowers are not λ^+-saturated. □

1.1.4 Filters and theories

The interaction between ultrafilters and theories, in both directions, is both coarse and subtle. This section discusses the sorts of dimensions in first-order theories to which regular ultrafilters are sensitive, by way of mapping the large territory between the minimum and maximum Keisler class, i.e. between ω^+- and λ^+-goodness. We describe three major properties: (1) the size of pseudofinite sets, (2) the size of the cut above ω in an ultrapower of $(\omega, <)$, (3) whether or not $\mathcal{D}$ has regularizing sets below every nonstandard integer. (1)-(2) are due to Shelah; (3) is new, and will be discussed further in Chapter 5.

1.1.5 Cardinalities of sets

Fact 1.14. *Let M be a model of signature $\mathcal{L}$, $\mathcal{L}_0 \subset \mathcal{L}$ and $\mathcal{D}$ an ultrafilter on λ. Then*

$$\left(M^\lambda/\mathcal{D}\right)|_{\mathcal{L}_0} = \left(M|_{\mathcal{L}_0}\right)^\lambda/\mathcal{D}$$

Corollary 1.15. *Let $\mathcal{D}$ be any ultrafilter on λ, not necessarily regular, with M countable and $N := M^\lambda/\mathcal{D}$. Let $X[t] \subset M[t]$ be infinite, and set $X := \prod_{t<\lambda} X[t]/\mathcal{D} \subset N$. Then $|X| = |N|$.*

Proof. Let $\mathcal{L}$ be the expansion of the language to include a new function symbol $\{f\}$, interpreted almost everywhere as a bijection $f[t] : X[t] \rightarrow M[t]$. Then $f = \prod_t f[t]/\mathcal{D}$ will remain a bijection in N by Łoś' theorem. □

For pseudofinite subsets of N, the story is different.

Definition 1.16. ([26] Definition III.3.5) *Let $\mathcal{D}$ be a regular ultrapower on λ.*

$$\mu(\mathcal{D}) := \min\left\{\prod_{t<\lambda} n[t]/\mathcal{D} \ : \ n[t] < \omega, \ \prod_{t<\lambda} n[t]/\mathcal{D} \geq \aleph_0\right\}$$

be the minimum value of the product of an unbounded sequence of cardinals modulo $\mathcal{D}$.

Theorem E. (Shelah, [26].VI.3.12) *Let $\mu(\mathcal{D})$ be as in Definition 1.16. Then for any infinite λ and $\nu = \nu^{\aleph_0} \leq 2^\lambda$ there exists a regular ultrafilter $\mathcal{D}$ on λ with $\mu(\mathcal{D}) = \nu$.*

See Theorem G below. This leads to obvious failures of saturation in theories which contain a parametrized family of sets of size n for all n (the finite cover property, Definition 0.9), because an ultrapower modulo $\mathcal{D}$ will contain nonstandard elements of the family whose size is precisely $\mu(\mathcal{D})$:

Corollary 1.17. *Let M be the standard model of the finite cover property (Example 1.2.3 above), i.e. an equivalence relation E with an equivalence class of size n for each $n < \omega$, and let $\mathcal{D}$ be a regular ultrafilter on λ. Then $M^\lambda/\mathcal{D}$ is λ^+-saturated iff $\mu(\mathcal{D}) \geq \lambda^+$.*

A deep and surprising theorem of Shelah shows that all failures of saturation in ultrapowers of stable theories come from pseudofinite sets which are too small. This will establish the identities of the only two known equivalence classes in the Keisler order: T without the finite cover property, and T stable with the finite cover property (Theorem H). We sketch the proof of this result:

Theorem F. (Shelah, [26]:VI.5) *If T is a countable stable theory, $M \models T$, $\mathcal{D}$ regular on λ and $\mu(\mathcal{D}) \geq \lambda^+$, then $M^\lambda/\mathcal{D}$ is λ^+-saturated.*

Proof. (Sketch) The proof rests on the following characterization of saturated models in stable theories: when T is stable, $N \models T$ is λ^+-saturated iff N is $\kappa(T)$-saturated and every maximal indiscernible set has size at least λ^+ (see [26]: Theorem III.3.10). Essentially, this is because any type $p \in S(C)$, $|C| \leq \lambda$ does not fork over a set C_0, $|C_0| < \kappa(T)$: so by $\kappa(T)$-saturation we can find a countable indiscernible set I of realizations of $p|_{C_0}$. Let $J \supset I$ be any indiscernible set extending I. Any element $a \in J$ which is free from C over C_0 will realize the unique nonforking extension of

$p|_{C_0}$ to C, which is p. Such an a will exist if $|J| > |C|$.

In our case, T is countable so $\kappa(T) = \aleph_1$; nonprincipal ultrapowers are automatically $\aleph_1$-saturated (Fact 1.8). So it would suffice to show that every countable indiscernible set in a regular ultrapower N of a countable stable theory can be extended to an indiscernible set of size λ^+. Expanding the language slightly to code $\in$, Shelah shows that this is the same as the problem of showing that every pseudofinite set in the ultrapower has cardinality $\geq \lambda^+$. □

1.1.6 Cuts above ω

A natural way of characterizing ultrafilters is to "sound out" the depth of their multiplicativity, using sufficiently complex theories. Let $M = (\omega, <)$ be a discrete linear order. Define the *lower cofinality of* ω *modulo* $\mathcal{D}$, written $\mathrm{lcf}(\omega, \mathcal{D})$, to be the reverse cofinality of the set of elements above the image of the diagonal embedding of ω in $(\omega, <)^\lambda/\mathcal{D}$ (Definition 1.23 below), i.e. the coinitiality of ω in the ultrapower.

If $\mathrm{lcf}(\omega, \mathcal{D}) \leq \lambda$, then any formula with the order property will omit a type (Theorem 6 below; for the order property, see Definition 0.9). The following theorem shows that filters with large $\mu(\mathcal{D})$ in the sense of Definition 1.16 can still have small $\mathrm{lcf}(\omega)$, which will establish a dividing line between stable and unstable theories.

Theorem G. *(Shelah, [26].VI.3.12) For any infinite* λ, $\nu = \nu^{\aleph_0} \leq 2^\lambda$ *and* $\aleph_0 < \kappa \leq \nu$ *there exists a regular ultrafilter* $\mathcal{D}$ *on* λ *such that* $\mathrm{lcf}(\omega, \mathcal{D}) = \kappa$ *and* $\mu(\mathcal{D}) = \nu$.

Proof. (Sketch) Here is a brief sketch of the proof (given in full in [26]:Theorem

VI.3.12, pps. 357-367). The construction uses Kunen's method of families of independent functions. A family $\mathcal{F} = \{f_j : j < \kappa\} \subset \omega^\lambda$ is said to be *independent* modulo some filter D on the index set λ if, for each finite set of elements $a_1, \ldots a_t \in \omega$ and $k_1, \ldots k_t \in \kappa$, $\{i \in I : (f_{k_1}[i] = a_1) \wedge \cdots \wedge (f_{k_t}[i] = a_t)\} \neq \emptyset$ modulo $\mathcal{D}$. The functions f_j keep track of the possible remaining partitions of I which the filter has not yet decided will be large or small. They can also be thought of as elements which have not yet decided to be nonstandard (since adding $\{i : f_t[i] = n\}$ to the filter is consistent). Thus at stage η of the construction, if α is an element which is nonstandard (i.e. $\{i : \alpha[i] \geq n\} \in D$ for all $n < \omega$), then $\{i : \alpha[i] > f_t[i]\} \in D$ for any f_t which remains in the independent family. Thus the construction of the filter can be visualized as working downwards towards the standard copy of ω, progressively reducing the independent family by adding sets to D saying that certain elements of $\mathcal{F}$ are nonstandard; meanwhile, the remaining elements of $\mathcal{F}$ remain in the gap between all known nonstandard elements and ω. Finish each stage η by ensuring D_η is a maximal filter modulo which the remaining family $\mathcal{F}_\eta$ is independent. Because of this maximality, the construction ensures that once $\mathcal{F}$ is used up the original D will have been extended to an ultrafilter.

If at some point in the construction the maximal independent family remaining is of size μ, then one can show that the cardinality of the set of elements below some nonstandard integer in $(\omega, <)^\lambda/\mathcal{D}$ is no more than μ. On the other hand, suppose these μ-many elements are partitioned into κ-many blocks, for $\kappa < \mu$, and added a block at a time. Then an artifact of the construction will be a descending κ-sequence of nonstandard integers approaching the cut above ω, giving $\text{lcf}(\omega, \mathcal{D}) = \kappa$. Certain restrictions on the values of κ, μ appear in the construction: $\kappa > \aleph_0, \mu = \mu^{\aleph_0}$. $\square$

1.1.7 Flexibility

Finally, we discuss a new property of regular ultrafilters which will become important in Chapter 5.

Definition 1.18. *(Flexibility) Let $\mathcal{D}$ be a regular ultrafilter.*

1. *A $\mathcal{D}$-*nonstandard integer *is any product modulo $\mathcal{D}$ of a $\mathcal{D}$-unbounded sequence of finite cardinals, i.e. $n^* = \prod_t n[t]/\mathcal{D}$ where each $n[t] < \aleph_0$ and $n^* \geq \aleph_0$.*
2. *Let $X = \{X_i : i < \lambda\}$ be a regularizing set. For each index $t < \lambda$, we define $\sigma[t] := |\{i : t \in X_i\}|$. Define the* shell size *$\sigma(X) := \prod_t \sigma[t]/\mathcal{D}$.*
3. *Say that $\mathcal{D}$ is* flexible *if for every $\mathcal{D}$-nonstandard integer n^*, $\mathcal{D}$ contains a regularizing set X such that $\{t : \sigma[t] \leq n[t]\} \in \mathcal{D}$, where $\sigma = \sigma(X)$.*

As in the case of goodness, in order to establish flexibility as a useful property for saturation of ultrapowers, we show that it is captured by types in some countable first-order theory. The following class of theories was first studied by Buechler [5]:

Definition 1.19. *The theory T is* low *if for every formula φ there exists $k < \omega$ such that for every instance $\varphi(x; a)$ of φ, $\varphi(x; a)$ divides iff it $\leq k$-divides. The* *theory is* simple *if for every formula φ and every $k < \omega$ there exists n_k such that $D(x = x, \varphi, k) < n_k$, that is, φ cannot sequentially k-divide more than n_k times.*

Remark 1.20. *Stability implies lowness; see Observation 5.8. Many simple theories are low, but not all [6].*

Lemma 1.21. *Let φ be a formula of T which is not low, $M \models T$ $\aleph_1$-saturated, $\mathcal{D}$ regular on λ. If $M^\lambda/\mathcal{D}$ is λ^+-saturated, then $\mathcal{D}$ must be flexible.*

Proof. Suppose we are given some nonstandard integer $n^* = \prod_t n[t]/\mathcal{D}$. Let us identify a small set $A \subset N$ and a consistent type $p \in S_\varphi(A)$ which is realized iff $\mathcal{D}$ has a regularizing set with shell size $\sigma \leq n^*$ modulo $\mathcal{D}$.

By choosing the index model M to be sufficiently saturated, we may assume by the hypothesis of non-lowness that M contains indiscernible sequences I_k, for each $k < \omega$, such that $\{\varphi(x; c) : c \in I_k\}$ is k-consistent but $(k+1)$-inconsistent. Write $I_k[t]$ for this sequence in the tth copy of the index model, $M[t]$.

Fix a regularizing set $X = \{X_i : i < \lambda\}$. To build our desired set A, it suffices to define each element a_i on X_i. Indeed, we will think of the element a_i as a representative of the set X_i, in a sense that will be clear from the construction: we would like a realization of the type $p \in S(A)$ to give the desired refinement of X. Let $Y[t] = \{t : t \in X_i\}$ be the indices of elements to be defined in $M[t]$, and let $m[t] = |Y[t]|$. We shall define the $a_j[t]$ (for $j \in Y[t]$) to be distinct elements of $I_{n[t]}[t]$, where recall that $n[t]$ is the shell size we are aiming for at index t. More precisely, we choose the a_j such that:

1. $j, k \in Y[t] \implies a_j[t] \neq a_k[t]$

2. for all $\rho \subseteq Y[t]$, because the elements are chosen along the indiscernible sequence $I_{n[t]}$,

$$M[t] \models \exists x \left(\bigwedge_{j \in \rho} \varphi(x; a_j[t]) \right) \iff |\rho| \leq n[t]$$

i.e., we have chosen our $m[t]$-many elements of $M[t]$ so that precisely the subsets of size $\leq n[t]$ are consistent. To finish, for each $i < \lambda$, set $a_i := \prod_t a_i[t]/\mathcal{D}$, and $p := \{\varphi(x; a_i) : i < \lambda\}$. To see that this is a consistent type, let $p_0 = \{\varphi(x; a_{i_1}), \dots \varphi(x; a_{i_k})\}$ be any finite subset. Then $X_{i_1} \cap \dots \cap X_{i_k} \cap \{t : n[t] > k\} \in \mathcal{D}$, and at every index in this large set the formula $\varphi(x; a_{i_1}[t]), \dots \varphi(x; a_{i_k}[t])$ have a common realization in $M[t]$, by condition (2).

On the other hand, suppose that $\alpha \models p$ for some $\alpha \in N$. Now the distinctness of the elements of A allows us to push down the shell size of the original regularizing set X. Namely, let $Z = \{Z_i : i < \lambda\}$ be given by $Z_i = \{t \ : \ t \in X_i, \ M[t] \models \varphi(\alpha[t]; a_i[t])\}$. By Łoś' theorem, $Z_i \in \mathcal{D}$. Because it refines X, Z remains a regularizing set. Because the a_i were chosen to be distinct,

$$\Big| \{i : t \in X_i, \ M[t] \models \varphi(\alpha[t]; a_i[t])\} \Big| \leq n[t]$$

so $\sigma(Z) \leq n^*$ modulo $\mathcal{D}$. $\square$

1.2 Keisler's order

We now define Keisler's order and state the known results.

Definition 1.22. Keisler's order *on countable theories is given by:* $T_1 \leq_\lambda T_2$ *if for any* $M_1 \models T_1, M_2 \models T_2$*, and* $\mathcal{D}$ *a regular ultrafilter on* λ*, if* $M_2^\lambda/\mathcal{D}$ *is* λ^+*-saturated then so is* $M_1^\lambda/\mathcal{D}$*. Say that* $T_1 \leq T_2$ *if for all infinite* λ*,* $T_1 \leq_\lambda T_2$*.*

That is, any regular ultrafilter on λ which produces λ^+-saturated ultrapowers of models of T_2 will produce λ^+-saturated ultrapowers of models of T_1. The condition of regularity ensures that the Keisler order is well-defined: any two elementarily equivalent λ-regular ultrapowers are back-and-forth equivalent in a game of length λ (Theorem D above), so the quantification over all models of T is justified. Weaker preorderings on theories have been investigated by Shelah and Džamonja [9] and Shelah and Usvyatsov [28].

The Keisler order is understood when T is stable and when T has the strict order property or SOP_3. Prior to this thesis, almost nothing was known about the Keisler order for T unstable with the independence property but without SOP_3 (a weakening of strict order).

1.2.1 Previous work

H. J. Keisler was responsible for the definition of, and the initial work on, the order in the 1960s [12]. He isolated several important elements of the smallest and largest class of theories (this is the origin of the finite cover property). Almost all subsequent work is due to Shelah; see [26] Chapter VI, sections 3-5, as well as [24], [27], and related work of Shelah and Džamonja [9] and Shelah and Usvyatsov [28].

See Section 1.1 above for a discussion of these results. This work may be summarized as follows:

Theorem H. (Shelah) *Let T be a countable first order theory, $\mathcal{D}$ a regular ultrafilter on λ, and the Keisler order as in Definition 1.22. Then:*

1. (Results on equivalence classes) *(for (a)-(b), see Theorem F, page 31)*

   (a) *The theories which do not have the finite cover property are minimal in the Keisler order, and form an equivalence class.*

   (b) *The theories which are stable and have the finite cover property form an equivalence class.*

   (c) *Theories which have the strict order property are maximal. SOP_3, a weakening of strict order, is also maximal ([27], and Theorem 3.32 below). A necessary model-theoretic condition for maximality is not known.*

2. (Results on dividing lines) *(for (a)-(b), see Theorem G, page 32)*

   (a) *For any cardinal $\mu = \mu^{\aleph_0}$, $\aleph_0 < \mu \leq \lambda^+$, there exists a regular ultrafilter $\mathcal{D}$ on λ such that for some unbounded sequence of finite cardinals $\langle n[t] : t < \lambda \rangle$, we have $\prod_t n[t]/\mathcal{D} = \mu$. This shows that stable theories break into two distinct classes (*nfcp *and* fcp*) when the ultrafilter is taken on $\lambda > \aleph_0$.*

   (b) *For any cardinal $\nu = \nu^{\aleph_0}$, $\aleph_0 < \nu \leq \lambda^+$, there exists a regular ultrafilter $\mathcal{D}$ on λ with the property that $\mathrm{lcf}(\omega, \mathcal{D}) = \nu$. This gives a dividing line between stable and unstable theories.*

(c) *Assuming $MA + 2^{\aleph_0} > \aleph_1$, there exists an ultrafilter on ω which saturates certain models of the random graph, but not those of any theory with the tree property. Thus, it is consistent that there is a dividing line between the random graph and theories with the tree property ([26]:Theorem VI.3.10).*

(d) *The ultrafilters which saturate maximal theories are precisely the "good" ultrafilters (Corollaries 1.11, 1.13 above). λ^+-good filters on λ exist by Kunen [16].*

Proofs can be found in [26]:VI.5, as well as in the previous section. As discussed in Chapter 0, notice first that the classification is quite coarse; second, that the known dividing lines are of independent model-theoretic interest; third, that the proof of clause 1(b), which is Theorem F above, in fact shows that if a stable theory is not saturated then a φ-type is omitted for any φ with the finite cover property. Theorem 1.33 below shows that this is paradigmatic: the Keisler order depends on an analysis of φ-types.

1.2.2 Some definitions

These are prolegomena for the next section.

Definition 1.23. (Traces of order)

1. *Let $X \subset M^\lambda/\mathcal{D}$ be a small set. X is* true *if there exists an assignment $f : X \rightarrow \mathcal{D}$ such that for any $t \in \lambda$, $t \in f(x) \cap f(y)$ implies $M[t] \models x[t] \neq y[t]$.*

2. *Let $X \subset (\kappa, <)^\lambda/\mathcal{D}$ be a small set. X is* order-true *if there exists an assignment $f : X \to \mathcal{D}$ such that for any x, y in X and any $t \in \lambda$, with $t \in f(x) \cap f(y)$, we have $M[t] \models x[t] < y[t]$ iff $N \models x < y$.*

3. *Write $N = (\gamma, <)^\lambda/\mathcal{D}$ for the ultrapower of a discrete linear order of order-type γ, and identify γ with its image under the diagonal embedding. The set of* γ-nonstandard elements *is $\{y \in N : g \in \gamma \to y > g\}$.*

4. *As in [26]:VI.4, let $\mathrm{lcf}(\gamma, \mathcal{D})$ be the cofinality of the set of γ-nonstandard elements considered with the reverse ordering, i.e. the coinitiality of γ.*

   We can extend this definition to γ-indexed sequences of the ultrapower (for our purposes, $|\gamma| = \aleph_0$). Let $\mathrm{lcf}^*(\gamma, \mathcal{D})$ be the minimal cardinality κ_Y, where $Y \subset (\omega, <)^\lambda/\mathcal{D}$ is of order-type γ and κ_Y is the reverse cofinality of the set of Y-nonstandard elements. See clause (3) of the next remark.

Remark 1.24. (on the previous definition)

1. *Note that all countable sets are true: refine any distribution so that the nth element is distinct from its finitely many predecessors. This need not be possible for uncountable sets.*

2. *In general, for any relation R one could define R-true; true is $=$-true.*

3. *As each κ_Y is a cardinal, the minimum is well defined. As $\mathcal{D}$ is regular, whether or not $\mathrm{lcf}^*(\omega, \mathcal{D}) \geq \lambda^+$ depends on $\mathcal{D}$ and not on the choice of $M = (\omega, <)$ (Theorem D above).*

The importance of $\mathrm{lcf}(\omega, \mathcal{D})$ for the Keisler order is given by:

Theorem I. (Shelah, [26]:Theorem VI.4.8) *If M is a model of an unstable theory T and $\mathcal{D}$ is a regular ultrafilter on λ with $\mathrm{lcf}(\omega, \mathcal{D}) = \kappa$, then $M^\lambda/\mathcal{D}$ is not κ^+-compact.*

When T is countable, one can replace "compact" with "saturated." The proof shows that any formula with the order property can code a "cut" which is not realized. (As the order is not on a definable set, the fact that one side of the cut is ω is crucial.) Thus any ultrafilter on λ which produces λ^+-saturated ultrapowers of some unstable theory can be assumed to have $\mathrm{lcf}(\omega, \mathcal{D}) \geq \lambda^+$. The construction of the next section uses this fact to extract "limit definitions" which make types principal in an expanded language.

1.3 Reduction to φ-types

The main result of this section is that the problem of realizing types in regular ultrapowers of countable theories reduces to that of realizing types in finite restrictions of the language (Theorem 1.33), i.e. φ-types for all formulas φ of $\mathcal{L}$.

The first lemma says that if $\mathrm{lcf}(\omega, \mathcal{D})$ is large, so is the reverse cofinality of the set above any ω-indexed sequence in the ultrapower: i.e., the minimum of Definition 1.23 is attained above the standard copy of ω.

Lemma 1.25. *For any regular ultrafilter $\mathcal{D}$ on λ, $\mathrm{lcf}^*(\omega, \mathcal{D}) = \mathrm{lcf}(\omega, \mathcal{D})$.*

Proof. Set $N := (\omega, <)^\lambda/\mathcal{D}$. Let $W = \langle w_n : n < \omega\rangle$ be an increasing ω-indexed sequence in N. Let $C^0 = \langle c^0_j : j < \kappa\rangle$ be any reverse cofinal subset of $\{b \in N : N \models b > w_n \text{ for each } w_n \in W\}$, with $\kappa = \mathrm{cof}(\kappa)$. Fix an assignment $f : W \cup C^0 \to \mathcal{D}$ whose image is a regularizing set and which is $<$-true on W. (The second clause is possible by countability of W: Remark 1.24). Write $W[t]$ for $\{w_n[t] : t \in f(w_n)\}$, i.e. the trace of W in the model $M[t]$.

Let us define a second sequence of elements, $C^1 \subset N$ such that:

- $|C^1| = \kappa$, and C^1 is cofinal in C^0.
- C^1 remains W-nonstandard, that is, for all $w \in W$ and all $c^1 \in C^1$, $N \models c^1 > w$.
- for each $i < \kappa$, $N \models c^0_i \geq c^1_i$.
- there is an assignment $g : C^1 \to \mathcal{D}$ such that for each $c^1 \in C^1$ and each $t \in \lambda$, we have $t \in g(c^1) \to c^1[t] \in W[t]$.

Construct C^1 by defining $c^1_i[t]$ to be 0 if $t \notin f(c^0_i)$, and otherwise:

$$c^1_i[t] = \max\{k : k \in W[t],\ M[t] \models c^0_i[t] \geq k\}$$

That is, in each index model, certain elements of C^0 fall into the gaps between elements of W; we move these down slightly until they enter $W[t]$. By Łoś' theorem, the desired conditions are satisfied.

To finish, we collapse W onto ω, which induces a map from C^1 onto a sequence of the same cardinality which is reverse-cofinal above the standard copy of ω. More precisely, in each index model $M[t]$, define $h[t] : W[t] \to |W[t]|$ to be an order-preserving bijection and set $h = \prod_t h[t]/\mathcal{D}$. Set $c_i = h(c_i^1)$ and $C = \langle c_i : i < \kappa \rangle$. Then h also gives an order-preserving bijection from $C^1 \to C$. Now:

1. For each n, $\{t \in \lambda : M[t] \models c_i[t] > n\} \supseteq \{t \in \lambda : M[t] \models c_i^1[t] > w_n[t]\}$, and

2. Suppose for some β, for all $c \in C$ and all $n \in \omega$, $N \models c > \beta > \omega$. Then $N \models h^{-1}(c) > h^{-1}(\beta) > h^{-1}(w_n)$. But h is an order-preserving bijection, so this schema of conditions would imply that C^1 was not reverse-cofinal above W, contradiction.

Thus $\mathrm{lcf}(\omega) \leq |C^1| = |C^0| = \mathrm{cof}(\kappa) = \kappa = \mathrm{lcf}^*(W)$, as desired. $\square$

Definition 1.26. (Induced predicates) *Let $N = M^\lambda/\mathcal{D}$. A set $C \subset N$ is* induced *if it is equivalent modulo $\mathcal{D}$ to the product of its projections to the index models, i.e. if*

$$\left[\prod_t C[t]\right]_{\mathcal{D}} = C$$

The predicate P is an induced predicate *if its interpretation in N is an induced set. "There exists an induced predicate P such that..." means: we can, assuming this is not redundant, expand $\mathcal{L}$ by adding a new predicate symbol P whose interpretation in N is an induced set with the desired property.*

Remark 1.27. (On induced sets) *The induced sets play a key role in the analysis of types in ultrapowers. Because they come from the index models in a concrete way, these sets are typically very large (the dimensional invariants of $\mathcal{D}$-pseudofinite sets,* *such as* $\mu(\mathcal{D})$*, apply), compared to the "small" base sets of types under analysis. At the same time, because induced sets are definable in an expanded language, they are very important carriers of structural information, by Fact 1.14. For instance, if one can establish that some first-order property holds on an induced subset of* N*, then Łoś' theorem applies.*

Note that any $\mathcal{L}$*-definable set is induced.*

The second lemma shows that a strictly descending countable sequence of induced predicates in a regular ultrapower can be regarded as describing intervals in some induced discrete linear order, so that the existence of an infinite induced predicate contained in the intersection (and itself containing some fixed small set X) will follow from realization of the analogous "X, ω" cut.

Lemma 1.28. *Let* $\mathcal{D}$ *be a regular ultrafilter on* λ *such that* $\mathrm{lcf}(\omega, \mathcal{D}) \geq \lambda^+$.

Fix M *w.l.o.g. countable,* $N := M^\lambda/\mathcal{D}$*, and* $X \subset N$ *small. Let* $\langle P_n : n < \omega \rangle$ *be a sequence of induced predicates such that* $P_n \supsetneq P_{n+1} \supset X$ *for all* $n < \omega$*. Then there exists an induced predicate* P_∞ *such that for all* $n < \omega$*,* $P_n \supsetneq P_\infty \supseteq A$*.*

Proof. First, "distribute the predicates" : choose an assignment $f : \langle P_n \rangle \rightarrow \mathcal{D}$ of the countably many predicates to index models such that for each $t < \lambda$, and all $m, n < \omega$,

$$(t \in f(P_m) \cap f(P_n)) \land (m < n) \rightarrow P_m[t] \supsetneq P_n[t]$$

As the number of predicates is countable, we can refine any assignment to have this property.

Next, define a total discrete linear order on each index model $M[t]$, using a new binary relation symbol $\leq_t$, in such a way that if $m < n < \omega$ and $t \in f(m) \cap f(n)$, we have

$$\forall x, y\, (x \in P_m[t] \wedge x \notin P_n[t] \ \wedge \ y \in P_n[t]) \rightarrow x \leq_t y$$

Without loss of generality, elements not in any predicate are placed below those in P_0. Let $\leq$ on N be the order induced by $\prod_t \leq_t /\mathcal{D}$.

Choose a countable sequence of elements $\alpha := \langle \alpha_n : n < \omega \rangle \subset N$ such that $P_n(\alpha_n)$ and $\neg P_{n+1}(\alpha_n)$. Then α is an increasing sequence of $\leq$-order-type ω. Furthermore, for every n, $X \subset P_{n+1}$ implies that for each $x \in X$, $x \geq \alpha_n$. That is, in N, the elements of X, considered with the reverse induced order, form a descending sequence above α of cofinality at most $|X| = \lambda$. By the previous Lemma and the hypothesis on $\mathcal{D}$, there exists c in the ultrapower which is above α and below X in the induced order. Given a distribution d of X, define P_∞ by expanding each index model $M[t]$ with a predicate which contains precisely the elements of X above $c[t]$. P_∞ contains X by construction. On the other hand, every element y in P_∞ is α-nonstandard, so it will be contained in P_n almost everywhere, for each n. $\square$

In order to apply these lemmas, we isolate an important class of induced predicates: those which give definitions for types.

Definition 1.29. (Almost principal types)

1. *Let $A, B \subset (M^\lambda/\mathcal{D})^n$ be disjoint small sets and $\psi(x; \overline{y})$, $l(y) = n$ a formula.*

Then (A, B) *is a* small ψ-partition of N *if the type*

$$p(x) = \{\psi(x; \overline{a}) : \overline{a} \in A\} \cup \{\neg\psi(x; \overline{b}) : \overline{b} \in B\}$$

is consistent. From now on, we will again suppress mention of the arity of y. The type p is called the type associated to the partition (A, B), *or its* associated ψ-type.

2. *In the other direction, every small ψ-type $p \in S_\psi(C), C \subset N$ has an associated small ψ-partition of N.*

3. *Let A, B be a small ψ-partition of N, and p its associated ψ-type. Say that p is* almost principal *if there exist induced predicates P, Q such that $P \supset A$, $Q \supset B$ and in the expanded language $\mathcal{L} \cup \{P, Q\}$, we have that*

$$N \models \exists x \left(\forall y \left(P(y) \rightarrow \ \psi(x; y)\right) \ \wedge \ \forall z \left(Q(z) \rightarrow \ \neg\psi(x; z)\right)\right) \tag{1.3}$$

A type which is almost principal has a consistent definition in an expanded language. Because it relies on induced sets, this is a property of the ultrafilter as well as the theory.

Remark 1.30. *The existential statement in Equation 1.3 is important. It may happen that for a small φ-partition (A, B) there are induced sets P, Q containing A, B respectively, but these leave no room for the resulting "definable" φ-type to be realized. For instance, take $\varphi(x; y) = xEy \wedge x \neq y$ where E is an equivalence relation with a class of size n for each $n < \omega$, choose $\mathcal{D}$ so that $\mu = \mu(\mathcal{D}) \leq \lambda$, and take $A \subset N$ to be* *the elements in some equivalence class of size μ, and $B = \emptyset$. Then, clearly, we can take P to name the equivalence class of A, but the associated type is omitted as there does not exist x such that $x \in P$ but $x \notin A$.*

Lemma 1.31. *Let $\mathcal{D}$ be a regular ultrafilter, $M \models T$ a countable model, $\varphi = \varphi(x; y)$ a formula of T, $A \subset N$ small, $p \in S_\varphi(A)$. Then p is realized iff it is almost principal.*

Proof. ($\Leftarrow$) Suppose p is almost principal, witnessed by the induced predicates P, Q. Let (A, B) be the associated small φ-partition of N. Write $P[t], Q[t]$ for the predicates in the index model $M[t]$. By Łoś' theorem,

$$\{t < \lambda \ : \ M[t] \models \exists x \left(\forall y \left(P(y) \rightarrow \ \varphi(x; y) \right) \ \wedge \ \forall z \left(Q(z) \Rightarrow \ \neg\varphi(x; z) \right) \right) \} \in \mathcal{D} \tag{1.4}$$

Let $c[t] \in M[t]$ be the witness x given by (1.4) and set $c = \prod_t c[t]/\mathcal{D}$. For any $a \in A$, a almost everywhere in P implies that $\varphi(c; a)$, and likewise for $b \in B$, b almost everywhere in Q implies that $\neg\varphi(c; b)$, so $c \models p$.

($\rightarrow$) Let $c \in N$ be an element realizing p, and d be a distribution of the sentences obtained when c is substituted into the formulas of p. Expand each index model $M[t]$ using P to name the parameters of positive instances of φ assigned to $M[t]$ by d, and Q the parameters of negative instances. By Łoś' theorem $P \supset A$ and $Q \supset B$, and by consistency of the full theory of each index model, the predicates are disjoint in each $M[t]$, thus also in the ultrapower. Furthermore, in the expanded language $\mathcal{L} \cup \{P, Q\}$, each index model $M[t]$ satisfies $\exists x((\overline{a} \in P \rightarrow \varphi(x; \overline{a})) \wedge (\overline{b} \in Q \rightarrow \neg\varphi(x; \overline{b})))$, as witnessed by c. $\square$

Corollary 1.32. *In general, if p is a small type in positive and negative instances of $\varphi_1, \dots \varphi_n$ with parameters from C, then p is realized in N iff there exist induced predicates $P_1, \dots, P_n, Q_1, \dots Q_n$ such that:*

$$N \models \exists x \bigwedge_{i \leq n} \Big(\forall y \left(P_i(y) \rightarrow \ \varphi_i(x;y) \right) \ \wedge \ \forall z \left(Q_i(z) \rightarrow \ \neg\varphi_i(x;z) \right) \Big)$$

and such that, for each $i \leq n$, $P_i \supseteq A_i$ and $Q_i \supseteq B_i$ where (A_i, B_i) is the φ_i-partition of C implied by p.

We can now reduce the problem of realizing small types in regular ultrapowers of a countable theory to the problem of realizing types in finite subsets of the language $\mathcal{L}$.

Theorem 1.33. *Let $\mathcal{D}$ be a regular ultrafilter on λ, T a countable theory in the language $\mathcal{L}$, $M \models T$, $N = M^\lambda/\mathcal{D}$. Suppose that for all formulas $\varphi \in \mathcal{L}$, N realizes all small consistent φ-types. Then N is λ^+-saturated.*

Proof. It suffices to consider unstable T, as by Shelah's classification of stable theories (see the proof of Theorem F above) saturation in stable theories depends on realizing φ-types for φ any formula with the finite cover property. So let T be unstable. By Theorem I, we may assume $\mathrm{lcf}(\omega, \mathcal{D}) \geq \lambda^+$, as otherwise some ψ-type would be omitted for a formula ψ with the order property.

Let $p \in S(C)$ be any small (partial) type in the ultrapower N and fix an enumeration $\langle \varphi_i : i < \omega \rangle$ of the formulas of $\mathcal{L}$. For any $\sigma \subset \omega$, $|\sigma| < \aleph_0$, write p_σ for p restricted to positive and negative instances of formulas $\{\varphi_i : i \in \sigma\}$.

In the spirit of the previous lemma, to realize p it would be enough to find induced predicates P_i^∞, Q_i^∞ $(i < \omega)$ such that the type

$$q(x) := \{\forall y \left(P_i^\infty(y) \to \ \varphi_i(x;y)\right) \ \wedge \ \forall z \left(Q_i^\infty(z) \to \ \neg\varphi_i(x;z)\right) : i < \omega\}$$

is consistent in N in the expanded language $\mathcal{L} \cup \{P_i^\infty, Q_i^\infty : i < \omega\}$ and extends p, because any consistent such q will be realized in N by $\aleph_1$-saturation. The appropriate predicates P_i^∞, Q_i^∞ will be found as limits of countably many approximations below by applying Lemma 1.28.

That is, we define a countable sequence of induced predicates $\langle P_i^n, Q_i^n : n < \omega\rangle$ for each $\varphi_i \in \mathcal{L}$, such that:

1. for each $n < \omega$, P_i^n, Q_i^n witness that $p_{\{i\}}$ is almost principal.

2. for all $j < \omega$, for all but finitely many $n < \omega$,

$$\{\varphi_i(x;a) : a \in P_i^n\} \ \cup \ \{\neg\varphi_i(x;b) : b \in Q_i^n\} \ \cup \\ \{\varphi_j(x;c) : c \in P_j^n\} \ \cup \ \{\neg\varphi_j(x;d) : d \in Q_j^n\}$$

   is consistent.

Fix an enumeration $f : \mathcal{P}_{\aleph_0}(\omega) \to \omega$. By hypothesis, for each $i < \omega$, $p_{\{i\}}$ is realized. At stage $n = 0$, applying Lemma 1.31 to $p_{\{i\}}$ gives induced predicates P_i^0, Q_i^0 witnessing that $p_{\{i\}}$ is almost principal. At stage $n + 1$, following Corollary

1.32, let $X^{n+1}_{i_1}, \ldots X^{n+1}_{i_k}, Y^{n+1}_{i_1}, \ldots Y^{n+1}_{i_k}$ be induced predicates witnessing the almost-principality of p_σ, where $\sigma = \{i_1, \ldots i_k\} = f^{-1}(n+1)$. Define

$$P^{n+1}_i = \begin{cases} P^n_i \cap X^{n+1}_i \text{ if } i \in \sigma \\ \\ P^n_i \text{ o.w.} \end{cases}$$

and analogously for Q^{n+1}_i, where Y^{n+1}_i replaces X^{n+1}_i.

Now for each $\varphi_i \in \mathcal{L}$ we have defined two descending sequences of induced predicates $P^0_i \cdots \supset P^n_i \cdots \supset A_i$ and $Q^0_i \cdots \supset Q^n_i \cdots \supset B_i$, where (A_i, B_i) is the small φ_i-partition of N associated to $p_{\{i\}}$. By construction, these sequences satisfy (1)-(2). Apply Lemma 1.28 to obtain the induced predicates P^∞_i, Q^∞_i such that for all $n < \omega$ $P^n_i \supset P^\infty_i \supset A_i$ and $Q^n_i \supset Q^\infty_i \supset B_i$.

The type q mentioned above is now well defined, and by Condition (2) must be consistent. This completes the proof. $\square$

Chapter 2

Persistence

This chapter develops a framework for analyzing the combinatorial complexity of formulas φ. To each formula we associate a countable sequence of hypergraphs, the *characteristic sequence*, which describe incidence relations on the parameter space of φ, Definition 2.2. The goal is to give an analysis of φ-types by describing the way that certain distinguished sets A (the complete P_∞-graphs, avatars of consistent partial φ-types) sit inside the ambient hypergraphs P_n. Motivated in part by applications to ultrapowers, we try to understand the distribution of "complex" graph-theoretic structure and whether it persists under progressive localizations as we close in on the distinguished set A. Localization is defined in 2.27, and persistence in 2.36. In Section 2.5 we show these methods can be used to characterize stability and simplicity.

2.1 The characteristic sequence

Definition 2.1. (Notation and conventions)

1. *Throughout this chapter, if a variable or a tuple is written x or a rather than $\overline{x}, \overline{a}$, this does not necessarily imply that $\ell(x), \ell(a) = 1$.*

2. *Unless otherwise stated, T is a complete theory in the language $\mathcal{L}$.*

3. *A graph in which no two elements are connected is called an* empty graph. *A pair of elements which are not connected is an* empty pair. *When R is an n-ary edge relation, to say that some X is an R*-empty graph *means that R does not hold on any n-tuple of elements of X. X is an R*-complete graph *if R holds on every n-tuple from X.*

4. *Write $\varphi_n(x; y_1, \ldots y_n)$ for the formula $\bigwedge_{i \leq n} \varphi(x; y_i)$.*

5. *P_∞ will be shorthand for the collection of predicates P_n when the context (of a given condition, not necessarily definable, which holds of P_n for all n) is clear, e.g. A is a P_∞-complete graph meaning A is a P_n-complete graph for all n.*

6. *The complete P_∞-graph A will be called a* positive base set *when the emphasis is on its identification with some consistent partial φ-type under analysis, Observation 2.4(5).*

7. *The sequence $\langle P_n \rangle$* has support *k if: $P_n(y_1, \ldots y_n)$ iff P_k holds on every k-element subset of $\{y_1, \ldots y_n\}$. See Remark 5.4.*

8. *In discussing graphs we will typically write concatenation for union, i.e.* Ac *for* $A \cup \{c\}$.

9. *The element* $a \in P_1$ *is a* one-point extension *of the* P_n*-complete graph* A *just in case* Aa *is also a* P_n*-complete graph. In most cases,* n *will be* ∞.

10. *A formula* $\psi(x; y)$ *of* $\mathcal{L}$ *will be called* dividable *if there exists an infinite set* $C \subset P_1$ *and* $k < \omega$ *such that* $\{\psi(x; c) : c \in C\}$ *is* 1*-consistent but* k*-inconsistent. (Thus, by compactness, some instance of* ψ *divides.) When it is important to specify the arity* k*, write* k-dividable.

11. *For* consistent, inconsistent, k-consistent, k-inconsistent *see Chapter 0, §4.*

Definition 2.2. (Characteristic sequences) *Let* T *be a first-order theory and* φ *a formula of the language of* T.

- *For* $n < \omega$, $P_n(z_1, \ldots z_n) := \exists x \bigwedge_{i \leq n} \varphi(x; z_i)$.
- *The* characteristic sequence *of* φ *in* T *is* $\langle P_n : n < \omega \rangle$.
- *Write* $(T, \varphi) \mapsto \langle P_n \rangle$ *for this association.*
- *Convention: We assume that* $T \vdash \forall y \exists z \forall x (\varphi(x; z) \leftrightarrow \neg\varphi(x; y))$*. If this does not already hold for some given* φ*, replace* φ *with* $\theta(x; y, z) = \varphi(x; y) \wedge \neg\varphi(x; z)$.

Convention 2.3. *Below, we will ask a series of questions about whether certain, possibly infinite, configurations appear as subgraphs of the* P_n*, or of the* P_n^f *in some finite* *localization (Definition 2.27). For our purposes, the existence of these configurations is a property of T. That is, we may, as a way of speaking, ask if some configuration X appears, or is persistent, inside of some P_n; however, we will always mean* whether or not it is consistent with T *that there are witnesses to X inside of P_n interpreted in some sufficiently saturated model. Certainly, one could ask the question of whether some* given *model of T, expanded to model of the P_n, must include witnesses to X; we will not do so here. Thus, the formulas P_n will often w.l.o.g. be identified with their interpretations in some monster model.*

Observation 2.4. (Basic properties) *Let $\langle P_n : n < \omega \rangle$ be the characteristic sequence of (T, φ). Then, regardless of the choice of T and φ, we will have:*

1. *(Reflexivity) $\forall x (P_1(x) \rightarrow P_n(x, \ldots x))$. In general, for each $\ell \leq m < \omega$,*

$$\forall z_1, \ldots z_\ell, y_1, \ldots y_m \Big((\{z_1, \ldots z_\ell\} = \{y_1, \ldots y_m\}) \\ \implies (P_\ell(z_1, \ldots z_\ell) \iff P_m(y_1, \ldots y_m)) \Big)$$

2. *(Symmetry) For any $n < \omega$ and any bijection $g : n \rightarrow n$,*

$$\forall y_1, \ldots y_n \Big(P_n(y_1, \ldots y_n) \iff P_n(y_{g(1)}, \ldots y_{g(n)}) \Big)$$

3. *(Monotonicity) For each* $\ell \leq m < \omega$*,*

$$\forall z_1, \ldots z_\ell, y_1, \ldots y_m \Big((\{z_1, \ldots z_\ell\} \subseteq \{y_1, \ldots y_m\}) \\ \implies (P_m(y_1, \ldots y_m) \implies P_\ell(z_1, \ldots z_\ell)) \Big)$$

   So in particular, if $\models P_m(y_1, \ldots y_m)$ *and* $\ell < m$ *then* P_ℓ *holds on all* ℓ*-element subsets of* $\{y_1, \ldots y_m\}$*. The converse is usually not true; see Remark 5.4 .*

4. *(Dividing) Suppose that for some* $n < \omega$*, it is consistent with* T *that there exists an infinite subset* $Y \subset P_n$ *such that* $Y^k \cap P_{nk} = \emptyset$*. Then in any sufficiently saturated model of* T*, some instance of the formula* $\varphi_n(x; y_1, \ldots y_n) = \bigwedge_{i<n} \varphi(x; y_i)$ k*-divides.*

5. *(Consistent types) Let* $A \subset P_1$ *be a set of parameters in some* $M \models T$*. Then* $\{\varphi(x; a) : a \in A\}$ *is a consistent partial* φ*-type iff* $A^n \subset P_n$ *for all* $n < \omega$*.*

Proof. (4) By compactness, there exists an infinite indiscernible sequence of n-tuples $C = \langle c^i_1, \ldots c^i_n : i < \omega \rangle$ such that $C^k \cap P_{nk} = \emptyset$. The set $\{\varphi_n(x; c^i_1, \ldots c^i_n) : i < \omega\}$ is therefore k-inconsistent. However, it is 1-consistent: for each $c^i_1, \ldots c^i_n \in C$, $M \models P_n(c^i_1, \ldots c^i_n)$, so $M \models \exists x \varphi_n(x; c^i_1, \ldots c^i_n)$. $\square$

Convention 2.5. (T_0-configurations) *Throughout this chapter, let* T_0 *denote the incomplete theory in the language* $\mathcal{L}_0 := \{P_n : n < \omega\} \cup \{=\}$ *which describes (1)-(3) of Observation 2.4. Blueprints for hypergraphs in the language* $\mathcal{L}_0$ *which are consistent* *with T_0 will be called T_0-configurations. That is: a* finite T_0-configuration *is a pair $X = (V_X, E_X)$ where $V_X = n < \omega$, $E_X \subseteq \mathcal{P}(n)$ and the following is consistent with T_0:*

$$(\exists x_1, \dots x_n)\ (\forall \sigma \subseteq n, |\sigma| = i, \sigma = \{\ell_1, \dots \ell_i\})\ \Big(P_i(x_{\ell_1}, \dots x_{\ell_i}) \iff \sigma \in E_X \Big) \quad (2.1)$$

In general, the domain of a T_0-configuration may be infinite; we simply require that its restriction to every finite subdomain satisfy (2.1). These are the graphs which can consistenly occur as finite subgraphs of some *characteristic sequence. That every such graph appears in some sequence follows from Example 2.14 below.*

Convention 2.6. (T_1-configurations) *Fix T, φ, and the associated sequence $\langle P_n : n < \omega \rangle$. Let $M \models T$; there is a unique expansion of M to $\mathcal{L}_0 = \{P_n : n < \omega\} \cup \{=\}$. Throughout this chapter, whenever $T, \varphi, \langle P_n \rangle$ are thus fixed, let T_1 denote the complete theory of M in the language $\mathcal{L}_0$. As the characteristic sequence is definable in T, when T is complete this will not depend on the model chosen.*

*Hypergraphs in the language $\mathcal{L}_0$ which are consistent with T_1 will be called T_1-*configurations.

2.2 Some examples

This section works out several key examples. Localization and persistence will be defined in Definitions 71 and 77; the general definitions of (η, ν)-arrays and trees will be given in Definition 2.18.

Example 2.7. *(The random graph)*

T is the theory of the random graph, and R its binary edge relation (see also Chapter 1, §2.2, Example 2). Let $\varphi(x; y, z) = xRy \wedge \neg xRz$, with $(T, \varphi) \mapsto \langle P_n \rangle$. Then:

- $P_1((y, z)) \iff y \neq z$.
- $P_n((y_1, z_1), \ldots (y_n, z_n)) \iff \{y_1, \ldots y_n\} \cap \{z_1, \ldots z_n\} = \emptyset$.

Notice:

1. The sequence has support 2.
2. There is a uniform finite bound on the size of an empty graph $C \subset P_1, C^2 \cap P_2 = \emptyset$: an analysis of the theory shows that φ is not dividable, and inspection reveals this bound to be 3.
3. P_n does not have the order property for any n and any partition of the $y_1, \ldots y_n$ into object and parameter variables. (Proof: The order property in P_n implies dividability of φ_{2n} by Observation 2.32. But none of the φ_ℓ are dividable, as inconsistency only comes from equality.)
4. Of course, the formula φ has the independence property in T. We can indeed find a configuration in P_2 which witnesses this: any C which models the T_0-configuration having $V_X = \omega$ and $\{i, j\} \notin E_X \iff \exists n(i = 2n \wedge j = 2n + 1)$.

Note that φ will have the independence property on any infinite P_2-complete subgraph of the so-called $(\omega, 2)$-array C (see Observation 2.47 below).

5. As φ is unstable, φ-types are not necessarily definable in the sense of stability theory. However, we can obtain a kind of definability "modulo" the independence property, or more precisely, definability over the name for a maximal consistent subset of an $(\omega, 2)$-array as follows:

Definable types modulo independence. Let $p \in S(M)$ be a consistent partial φ-type presented as a positive base set $A \subset P_1$. Let us suppose $p \vdash \{xRc : c \in C\} \cup \{\neg xRd : d \in D\} \vdash p$, so that $A \subset M^2$ is a collection of pairs of the form (c, d) which generate the type.

There is no definable (in T with or without parameters, so in particular not from P_2) extension of the type A, so we cannot expect to find a localization of P_1 around A which is a P_2-complete graph. However:

Claim 2.8. *In the theory of the random graph, with $\varphi(x; y, z) = xRy \wedge \neg xRz$ as above, for any positive base set $A \subset P_1$ there exist a definable $(\omega, 2)$-array $W \subset P_1$, a solution S of W and an S-definable P_∞-graph containing A.*

Proof. Work in P_1. Fix any element (a, b) with $a, b \notin C, D$ and set $W_0 := \{(y, z) \in P_1 : \neg P_2((y, z), (a, b))\}$. Thus $W_0 = \{(b, z) : z \neq b\} \cup \{(y, a) : y \neq a\}$. So the only P_2-inconsistency among elements of W_0 comes from pairs of the form $(b, c), (c, a)$;

thus, writing Greek letters for the elements of P_1,

$$(\forall \eta \in W_0)(\exists \nu \in W_0)(\forall \zeta \in W_0)\left(\neg P_2(\eta, \zeta) \rightarrow \zeta = \nu\right)$$

In other words, $W := W_0 \setminus \{(b, a)\}$ is an $(\omega, 2)$-array (Definition 2.18). Moreover:

1. $(y, z), (w, v) \in W$ and $\neg P_2((y, z), (w, v))$ implies $y = v$ or $z = w$, and

2. for any $c \neq a, b$, there are $d, e \in M$ such that $(d, c), (c, e) \in W$. Thus:

3. we may choose a maximal complete P_2-subgraph C of W such that CA is a complete P_∞-graph. For instance, let C be any maximal complete extension of $\{(b, d) : d \in D\} \cup \{(c, a) : c \in C\}$. Call any such C a *solution* of the array W.

Let S be a new predicate which names this solution C of W. Then $\{y \in P_1 : z \in S \rightarrow P_2(y, z)\} \supset A$ is a P_2-complete graph, definable in $\mathcal{L} \cup \{S\}$. Support 2 implies that it is a P_∞-graph. Notice that by (2), we have in fact chosen a maximal consistent extension of A (i.e. a complete global type). $\square$

Remark 2.9. *The idiosyncracies of this proof, e.g. the choice of a* definable *$(\omega, 2)$-array, reflect an interest in structure which will be preserved in ultrapowers.*

Example 2.10. *(Coding complexity into the sequence)*

It is often possible to choose a formula φ so that some particular configuration appears in its characteristic sequence. For instance, by applying the template below when φ has the independence property, we may choose a simple unstable θ whose P_2 is universal for finite bipartite graphs (X, Y), provided we do not specify whether or not edges hold between $x, x' \in X$ or between $y, y' \in Y$. Nonetheless, Conclusion 2.33 below will show this is "inessential" structure in the case of simple theories: whatever complexity was added through coding can be removed through localization.

The construction. Fix a formula φ of T. Let $\theta(x; y, z, w) := (z = w \wedge x = y) \vee (z \neq w \wedge \varphi(x; y))$. Write $(y, *)$ for (y, z, w) when $z = w$, and $(y, -)$ for (y, z, w) when $z \neq w$. Let $\langle P_n \rangle$ be the characteristic sequence of θ, $\langle P^{\varphi}_n \rangle$ be the characteristic sequence of φ, and $\langle P^{=}_n \rangle$ be the characteristic sequence of $x = y$. Then P_n can be described as follows:

- $P_n((y_1, -), \dots (y_n, -)) \leftrightarrow P^{\varphi}_n(y_1, \dots y_n)$.
- $P_n((y_1, *), \dots (y_n, *)) \leftrightarrow P^{=}_n(y_1, \dots y_n)$.
- Otherwise, the n-tuple $y := ((y_1, z_1), \dots (y_n, z_n))$ can contain (up to repetition) at most one $*$-pair, so $z_i = z_j = * \rightarrow y_i = y_j$. In this case the unique y_* in the $*$-pair is the realization of some φ-type in the original model M of T, and $P_{n+1}((y_*, *), (y_1, -), \dots (y_n, -))$ holds iff $M \models \bigwedge_{j \leq n} \varphi(y^*; y_j)$.

Remark 2.11. *Of course, not all structure can be coded in; see Remark 3.24. Conclusion 2.33, §5.1, says that we can localize to avoid the order property, and thus the random bipartite graph, in P_2 when φ is simple. In fact, we can do this by localizing with parameters from the instance of the order property. Notice that, by contrast,* *if one or many of the P_n contain infinite random n-ary hypergraphs (not n-partite), we cannot localize to avoid this using parameters from the random hypergraph. See §3.3.5.*

Example 2.12. *(A theory with TP_2)*

For TP_2, see Definition 0.11. Let T be the model completion of the following theory [28]. There are two infinite sorts X, Y and a single parametrized equivalence relation $E_x(y, z)$, where $x \in X$, and $y, z \in Y$. Let $\varphi_{eq} := \varphi(y; xzw) = E_x(y, z) \wedge \neg E_x(z, w)$. Then:

- $P_1((xzw)) \iff z \neq w$.
- $P_2((x_1z_1w_1), (x_2z_2w_2)) \iff$ each triple is in P_1 and furthermore:

$$(x_1 = x_2) \rightarrow (\ E_x(z_1, z_2) \wedge \bigwedge_{i \neq j \leq 2} \neg E_x(w_i, z_j))$$

The sequence has support 2. There are many empty graphs; these persist under localization (Theorem 2.60). Fixing α, choose a_i $(i < \omega)$ to be a set of representatives of equivalence classes in E_α, and choose b such that $\neg E_\alpha(a_i, b)$ $(i < \omega)$. Then $\{(\alpha, a_i, b) : i < \omega\} \subset P_1$ is a P_2-empty graph. We in fact have arrays $\{(\alpha^t, a^t_i, b^t) : i < \omega, t < \omega\}$ whose "columns" (fixing t) are P_2-empty graphs and where every path which chooses exactly one element from each column is a P_2-complete graph, thus a P_∞-complete graph. The parameters in this so-called (ω, ω)-array describe TP_2 for φ_{eq} (Claim 2.22).

Claim 2.13. *P_2 does not have the order property.*

Proof. This is essentially because inconsistency requires the parameters x to coincide. Suppose that $\langle a_i, b_i : i < \omega \rangle$ were a witness to the order property for P_2. Fix any $a_i = (\alpha_s, a_s, d_s)$. Now $\neg P_2(b_j, a_i)$ for $j < i$, where $b_j = (\beta_t, b_t, c_t)$. P_2-inconsistency requires $\alpha_s = \beta_t$. As this is uniformly true, $\alpha_s = \alpha_t = \beta_s = \beta_t$ for all $s, t < \omega$ in the sequence. But now that we are in a single equivalence relation E_α, transitivity effectively blocks order: $\neg P_2(b_j, a_i) \leftrightarrow \neg E_\alpha(a_s, b_t)$. Depending on whether at least one of the a- or b-sequences is an empty graph, we can find a contradiction to the order property with either three or four elements. $\square$

Example 2.14. *(A maximally complicated theory)*

In this example the sequence is universal for finite T_0-configurations (Convention 2.5).

Let the elements of M be all finite subsets of ω; the language has two binary relations, $\subseteq$ and $=$, with the natural interpretation. Set $T = Th(M)$.

Choose $\varphi_\subseteq := \varphi(x; y, z) = x \subseteq y \wedge x \not\subseteq z$. Then:

- $P_1((y, z)) \iff \emptyset \subsetneq y \not\subseteq z$.
- $P_n((y_1, z_1), \dots (y_n, z_n)) \iff \emptyset \subsetneq \bigcap_{i \leq n} y_i \not\subseteq \bigcup_{i \leq n} z_i$.

The sequence does not have finite support. Moreover:

Claim 2.15. *Let $\langle P_n \rangle$ be the characteristic sequence of $\varphi_\subseteq$, $k < \omega$, and let X be a finite T_0-configuration. Then there exists a finite $A \subseteq P_1$ witnessing X.*

Proof. This is just the proof of sensistivity to goodness, Observation 1.12. Write the elements of P_1 as $w_i = (y_i, z_i)$; it suffices to choose the positive pieces y_i first, and afterwards take the z_i to be completely disjoint. More precisely, suppose X is given by $V_X = m$ and $E_X \subset \mathcal{P}(m)$. We need simply to choose $y_1, \dots y_m$ such that for all $\sigma \subseteq m$,

$$\left(\bigcap_{j \in \sigma} y_j \neq \emptyset \right) \iff \sigma \in E_X$$

which again, is possible by the downward closure of E_X. $\square$

Corollary 2.16. *This characteristic sequence is universal for finite T_0-configurations.*

Remark 2.17. *That the sequence is universal for finite T_0-configurations is sufficient, though not necessary, for maximal complexity in the Keisler order. By [26].VI.3, $\varphi(x; y, z) = y < x < z$ in $Th(\mathbb{Q}, <)$ is maximal. Its characteristic sequence has support 2, but its P_2 is clearly not universal.*

2.3 Static configurations

This section asks: what can we tell about φ from the T_1-configurations which appear in its characteristic sequence, without yet appealing to localization (Definition 2.27) or to invariance under localization? Some answers are given: we describe configurations which signal that φ has the order property, the independence property, the tree property and SOP_2 (Chapter 0, Definitions 0.9-0.11). More extensive analysis will come in Chapter 3. Recall that:

Theorem J. (Shelah; see Chapter 0, §0.4)

- *T is simple iff no formula φ of T has the tree property, iff no φ has the 2-tree property.*
- *If φ has the 2-tree property then either φ has TP_1 or φ has TP_2.*

We fix a monster model M from which the parameters are drawn; see Convention 2.3.

Definition 2.18. (Diagrams, arrays, trees) *Let $\lambda \geq \mu$ be finite cardinals or ω. Write $\subseteq$ to indicate initial segment. The sequence $\langle P_n \rangle$ has:*

1. an $(\omega, 2)$-diagram *if there exist elements $\{a_\eta : \eta \in 2^{<\omega}\} \subseteq P_1$ such that*
   - *for all $\eta \in 2^{<\omega}$, $\neg P_2(a_{\eta^\frown 0}, a_{\eta^\frown 1})$, and*
   - *for all $n < \omega$ and $\eta_1, \dots \eta_n \in 2^{<\omega}$, we have that $\eta_1 \subseteq \dots \subseteq \eta_n \implies P_n(a_{\eta_1}, \dots a_{\eta_n})$*

   That is, sets of pairwise comparable elements are P_∞-consistent, while immediate successors of the same node are P_2-inconsistent.

2. a $(\lambda, \mu, 1)$-array *if there exists $X = \{a^m_l : l < \lambda, m < \mu\} \subset P_1$ such that:*
   - *$P_2(a^{m_1}_{l_1}, a^{m_2}_{l_2}) \iff (l_1 = l_2 \rightarrow m_1 = m_2)$*
   - *For all $i < \omega$,*

$$P_n(a^{m_1}_{l_1}, \dots a^{m_n}_{l_n}) \iff \bigwedge_{1 \leq i,j \leq n} P_2(a^{m_i}_{l_i}, a^{m_j}_{l_j})$$

That is, any $C \subset X$, possibly infinite, is a P_∞-graph iff it contains no more than one element from each column. (We will relax this last condition in the more general Definition 2.46 below.)

3. a (λ, μ)-tree *if there exist elements* $\{a_\eta : \eta \in \mu^{<\lambda}\} \subset P_1$ *such that*

   - *for all* $\eta_2, \eta_2 \in \mu^{<\lambda}$,

     $$P_2(a_\eta, a_\nu) \iff (\eta_1 \subseteq \eta_2 \vee \eta_2 \subseteq \eta_1)$$

     i.e. only if *the nodes are comparable; and*

   - *for all* $n < \omega, \eta_1, \dots \eta_n \in \mu^{<\lambda}$,

     $$\eta_1 \subseteq \dots \subseteq \eta_n \implies P_n(a_{\eta_1}, \dots a_{\eta_n})$$

Remark 2.19. *Diagrams are prototypes which can give rise to either arrays or trees, in the case where the unstable formula φ has the independence property or SOP_2, respectively.*

The arrays will be revisited in Definitions 2.41 and 2.46.

Claim 2.20. *Let φ be a formula of T and set $\theta(x; y, z) = \varphi(x; y) \wedge \neg\varphi(x; z)$. Let $\langle P_n \rangle$ be the characteristic sequence of (T, θ). The following are equivalent:*

1. *$\langle P_n \rangle$ has an $(\omega, 2)$-diagram.*

2. *$R(x = x, \varphi(x; y), 2) \geq \omega$, i.e. φ is unstable.*

3. *$R(x = x, \theta(x; yz), 2) \geq \omega$, i.e. θ is unstable.*

Proof. (2) $\rightarrow$ (1): We have in hand a tree of partial φ-types $\mathcal{R} = \{p_\nu : \nu \in 2^\omega\}$, partially ordered by inclusion, witnessing that $R(x = x, \varphi, 2) \geq \omega$. Let us show that we can build an $(\omega, 2)$-diagram. That is, we shall choose parameters $\{a_\eta : \eta \in 2^{<\omega}\} \subset P_1$ satisfying Definition 2.18(1).

First, by the definition of the rank R, which requires the partial types to be explicitly contradictory, we can associate to each ν an element $c_\nu \in M$, $\ell(c_\nu) = \ell(y)$ such that:

- $\varphi(x; c_\nu) \in p_{\nu^\frown 1} \setminus p_\nu$, and
- $\neg\varphi(x; c_\nu) \in p_{\eta^\frown 0} \setminus p_\eta$.

i.e., the split after index ν is explained by $\varphi(x; c_\nu)$.

Second, choose a set of indices $\mathcal{S} \subseteq 2^{<\omega}$ such that:

- $(\forall \eta \in 2^{<\omega})\ (\exists s \in \mathcal{S})(\eta \subsetneq s)$
- $(\forall s_1 \subsetneq s_2 \in \mathcal{S})\ (\exists \eta \notin \mathcal{S})\ (s_1 \subsetneq \eta \subsetneq s_2)$

It will suffice to define $a_{s^\frown i}$ for $s \in \mathcal{S}$, $i \in \{0, 1\}$. (The sparseness of S ensures the chosen parameters for φ won't overlap, which will make renumbering straightforward.) Recall that the a_η will be parameters for $\theta(x; y, z) = \varphi(x; y) \wedge \neg(x; z)$. So we define:

- $a_{s^\frown 0} = (c_{s^\frown 0}, c_s)$;
- $a_{s^\frown 1} = (c_s, c_{s^\frown 1})$.

The consistency of the paths through our $(\omega, 2)$-diagram is inherited from the tree $\mathcal{R}$ of consistent partial types. However, $\neg P_2(a_{s^\frown 0}, a_{s^\frown 1})$ because these contain an explicit contradiction:

$$\neg \exists x \Big((\varphi(x; c_{s^\frown 0}) \wedge \neg \varphi(x; c_s)) \wedge (\varphi(x; c_s) \wedge \neg(\varphi(x; c_{s^\frown 1}) \Big)$$

(1) $\rightarrow$ (3): Reading off the parameters from the diagram we obtain a tree of consistent partial θ-types $\{p_\eta : \eta \in 2^{<\omega}\}$, partially ordered by inclusion. For any $\eta \in 2^{<\omega}$, $\neg P_2(a_{\eta^\frown 0}, a_{\eta^\frown 1})$, i.e. $\neg \exists x (\theta(x; a_{\eta^\frown 0}) \wedge \theta(x; a_{\eta^\frown 1}))$. Furthermore, $\theta(x; a_{\eta^\frown 0}) \in p_{\eta^\frown 0} \setminus p_\eta$, while $\theta(x; a_{\eta^\frown 1}) \in p_{\eta^\frown 1} \setminus p_\eta$. So there is no harm in making the types explicitly inconsistent, as the rank R requires, by adding $\neg\ \theta(x; a_{\eta^\frown i})$ to $p_{\eta^\frown j}$ for $i \neq j < 2$.

(2) $\leftrightarrow$ (3): for all A, $|A| \geq 2$, $|S_\varphi(A)| = |S_\theta(A)|$. $\square$

Claim 2.21. *Let φ be a formula of T and set $\theta(x; y, z) = \varphi(x; y) \wedge \neg \varphi(x; z)$. Let $\langle P_n \rangle$ be the characteristic sequence of (T, θ). The following are equivalent:*

1. *$\langle P_n \rangle$ has an $(\omega, 2, 1)$-array.*

2. *φ has the independence property.*

3. *θ has the independence property.*

Proof. (1) $\rightarrow$ (3): This is Observation 2.47. (Essentially, let A_0 be the top row of the array A, and $\sigma, \tau \subset A$ finite disjoint; let $B \subset A$ be a maximal positive base set, i.e. a maximal P_∞-complete graph, in A containing σ and avoiding τ. Then any realization of the type corresponding to B is a witness to this instance of independence.)

(2) $\rightarrow$ (1): Let $\langle i_\ell : \ell < \omega \rangle$ be a sequence over which φ has the independence property. For $t < 2, j < \omega$ set $a^0_j = (i_\ell, i_{\ell+1})$, $a^1_j = (i_{\ell+1}, i_\ell)$. Then $\{a^t_j : t < 2, j < \omega\}$ is an $(\omega, 2, 1)$-array for P_∞.

(3) $\rightarrow$ (2): For any infinite A, $|S_\varphi(A)| = |S_\theta(A)|$, as any type on one side can be presented as a type on the other. The independence property can be characterized in terms of the cardinality of the space of types over finite sets ([26] Theorem II.4.11). $\square$

Claim 2.22. *Let φ be a formula of T and set $\theta(x; y, z) = \varphi(x; y) \wedge \neg\varphi(x; z)$. Let $\langle P_n \rangle$ be the characteristic sequence of (T, θ). Suppose that T does not have SOP_2. Then the following are equivalent:*

1. *$\langle P_n \rangle$ has an $(\omega, \omega, 1)$-array.*

2. *φ has the 2-tree property.*

Proof. (1) $\rightarrow$ (2) Each column (=empty graph) of the array witnesses that φ is 2-dividable, and the condition that any subset of the array containing no more than one element from each column is a P_∞-complete graph ensures that the dividing can happen sequentially.

(2) $\rightarrow$ (1) By Theorem J above, $NSOP_2$ implies φ has TP_2. That is, there is a tree of instances $\{\varphi(x; a_\eta) : \eta \in \omega^{<\omega}\}$ such that first, for any finite n, $\eta_1 \subseteq \cdots \subseteq \eta_n$ implies that the partial type $\{\varphi(x; a_{\eta_1}), \ldots \varphi(x; a_{\eta_n})\}$ is consistent; and second,

$$\neg\exists x \left(\varphi(x; a_\eta) \wedge \varphi(x; a_\nu)\right) \iff (\exists \rho \in \omega^{<\omega})(\exists i \neq j \in \omega) \left(\eta = \rho^\frown i \ \wedge \ \nu = \rho^\frown j\right)$$

Thus the parameters $\{a_\eta : \eta \in \omega^{<\omega}\} \subset P_1$ form an $(\omega, \omega, 1)$-array for P_∞. $\square$

It is straightforward to characterize the analogous k-tree properties in terms of arrays whose columns are k-consistent but $(k+1)$-inconsistent.

Claim 2.23. *The following are equivalent:*

1. *$\langle P_n \rangle$ has an $(\omega, 2)$-tree.*

2. *φ has SOP_2.*

Proof. (2) $\rightarrow$ (1) This is a direct translation of Definition 0.11.

(1) $\rightarrow$ (2) It suffices to show that $\langle P_n \rangle$ has an (ω, ω)-tree, which is true by compactness, using the strictness of the tree. □

In Chapter 3, §3.5 we will consider a distinguished configuration:

Definition 2.24. *P_∞ has the* compatible order property *if there exists a sequence $C = \langle a_i, b_i : i < \omega \rangle \subset P_1$ such that for any $n < \omega$ and any $a_1, b_1, \ldots a_n, c_n \subset C$,*

$$P_n((a_1, b_1), \ldots (a_n, c_n)) \iff (\max\{a_1, \ldots a_n\} < \min\{b_1, \ldots b_n\})$$

Say that P_m has the compatible order property to indicate that this holds for all P_n, $n \leq m$.

Observation 2.25. *Suppose $(T, \varphi) \mapsto \langle P_n \rangle$, and that $\langle P_n \rangle$ has the compatible order property. Then φ_2 has the tree property, and in particular, SOP_2.*

Proof. Let us build an SOP_2-tree $\{\varphi_2(x; a_\eta, b_\eta) : \eta \in \omega^{<\omega}\}$ following Definition 0.11 above by specifying the corresponding tree of parameters $\{c_\eta : \eta \in \omega^{<\omega}\} \subset P_1$, where each c_η is a pair (a_η, b_η). Let $S = \langle a_i b_i : i < \mathbb{Q}\rangle$ be an indiscernible sequence witnessing the compatible order property. We will use two facts in our construction:

1. Let $\langle a_{i_\ell} b_{j_\ell} : \ell < \omega\rangle$ be any subsequence of S such that $\ell < k \implies a_{i_\ell} < b_{j_\ell} < a_{i_k} < b_{j_k}$. Then $\{\varphi_2(x; a_{i_j}, b_{i_j}) : j < \omega\}$ 2-divides by Observation 2.30.

2. Let $a_{i_1}, b_{j_1}, \ldots a_{i_n}, b_{j_n} \in S$. Then

$$P_n((a_{i_1}, b_{j_1}), \ldots (a_{i_n}, b_{j_n})) \iff \max\{i_1, \ldots i_n\} < \min\{j_1, \ldots j_n\}$$

so in particular

$$P_2((a_{i_1}, b_{j_1}), (a_{i_2}, b_{j_2})) \iff \max\{i_1, i_2\} < \min\{j_1, j_2\}$$

Let $\eta \in \omega^{<\omega}$ be given and suppose that either c_η has been defined or $\eta = \emptyset$. If c_η has been defined, it will be (a_i, b_j) for some $i < j \in \mathbb{Q}$. Let $\langle k_\ell : \ell < \omega\rangle$ be any ω-indexed subset of $(i, j) \cap \mathbb{Q}$, or of $\mathbb{Q}$ if $\eta = \emptyset$. Define $c_{\eta^\frown \ell} = (a_{k_\ell}, b_{k_{\ell+1}})$. Now suppose we have defined the full tree of parameters c_η in this way. By fact (1) we see that immediate successors of the same node are P_2-inconsistent. By fact $(2)_n$, paths are consistent, while by fact $(2)_2$, any two *incomparable (Definition 0.11) elements c_ν, c_η are P_2-inconsistent. $\square$

Remark 2.26. *The significance of the compatible order property is suggested by Chapter 3, Observation 3.28, along with Theorem 3.32.*

2.4 Localization and persistence

A major goal of these methods is to analyze φ-types, and thus to concentrate on the combinatorial structure which is "close to" or "inseparable from" the complete graph A representing a consistent partial φ-type under analysis.

P_n asks about incidence relations on a set of parameters; it will be useful to definably restrict the witness and parameter sets. For instance:

- we may ask that the witnesses lie inside certain instances of φ, e.g. by setting $P_1'(y) = \exists x(\varphi(x;y) \wedge \varphi(x;a))$, i.e. $P_1' = P_2(y,a)$.
- we may ask that the parameters be consistent 1-point extensions (in the sense of some P_n) of certain finite graphs C. For instance, we might define $P_1''(y) = P_1(y) \wedge P_2(y,c_1) \wedge P_3(y,c_2,c_2)$.

The next definition gives the general form.

Definition 2.27. (Localization) *Fix a characteristic sequence $(T,\varphi) \rightarrow \langle P_n \rangle$, and choose $B, A \subset M \models T$ with A a positive base set and $A = \emptyset$ possible.*

1. (the localized predicate P_n^f) *A* localization P_n^f *of the predicate $P_n(y_1, \dots y_n)$ around the positive base set A with parameters from B is given by a finite sequence of triples $f : m \rightarrow \omega \times \mathcal{P}_{\aleph_0}(y_1, \dots y_n) \times \mathcal{P}_{\aleph_0}(B)$ where $m < \omega$ and:*
   - *writing $f(i) = (r_i, \sigma_i, \beta_i)$ and $\check{s}$ for the elements of the set s, we have:*

$$P_n^f(y_1, \dots y_n) := \bigwedge_{i \leq m} P_{r_i}(\check{\sigma}_i, \check{\beta}_i)$$

- *for each $\ell < \omega$, T_1 implies that there exists a P_ℓ-complete graph C_ℓ such that P^f_n holds on all n-tuples from C_ℓ. If this last condition does not hold, P^f_n is a* trivial localization. *By* localization *we will always mean non-trivial localization.*

- *In any model of T_1 containing A and B, P^f_n holds on all n-tuples from A.*

Write $\mathrm{Loc}^B_n(A)$ *for the set of localizations of P_n around A with parameters from B (i.e. nontrivial localizations, even when $A = \emptyset$).*

2. (the localized formula φ^f) *For each localization P^f_n of some predicate P_n in the characteristic sequence of φ, define the corresponding formula*

$$\varphi^f_n(x; y_1, \dots y_n) := \varphi_n(x; y_1, \dots y_n) \wedge P^f_n(y_1, \dots y_n)$$

When $n = 1$, write $\varphi^f = \varphi^f_1$. Let $S^f_\varphi(N)$ denote the set of types $p \in S_\varphi(N)$ such that for all $\{\varphi_{i_1}(x; c_{i_1}), \dots \varphi_{i_n}(x; c_{i_n})\} \subset p$, $P^f_n(c_{i_1}, \dots c_{i_n})$. Then there is a natural correspondence between the sets of types

$$S^f_\varphi(N) \leftrightarrow S_{\varphi^f}(N)$$

3. (the *localized formula $\varphi^{f+\bar{a}}$) *We have thus far described localizations of the parameters of φ. We will also want to consider restrictions of the possible witnesses to φ by adjoining instances of φ_k. That is, set*

$$\varphi^{f+\bar{a}}(x; y) = \varphi^{f+a_1,\dots a_k}(x; y) := \varphi(x; y) \wedge P^f_1(y) \wedge \varphi_k(x; a_1, \dots a_k)$$

where, as indicated, $k = \ell(\overline{a})$. The * is to emphasize that this is really the construction from φ of a new, though related, formula, which will have its own characteristic sequence, given by:

4. (the *localized characteristic sequence $\langle P_n^{f+\overline{a}} : n < \omega \rangle$) *The sequence $\langle P_n^{f+\overline{a}} : n < \omega \rangle$ associated to the formula $\varphi^{f+\overline{a}}$ is given by, for each $n < \omega$,*

$$P_n^{f+\overline{a}}(y_1, \ldots y_n) = \bigwedge_{i \leq n} P_1^f(y_i) \wedge P_{n+k}(y_1, \ldots y_n, a_1, \ldots a_k)$$

When f or $\overline{a}$ are empty, we will omit them.

Remark 2.28. *Convention 2.3 applies: that is, localization is not essentially dependent on the choice of model M. See Definition 2.36 (Persistence) and the observation following.*

As a first example of the utility of localization, notice that when φ is simple we can localize to avoid infinite empty graphs.

Observation 2.29. *Fix a positive base set A for the formula ψ, possibly empty. When ψ does not have the tree property, then for each $k < \omega$ there is a finite set C over which ψ is not k-dividable. As a consequence, if ψ does not have the tree property, then for each predicate P_n there is a localization around A on which there is a uniform finite bound on the size of a P_n-empty graph. We can clearly also choose the localizations so that none of $\psi_1, \ldots \psi_\ell$ are k-dividable for any finite k, ℓ fixed in advance. When ψ is low, we can do this for all k at once (Definition 1.19).*

Proof. This is the proof that $D(x = x, \psi, k) < \omega$ for any simple formula ψ; see for instance [13]. $\square$

2.4.1 Stability in the parameter space

The classification-theoretic complexity of the formulas P_n is often strictly less than that of the original theory T. Note that the results here refer to the *formulas* $P_n(y_1, \ldots y_n)$, not necessarily to their full theory T_1.

Observation 2.30. *Suppose* $(T, \varphi) \mapsto \langle P_n \rangle$. *If* $P_2(x; y)$ *has the order property then* $\varphi(x; y) \wedge \varphi(x; z)$ *is 2-dividable.*

Proof. Let $\langle a_i, b_i : i < \omega \rangle$ be a sequence witnessing the order property for P_2, so $P_2(a_i, b_j)$ iff $i < j$. This means that $\exists x(\varphi(x; a_i) \wedge \varphi(x; b_j))$ iff $i < j$. So $\varphi(x; a_i) \wedge \varphi(x; b_{i+1})$ are consistent for each i, but the set $\{\varphi(x; a_i) \wedge \varphi(x; b_{i+1}) : i < \omega\}$ is 2-inconsistent. $\square$

Remark 2.31. *By compactness, without loss of generality the sequence of Observation 2.30 can be chosen to be* $(T-)$*indiscernible, and so actually witnesses the dividing of some instance of* φ_2.

Note that the converse of Observation 2.30 fails: for $\varphi(x; y) \wedge \varphi(x; z)$ *to divide it is sufficient to have a disjoint sequence of "matchsticks" in* P_2 *(i.e.* $(a_i, b_i) : i < \omega$ *such that* $P_2(a_i, b_j)$ *iff* $i = j$*), without the additional consistency which the order property provides.*

But see Chapter 3, Theorem 3.19, which illuminates the role of the order.

Observation 2.32. *Suppose that $(T, \varphi) \mapsto \langle P_n \rangle$, and for some n, k and some partition of $y_1, \ldots y_n$ into k object and (n-k) parameter variables, $P_n(y_1, \ldots y_k; y_{k+1}, \ldots y_n)$ has the order property. Then $\varphi_n(x; y_1, \ldots y_n)$ is 2-dividable.*

Proof. The proof is analogous to that of Observation 2.30, replacing the a_i by k-tuples and the b_j by (n-k)-tuples. $\square$

Thus in cases where we can localize to avoid dividing of φ, we can assume any initial segment of the associated predicates P_n are stable:

Conclusion 2.33. *For each formula φ and for all $m < \omega$, if φ_{2n} does not have the tree property, then for each positive base set A there are a finite B and $P_1^f \in \mathrm{Loc}_1^B(A)$ over which $P_2, \ldots P_n$ do not have the order property. In particular, this holds if T is simple.*

By way of motivating the next subsection, let us prove the contrapositive: If the order property in P_2 persists under repeated localization, then φ has the tree property. Compare the proof of Observation 2.25 above. Without the compatible order property, we cannot ensure the tree is strict. While that argument built a tree out of a set of parameters which were given all at once (a so-called "static" argument), the following "dynamic" argument must constantly localize to find subsequent parameters, so cannot ensure that elements in different localizations are inconsistent.

Lemma 2.34. *Suppose that in every localization of P_1 (around $A = \emptyset$), P_2 has the order property. Then φ_2 has the tree property.*

Proof. Let us describe a tree with nodes (c_η, d_η), $(\eta \in \omega^{<\omega})$, such that:

1. for each $\rho \in \omega^\omega$, $\{c_\eta, d_\eta : \eta \subseteq \rho\}$ is a complete P_∞-graph, where $\subseteq$ means initial segment.

2. for any $\nu \in \omega^{<\omega}$, $P_2(c_{\eta^\frown i}, d_{\eta^\frown j}) \iff i \leq j$.

For the base case ($\eta \in \omega^1$), let $\langle c_i, d_i : i \in \omega\rangle$ be an indiscernible sequence witnessing the order property (so $P_2(c_i, d_j) \iff i \leq j$) and assign the pair (c_i, d_i) to node i.

For the inductive step, suppose we have defined (c_η, d_η) for $\eta \in \omega^n$. Write $E_\eta = \{(c_\nu, d_\nu) : \nu \leq \eta\}$ for the parameters used along the branch to (c_η, d_η). Using $\check{x}$ to mean the elements of the set x, let $P_1^{f_\eta}$ be given by $P_{n+1}((y,z), \check{E}_\eta)$. Let $\langle a_j, b_j : j \in \omega\rangle$ be an indiscernible sequence witnessing the order property inside this localization, and define $(c_{\eta^\frown i}, d_{\eta^\frown i}) := (a_j, b_j)$.

Finally, let us check that this tree of parameters witnesses the tree property for φ_2. On one hand, the order property in P_2 ensures that for each $n \in \omega^{<\omega}$, the set

$$\{\varphi_2(x; c_{\eta^\frown i}, d_{\eta^\frown i}) : i \in \omega\}$$

is 1-consistent but 2-inconsistent. On the other hand, the way we constructed each localization $P_1^{f_\eta}$ ensured that each path was a complete P_∞-graph, thus naturally a complete P'_∞-graph, where $\langle P'_n \rangle$ is the characteristic sequence of the conjunction φ_2. $\square$

Remark 2.35. *Example 2.12 (page 61) shows that the condition that φ has the tree property is necessary, but not sufficient, for the order property in P_2 to be* persistent*, Definition 2.36 below.*

Question 2.1. *Is SOP_2 sufficient?*

Compare Theorem 3.32.

2.4.2 Persistence

Definition 2.36. (Persistence) *Fix $(T, \varphi) \mapsto \langle P_n \rangle$, $M \models T$ sufficiently saturated, and a positive base set A, possibly $\emptyset$. Let X be a T_0-configuration, possibly infinite. Then X is* persistent *around the positive base set A if for all finite $B \subset M$ and for all $P_1^f \in Loc_1^B(A)$, P_1^B contains witnesses for X.*

We will write X is A-persistent to indicate that X is persistent around A.

Note 2.37. Persistence asks whether all finite localizations around A contain witnesses for some T_0-configuration X. The predicates P_n mentioned in X are, however, not the localized versions. We have simply restricted the set from which witnesses can be drawn. This is an obvious but important point: for instance, in the proof of Lemma 2.45 below it is important that the sequence of P_2-inconsistent pairs found inside of successive localizations $P_1^{f_n}$ are P_2-inconsistent in the sense of T_1.

Observation 2.38. (Persistence is a property of T_A) *The following are equivalent, fixing $T, \varphi, \langle P_n \rangle$, $A \subset \mathcal{M}$ a small positive base set in the monster model, and a T_0-configuration X. Write $P_1^f(M)$ for the set which P_1^f defines in the model M.*

1. *In some sufficiently saturated model $M \models T_1$ which contains A, X is persistent around A in M. That is, for every finite $B \subset M$ and every localization $P_1^f \in \mathrm{Loc}_1^B(A)$, there exist witnesses to X in $P_1^f(M)$.*

2. *In every model $N \models T_1$, $N \supset A$, for every finite $B \subset N$, every localization $P_1^f \in Loc_1^B(A)$, and every finite fragment X_0 of X, $P_1^f(N)$ contains witnesses for X_0.*

Proof. (2) $\implies$ (1) Compactness.

(1) $\implies$ (2) Suppose not, letting $P_1^f \in Loc_1^B(A)$ and X_0 witness this. To this P_1^f we can associate a T_1-type $p(y_1, \ldots y_{|B|}) \in S(A)$ which says that the localization given by f with parameters $y_1, \ldots y_{|B|}$ contains A but implies that X_0 is inconsistent. But any sufficiently saturated model containing A will realize this type, and thus contain such a localization. □

To reiterate Convention 2.3, then, we may, as a way of speaking, call a configuration "persistent" while working in some fixed sufficiently saturated model, but we always refer to the corresponding property of T.

Corollary 2.39. *Persistence around the positive base set A remains a property of T in the language with constants for A.*

In Section 2.5 below we will see that stability and simplicity can be naturally characterized in terms of persistence.

Finally, let us check the (easy) fact that persistence of some T_0-configuration around $\emptyset$ in some given sequence $\langle P_n \rangle$ implies its persistence around any positive base set A for that sequence.

Fact 2.40. *Suppose that X is an $\emptyset$-persistent T_0-configuration in the characteristic sequence $\langle P_n \rangle$ and A is a positive base set for $\langle P_n \rangle$. Then X remains persistent around A.*

Proof. Let $p(x_0, \dots) \in \mathcal{L}(=, P_1, P_2, \dots)$ describe the type, in V_X-many variables, of the configuration $X = (V_X, E_X)$. Let $q(y) \in S(A)$ be the type of a 1-point P_∞-extension of A in the language $\mathcal{L}_0 = \{P_n : n < \omega\} \cup \{=\}$. We would like to know that $q(x_0), q(x_1), \dots, p(x_0, \dots)$ is consistent, i.e., that we can find, in some given localization, witnesses for X from among the elements which consistently extend A. If not, for some finite subset $A' \subset A$, some $n < \omega$, and some finite fragments q' of $q|_{A'}$ and p' of p,

$$q'(x_0) \cup \cdots \cup q'(x_n) \vdash \neg p'(x_0, \dots x_n)$$

But now localizing P_1 according to the conditions on the lefthand side (which are all positive conditions involving the P_n and finitely many parameters A') shows that X is not persistent, contradiction. $\square$

2.5 Dividing lines: Stability and simplicity

In this section we will see that the language of localization and persistence naturally characterizes the model-theoretic dividing lines of stability and simplicity.

2.5.1 Stability: the case of P_2

The argument in this technically much simpler case will generalize without too much difficulty. We first revisit an avatar of the independence property.

Definition 2.41. ($(\omega, 2)$-arrays revisited)

1. *The predicate P_n is $(\omega, 2)$ if there is $C := \{a^t_i : t < 2, i < \omega\}$ such that for all $\ell \leq n$, any ℓ-element subset C_0,*

$$P_\ell(C_0) \quad \iff \quad \big(a^t_i, a^s_j \in C_0 \implies (i \neq j) \vee (t = s)\big)$$

2. *If for all $n < \omega$, P_n is $(\omega, 2)$, we say that P_∞ is $(\omega, 2)$.*

3. *A* path *through the $(\omega, 2)$-array A is a set $X \subset A$ which contains no more than one element from each column. Paths will be positive base sets.*

Remark 2.42. *If P_∞ is $(\omega, 2)$, then φ has the independence property.*

Proof. This is a special case of Observation 2.47 below. □

Lemma 2.43. *If φ is stable then there is a finite localization P^f_1 for which TFAE:*

1. *There exists $X \subset P^f_1$, X an $(\omega, 2)$-array wrt P_2*

2. There exists $Y \subset P_1^f$, Y an $(\omega, 2)$-array wrt P_∞

Proof. Choose the localization P_1^f according to Observation 2.29 so that neither φ nor φ_2 divide on parameters from P_1^f. This is possible because stable formulas are simple and low (Definition 1.19). Let $Z = \langle c_i^t : t < 2, i < \omega \rangle \subset P_1^f$ be an indiscernible sequence of pair which is an $(\omega, 2)$-array for P_2. Each of the sub-sequences $\langle c_i^0 : i < \omega \rangle$, $\langle c_i^1 : i < \omega \rangle$ is indiscernible, so will be either P_2-complete or P_2-empty; by choice of P_1^f, they cannot be empty.

It remains to show that any path $X \subset Z$ is a P_∞-complete graph. Suppose not, and let n be minimal so that the n-type of some increasing sequence of elements $z_1^{t_1}, \dots z_n^{t_n}$ implies $\neg \exists x (\bigwedge_{i<n} \varphi(x; z_i^{t_i}))$. Choose an infinite indiscernible subsequence of pairs $Z' \subset Z^2$ of the form $\langle c_i^0, c_{i+1}^1 : i \in W \subset \omega \rangle$. Then the set $\{\varphi(x; c_i^0) \wedge \varphi(x; c_{i+1}^1) : i \in W\}$ will be 1-consistent by definition but n-inconsistent by assumption, contradicting the hypothesis that φ_2 is not dividable in the localization P_1^f. $\square$

Lemma 2.44. *Suppose every localization P_1^f around some fixed positive base set A contains elements y, z such that $\neg P_2(y, z)$. Then P_∞ is $(\omega, 2)$.*

Proof. Choose $P_1^{f_0}$ to be any localization given by the previous lemma. We construct an $(\omega, 2)$-array as follows.

At stage 0, let c_0^0, c_0^1 be any pair of P_2-incompatible elements each of which is a consistent 1-point extension of A in $P_1^{f_0}$. At stage $n+1$, write C_n for $\{c_i^t : t < 2, i \leq n\}$ and suppose we have defined $P_1^{f_n} \in \mathrm{Loc}_1^{C_n}(A)$. By hypothesis, there are $c_{n+1}^0, c_{n+1}^1 \in$

$P_1^{f_n}$ such that $\neg P_2(c^0_{n+1}, c^1_{n+1})$ and such that each c^i_{n+1} is a consistent 1-point extension of A (Fact 2.40). Let $C_{n+1} = C_n \cup \{c^0_{n+1}, c^1_{n+1}\}$ and define $P_1^{f_{n+1}} \in \mathrm{Loc}_1^{C_{n+1}}(A)$ by

$$P_1^{f_{n+1}}(y) = P_1^f(y) \wedge P_2(y; c^0_{n+1}) \wedge P_2(y; c^1_{n+1})$$

Thus we construct an $(\omega, 2)$-array for P_2, as desired. Applying Lemma 2.43 we obtain an $(\omega, 2)$-array for P_∞. □

In the next Corollary, an "empty pair" is the T_0-configuration given by $V_x = 2, E_x = \{\{1\}, \{2\}\}$, in the notation of Convention 2.5.

Corollary 2.45. *Suppose an empty pair is persistent around A. Then φ has the independence property, so it follows that φ is unstable.*

Proof. By Lemma 2.44, P_∞ is $(\omega, 2)$, which by Remark 2.42 implies that φ has the independence property. □

2.5.2 Stability: the case of k

The generalization will be straightforward once the right definitions are in place. We will show that, in each P_n $(n \geq 2)$, a persistent n-tuple of P_n-inconsistent elements produces a so-called (ω, n)-array, which will imply the independence property for a formula close to φ and thus for φ itself. We first establish a "sharpness lemma" which gives a normal form for the arrays when $n > 2$. Recall that the $P_n^{\overline{a}}$ are *localized predicates from Definition 2.27.

Definition 2.46. ((ω, n)-arrays revisited) *Assume $n \leq \ell$. Compare Definition 2.18; here, the possible ambiguity of the amount of consistency will be important.*

1. *The predicate P_ℓ is (ω, n) if there is $C = \{c^t_i : t < n, i < \omega\} \subset P_1$ such that, for all $c^{t_1}_{i_1}, \dots c^{t_\ell}_{i_\ell} \in C$,*

   - *ℓ-tuples from ℓ distinct columns are consistent, i.e.*

   $$\bigwedge_{j,k \leq \ell} i_j \neq i_k \implies P_\ell(c^{t_1}_{i_1}, \dots c^{t_\ell}_{i_\ell})$$

   - *and no column is entirely consistent, i.e. for all $\sigma \subset \ell$, $|\sigma| = n$,*

   $$\bigwedge_{j,k \in \sigma} i_j = i_k \implies \neg P_\ell(c^{t_1}_{i_1}, \dots c^{t_\ell}_{i_\ell})$$

   Any such C is an (ω, n)-array. *The precise arity of consistency is not specified, see condition (4).*

2. *If for all $n \leq \ell < \omega$, P_ℓ is (ω, n), say that P_∞* is *(ω, n).*

3. *A* path *through the (ω, n) array C is a set $X \subset C$ which contains no more than n-1 elements from each column.*

4. *P_ℓ is* sharply *(ω, n) if it contains an (ω, n)-array C on which, moreover, for all $\{c^{t_1}_{i_1}, \dots c^{t_\ell}_{i_\ell}\} \subset C$*

   $$P_\ell(c^{t_1}_{i_1}, \dots c^{t_\ell}_{i_\ell}) \iff \bigwedge_{\sigma \subset \ell, |\sigma| = n} \left(\bigwedge_{j,k \in \sigma} i_j = i_k \implies \bigvee_{j \neq k \in \sigma} t_j = t_k \right)$$

   i.e., if every path is a P_ℓ-complete graph.

5. *P_∞ is sharply (ω, n) if P_ℓ is sharply (ω, n) for all $n \leq \ell < \omega$.*

Observation 2.47. *If P_∞ is sharply (ω, k) then φ_{k-1} has the independence property.*

Proof. Let $X = \langle a^1_i, \dots a^k_i : i < \omega \rangle$ be the array in question; then φ_{k-1} has the independence property on any maximal path, e.g. $B := \langle a^1_i, \dots a^{k-1}_i : i < \omega \rangle$. To see this, fix any $\sigma, \tau \subset \omega$ finite disjoint; then by the sharpness hypothesis $\{a^1_i, \dots a^{k-1}_i : i \in \sigma\} \cup \{a^2_j, \dots a^k_j : j \in \tau\}$ is a P_∞-complete graph and thus corresponds to a consistent partial φ-type q. But any realization α of q cannot satisfy $\varphi(x; a^1_j)$ for any $j \in \tau$, because P_k does not hold on the columns. A fortiori $\neg\varphi_k(\alpha; a^0_j, \dots a^{k-1}_j)$. $\square$

Let us write down some conventions for describing types in an array.

Definition 2.48. *Let x^t_i, x^s_j be elements of some (ω, n)-array X.*

1. *Let $I(x^t_i) = \{x^s_j \in X : j = i\}$, i.e. the elements in the same column as x^t_i.*

2. *Let $X_0 = \{x^{t_1}_{i_1}, \dots x^{t_\ell}_{i_\ell}\} \subset X$ be a finite subset. The* column count *of $\{x^{t_1}_{i_1}, \dots x^{t_\ell}_{i_\ell}\}$ is the unique tuple $(m_1, \dots m_\ell) \in \omega^\ell$ such that:*

   - *$m_i \geq m_{i+1}$ for each $i \leq \ell$*
   - *$\Sigma_i\ m_i = \ell$*
   - *if $Y_0 = \{y_1, \dots y_r\}$ is a maximal subset of X_0 such that $y, z \in Y_0, y \neq z \rightarrow y \notin I(z)$, then some permutation of*

   $$(|I(y_1) \cap X_0|, \dots, |I(y_r) \cap X_0|)$$

   is equal to $(m_1, \dots m_\ell)$.

In other words, we count how many elements have been assigned to each column, and put these counts in descending order of size. Write col-ct$(\overline{x})$ *for this tuple.*

3. *Let $\leq$ be the lexicographic order on column counts, i.e. $(1, \dots 1) < (2, 1, \dots)$. This is a discrete linear order, so we can define $(m_1, \dots m_\ell)^+$ to be the immediate successor of $(m_1, \dots m_\ell)$ in this order. Define* $\operatorname{gap}((m_1, \dots m_\ell)) = m_i$ *where* $((n_1, \dots n_\ell)^+ = (m_1, \dots m_\ell)$ *and* $\forall j \neq i \; m_j = n_j$, *i.e. the value which has just incremented.*

Recall that if φ is stable then φ is simple and low. By analogy to Lemma 2.43,

Lemma 2.49. *(Springboard lemma) Suppose that φ is simple and low, and let $\langle P_n \rangle$ be the characteristic sequence of (T, φ). For each $n < \omega$, there exists a localization P_1^f of P^1 in which the following are equivalent:*

1. *P_1^f contains a sharp (ω, n)-array for P_{2n-2}.*

2. *P_1^f contains a sharp (ω, n)-array for P_∞.*

Proof. Assume (1), so let $C = \{c_i^t : t < n, i < \omega\} \subset P_1^f$ be sharply (ω, n) for P_{2n-2}, chosen without loss of generality to be an indiscernible sequence of n-tuples. Fix a path $Y = y_1, \dots y_m$ of minimal size $m > n$ such that $\neg P_m(y_1, \dots y_m)$. Let $S := \{c_i^0, \dots c_i^{n-1}, c_{i+1}^1, \dots c_{i+1}^n : i < \omega\} \subset C^{2n-2}$ be a sequence of pairs of offset $(n-1)$-tuples.

Note that S is 1-consistent as we assumed (1).

On the other hand, C is indiscernible, so any increasing sequence of m elements from S will cover all the possible m-types from C. Since Y is inconsistent, this implies that S is m-inconsistent. Thus φ_{2n-2} is m-dividable.

The appropriate localization is thus one in which (for good measure) none of the finite set of formulas $\{\varphi_{2\ell-2} : 1 \leq \ell \leq n\}$ are dividable. By lowness, each of these formulas has a uniform finite bound n_ℓ on the arity of its dividing. By simplicity, the $D(x = x, \varphi_{2\ell-2}, n_\ell)$-ranks are finite, so the desired localization exists. □

Remark 2.50. *Suppose that φ is simple and low, and let P_1^f be the localization from Lemma 2.49. Suppose that there is an (ω, n)-array C for P_n, and $a_1, \ldots a_r \in C$, with $r < n - 1$ and $\ell + r = n$. Then the following are equivalent:*

1. *P_1^f contains a sharp (ω, ℓ)-array $D \subset C$ for $P_{2\ell-2}^{\{a_1,\ldots a_r\}}$.*

2. *P_1^f contains a sharp (ω, ℓ)-array for $P_\infty^{\{a_1,\ldots a_r\}}$, i.e. for $\langle P_\ell^{\{a_1,\ldots a_r\}} : \ell < \omega \rangle$.*

Proof. The remark asks us to reprove the lemma for the formula $\varphi^{\overline{a}}$ and its associated characteristic sequence; because of the additional hypotheses about c and ℓ, this proof is contained in the previous argument. □

Lemma 2.51. (Sharpness lemma) *Suppose that φ is simple and low and that in some localization P_1^f, P_n is (ω, n). Then there exists $0 \leq r < n - 1$ and a finite tuple $a_1, \ldots a_r$ (empty if $r = 0$) such that the predicate $P_{2\ell-2}^{\overline{a}}$ is sharply (ω, ℓ), where $\ell = n - r$ and $P_{2\ell-2}^{\overline{a}}$ is the *localized predicate in the sense of Definition 2.27. (When $r = 0$, this is just P_{2n-2}.)*

Proof. Let $C = \{c_i^t : t < n, i < \omega\} \subset P_1^f$ be an indiscernible sequence of n-tuples which is an (ω, n)-array for P_n. We will systematically check all k-tuples (for $k \leq 2n-2$, and which do not contain an entire column) for consistency, by inducting on column count. Column count does not entirely determine the type of a tuple, of course, but it is close enough for our purposes. For notation, let g be an enumeration, in lexicographic order, of all possible values of col-ct(X_0) which occur on paths $X_0 \subset C$, $|X_0| \leq 2n-2$.

At stage $s+1$, suppose we are considering r-tuples with column count $f(s)$, for some $1 < r \leq 2n-2$. If for all $c_1, \ldots c_r \subset C$, col-ct$(c_1, \ldots c_r) = f(s) \implies P_r(c_1, \ldots c_r)$, then we continue to stage $s+2$; and if we reach $2n-2$ this way, we have proved the lemma with $r = 0$. Otherwise, $f(s)$ is the first column-count to produce an inconsistency, witnessed by $c_1, \ldots c_r \subset C$. Let $\ell = \text{gap}(g(s+1))$, in the notation of Definition 2.48, be the size of the column just incremented, and w.l.o.g. suppose c_1 is in this column. Then we can set $\overline{a} := \{c_i : i \leq r, c_i \notin I(c_1)\}$ to be the elements in the other columns. Notice that by indiscernibility we could have chosen the columns containing these r inconsistent elements to be as far apart as desired.

Let $D_0 \subset C$ be the set of ℓ rows (where $\ell = \text{gap}(g(s+1)) = |I(c_1)|$) containing the elements of $I(c_1)$. Then, by indiscernibility of the array C, $P_\ell^{\overline{a}}$ is w.l.o.g. (ω, ℓ) on some infinite subset $D \subset D_0$; and by inductive hypothesis, it is sharply (ω, ℓ). $\square$

Fact 2.52. *The following are equivalent for a formula $\varphi(x; y)$.*

1. *φ has the independence property.*

2. For some $n < \omega$, φ_n has the independence property.

3. For every $n < \omega$, φ_n has the independence property.

*4. Some *localization $\varphi^{\overline{a}}$ has the independence property.*

Proof. (1) $\rightarrow$ (3) $\rightarrow$ (2) $\rightarrow$ (1) $\rightarrow$ (4) are straightforward: use the facts that the formulas φ_i, φ_j generate the same space of types, and that the independence property can be characterized in terms of counting types over finite sets ([26]:II.4). Finally, (4) $\rightarrow$ (2) as we have simply specified some of the parameters. $\square$

Lemma 2.53. *Suppose that for some $n < \omega$, every localization of P_1 around some fixed positive base set A contains an n-tuple on which P_n does not hold. Then P_n is (ω, n).*

Proof. We proceed just as in the case of $n = 2$. Let P_1^f be any localization, for instance that of Lemma 2.49.

At stage 0, let $c_0^0, c_0^1, \ldots c_0^{n-1} \subset P_1^{f_0} := P_1^f$ be an n-tuple of elements on which P_n does not hold, chosen by Fact 2.40 so that each c_0^i is a consistent 1-point extension [in the sense of P_n] of A. Let $X_0 = \{\{c_0^i\} : i \leq n\}$ be the set of these singletons. Recall that $\check{x}$ denotes the elements of x. Define

$$P_1^{f_1}(y) = P_1^{f_0}(y) \wedge \bigwedge_{x \in X_0} P_2(y; \check{x})$$

which includes A by construction.

At stage $m+1$, write C_m for $\{c_i^t : t < n, i \leq m\}$ and consider the localized set of elements $P_1^{f_m} \in \mathrm{Loc}_1^{C_m}(A)$. Let

$$X_m := \{x \subset C_m : \ |x| = n-1 \text{ and for all i} < \text{m}, \ |x \cap (C_{i+1} \setminus C_i)| \leq 1\}$$

i.e. sets which choose no more than one element from each stage in the construction.

By hypothesis, there are $c_{m+1}^0, \dots c_{m+1}^{n-1} \in P_1^{f_m}$ such that $\neg P_n(c_{m+1}^0, \dots c_{m+1}^{n-1})$ and such that for all $x \in X_m$, each c_{m+1}^i is a consistent 1-point extension of $A \cup x$, in the sense of P_n. Let $C_{m+1} = C_m \cup \{c_{m+1}^0, \dots c_{m+1}^{n-1}\}$, and let X_{m+1} be the sets from C_{m+1} which choose no more than one element from each stage in the construction. We now define $P_1^{f_{m+1}} \in \mathrm{Loc}_1^{C_{m+1}}(A)$ by

$$P_1^{f_{m+1}}(y) = P_1^{f_m}(y) \wedge \bigwedge_{x \in X_{m+1}} P_n(y; \check{x})$$

(If $m < n$, the parameters from $\check{x}$ need not necessarily be distinct.) Again, this localization contains A by construction. Thus we construct an (ω, n)-array for P_n, as desired. $\square$

We are now in a position to prove:

Theorem 2.54. *Fix $(T, \varphi) \mapsto \langle P_n : n < \omega \rangle$ and a positive base set A. Suppose that for some $n < \omega$, there is a T_0-configuration for which $X = n$, $\{1, \dots n\} \notin E_x$ (i.e. a P_n-empty tuple) and X is persistent around A. Then φ has the independence property.*

Proof. We work inside the localization $P_1^f \supset A$ given by Lemma 2.49. As a tuple of P_n-inconsistent elements is persistent, we apply Lemma 2.53 to obtain $Y \subset P_1^f$ on which P_n is (ω, n). By the sharpness lemma (2.51), we obtain a sharp (ω, ℓ)-array for the predicate $P^{\overline{a}}_{2\ell-2}$. (Note that in that lemma, we obtained sharpness at the cost of adding parameters; we can't do this indefinitely, but luckily P_1^f says that $2n-2$ is enough.) Recall that $0 < |\overline{a}| + 1 < n$.

By choice of the localization P_1^f, as we now satisfy condition (1) of Lemma 2.49, we obain a sharp (ω, ℓ)-array for the entire sequence $\langle P^{\overline{a}}_n \rangle$. By Observation 2.47 applied to the *localized formula $\varphi^{\overline{a}}$, a sharp (ω, ℓ)-array means that $\varphi^{\overline{a}}_{\ell-1}$ has the independence property. By Fact 2.52 applied to $\varphi^{\overline{a}}$, $\varphi^{\overline{a}}$ has the independence property. By (4) $\rightarrow$ (1) of the same fact, φ must also have the independence property, so we finish. $\square$

The property that for all n, no nontrivial T_0-configuration is persistent for P_n around the positive base set A characterizes the class of formulas which are stable in some localization around the type corresponding to A:

Theorem 2.55. *Let φ be a formula of T and $\langle P_n : n < \omega \rangle$ its characteristic sequence. In the notation of Definition 2.27,*

1. *If the localization φ^f of φ is stable, then for each P_∞-graph $A \subset P_1^f$ and for each $n < \omega$, there exists a localization $P_1^{f_n} \supset A$ of P_1^f which is a P_n-complete graph, i.e. $\{y_1, \ldots y_n\} \subset P_1^{f_n} \rightarrow P_n(y_1, \ldots y_n)$.*

2. *If the localization φ^g of φ is not stable, then for all $n < \omega$, P_1^g contains a tuple*

of P_n-inconsistent elements.

In other words, the following are equivalent for any positive base set A:

(i) There exists a localization φ^f of φ (with $\varphi^f = \varphi$ possible) such that φ^f is stable and $P_1^f \supset A$.

(ii) For every $n < \omega$, there exists a localization $P_1^{f_n} \supset A$ which is a P_n-complete graph.

Proof. It suffices to prove the first two statements. (1) is Theorem 2.54. On the other hand, if φ^f has the order property its associated P_1^f contains a diagram in the sense of Definition 2.18. Thus it contains two P_2-inconsistent elements, and so a fortiori an n-tuple of P_n-inconsistent elements, for each n. □

Remark 2.56. (Instability and order) *In order to get the independence property for φ we need an (ω, k)-array for P_∞. The construction of Lemma 2.53 produces an (ω, k) array for P_n whenever a P_n-empty tuple is persistent, but certainly when φ is not low there is no reason to expect analogues of the springboard lemmas. However, if the sequence has finite support then this argument can indeed be used to extract independence from persistence of an empty tuple.*

This suggests that empty tuples are not persistent (under all *localizations) in sequences with finite support in theories without the independence property. In fact this is easy to see: in $(\mathbb{Q}, <)$ for instance, let $\varphi(x; y, z) = y > x > z$ and let the positive base set A be given by some concentric sequence of intervals $\{(a_i, b_i) : i < \kappa\} \subset P_1$.*

Then there is indeed a P_2-empty pair $(c_1, c_2), (d_1, d_2)$ which are each consistent 1-point extensions of A – namely, any pair of disjoint intervals lying in the cut described by the type corresponding to A. Localizing to require consistency with any such pair amounts to giving a definable complete graph containing A, i.e. realizing the type.

One could, of course, consider less powerful notions of persistence, for instance by restricting the form of allowed localizations.

2.5.3 Simplicity

We have seen that the natural first question for persistence, whether there exist persistent empty tuples, characterizes stability: Theorem 2.55. Here we will show that a natural next question, whether there exist persistent infinite empty graphs, characterizes simplicity. Recall that a formula φ is simple if it does not have the tree property.

Notice that we have an immediate proof of this fact by Observation 2.29, which appealed to finite $D(\varphi, k)$-rank for simple formulas to conclude that infinite empty graphs are not persistent. Let us sketch the framework for a different proof by analogy with the previous section. This amounts to deriving Observation 2.29 directly in the characteristic sequence.

Remark 2.57. *In the case of stability, much of the work came in establishing sharpness of the (ω, ℓ)-array. Here, since the persistent configuration is infinite, we have compactness on our side; we may in fact always choose the persistent empty graphs* *to be indiscernible and uniformly k-consistent but $(k+1)$-inconsistent, for some given $k < \omega$. This greatly simplifies matters. The role of the springboard lemma is played by compactness as well; this could have been done in the stable case, but lowness allowed us to give the stronger derivation obtained above. (That argument would remain valid in cases where e.g. $A = \emptyset$ and we restrict the class of allowed localizations to ones which involve instances of some fixed initial segment of the characteristic sequence.)*

Observation 2.58. *Suppose that $(T, \varphi) \mapsto \langle P_n \rangle$. Then the following are equivalent:*

1. *there is a set $T = \{a_\eta : \eta \in 2^{<\omega}\} \subset P_1$ such that, writing $\subseteq$ for initial segment:*
   (a) *For each $\nu \in 2^\omega$, $\{a_\eta : \eta \subset \nu\}$ is a complete P_∞-graph.*
   (b) *For some $k < \omega$, and for all $\rho \in \omega^{<\omega}$, the set $\{a_{\rho^\frown i} : i < \omega\} \subset P_1$ is a P_k-empty graph.*
2. *φ has the k-tree property.*

Proof. This is a direct translation of Definition 0.11. $\square$

Lemma 2.59. *Let X_k be the T_0-configuration describing a strict $(k+1)$-inconsistent sequence, i.e. $V_{X_k} = \omega$ and $E_{X_k} = \{\sigma : \sigma \subset \omega, |\sigma| \leq k\}$. Suppose that for some fixed $k < \omega$ and some formula φ, X_k is persistent in the characteristic sequence $\langle P_n \rangle$ of φ. Then φ is not simple.*

Proof. Let us show that φ has the tree property, around some positive base set A if one is specified. At stage 0, by hypothesis there exists an infinite indiscernible sharply

$(k+1)$-inconsistent sequence $Y_0 \subset P_1$, each of whose elements can be chosen to be a consistent 1-point extension of A in the sense of P_∞ by Fact 2.40. Set a_i to be the ith element of this sequence, for $i < \omega$.

At stage $t+1$, suppose we have constructed a tree of height n, $T_n = \{a_\eta : \eta \in \omega^{\leq n}\}$ such that, writing $\subseteq$ for initial segment:

- every path is a consistent n-point extension of A, i.e. $A \cup \{a_\eta : \eta \subseteq \nu\}$ is a complete P_∞-graph, for each $\nu \in \omega^n$;
- for all $0 \leq k < n$ and all $\eta \in \omega^k$, $\{a_{\eta^\frown i} : i < \omega\}$ is P_k-complete but P_{k+1}-empty.

We would like to extend the tree to level $n+1$, and it suffices to show that the extension of any given node a_ν (for $\nu \in \omega^n$) can be accomplished. But this amounts to repeating the argument for stage 0 in the case where $A = A \cup \{a_\eta : \eta \subseteq \nu\}$. By assumption and Fact 2.40, this remains possible, so we continue.

Notice that the threat of all possible localizations is what makes continuation possible. That is, the schema which says that "x is a 1-point extension of A" simply says that x remains (along with witnesses for X_k) in each of an infinite set of localizations of P_1 with parameters from A. If this schema is inconsistent, there will be a localization contradicting the hypothesis. $\square$

Theorem 2.60. *Let φ be a formula of T and $\langle P_n \rangle$ its characteristic sequence.*

1. *If the localization φ^f of φ is simple, then for each P_∞-graph $A \subset P_1^f$ and for each $n < \omega$, there exists a localization $P_1^{f_n} \supset A$ of P_1^f in which there is a*

uniform finite bound on the size of a P_n-empty graph, i.e. there exists m_n such that $X \subset P_1^f$ and $X^n \cap P_n = \emptyset$ implies $|X| \leq m_n$.

2. *If localization φ^g of φ is not simple, then for all but finitely many $r < \omega$, P_1^g contains an infinite $(r+1)$-empty graph.*

In other words, the following are equivalent for any positive base set A:

(i) There exists a localization φ^f of φ (with $\varphi^f = \varphi$ possible) such that φ^f is simple and $P_1^f \supset A$.

(ii) For each $n < \omega$, there exists a localization $P_1^{f_n} \supset A$ in which there is a uniform finite bound on the size of a P_n-empty graph.

Proof. Once again, it suffices to show the first two statements. (1) is Lemma 2.59 applied to the formula φ^f. Notice that by Theorem 2.55, if φ^f is simple unstable, then there is at least one empty pair. (2) is the second clause of Observation 2.58, where "almost all" means for r above k, the arity of dividing. $\square$

The arguments in this section use persistence to magnify the power of the T_0-configuration, but in a very compatible way. Genericity is inherent in the dynamic arguments, and this restricts our ability to produce genuinely more complex instances of inconsistency (e.g. TP_1) than those already in the T_0-configuration. Thus our project of looking for complex structure in the characteristic sequence $\langle P_n \rangle$ must pass through a more "static" analysis of the T_0-configurations which occur inside some fixed localization. This is the subject of the next chapter.

Chapter 3

Regularity

In this chapter, building on work in Chapter 2, we consider how the distribution and density of edges and of finite configurations in some given localization relate to the the classification-theoretic complexity of φ. The framework of characteristic sequences allows us to bring a deep collection of graph-theoretic structure theorems to bear on our investigations. The organizing principle of the chapter is the question of how arbitrarily large subsets of P_1 can generically interrelate, in the sense of Szemerédi regularity (Theorem L below). That is, we ask which properties of T affect the density δ attained between arbitrarily large ϵ-regular subsets $A, B \subset P_1$ (after localization), where the edge relation is given by P_2.

The picture we obtain is as follows. When φ is stable, by Theorem 2.55, the density (after localization) is always 1. When φ is simple unstable, after localization, there will be an infinite number of missing edges but we can say something strong about their distribution: the density between arbitrarily large ϵ-regular pairs must tend towards 0 or 1 as the graphs grow. In the simple unstable case, a finer function counting the number of edges omitted over finite subgraphs of size n is meaningful, and we give a preliminary description of its possible values in Proposition 3.7. In Section 3.4, we describe the property of having arbitrarily large ϵ-regular subsets of P_1 with edge density bounded away from 0 and 1 in terms of the order property for P_2, Theorem 3.19. In Section 3.5 we relate these issues to complexity in the sense of the Keisler order.

3.1 Preliminaries

Convention 3.1. (Reminders and conventions)

1. *Throughout this chapter $\langle P_n : n < \omega \rangle$ will be the characteristic sequence associated to (T, φ), Definition 2.2. The predicates P_n will always refer to the characteristic sequence, and A will be a positive base set as in Definition 2.1. P_1^f will indicate the localization given by f, i.e. a definable finite restriction of P_1 of a certain fixed form: Definition 2.27.*

2. *Convention 2.3 applies: when we ask whether certain infinite configurations exist in the characteristic sequence, we will always mean whether this is consistent with T.*

3. *For stability, simplicity, and other model-theoretic properties, see §0.4. Recall*

that φ *is* simple *iff* φ *does not have the* tree property*, and* φ *is* stable *iff* φ *does not have the* order property*.* NSOP *means that* T *does not have the strict order property, Definition 0.9.*

4. *We will use the phrase "after localization [X is true]" to mean: "there exists a localization* P_1^f *(consistent with the positive base set* A*, if one has been specified) in which [X is true]", in the sense of Definition 2.27.*

5. ϵ, δ *are real numbers, with* $0 < \epsilon < 1$ *and* $0 \leq \delta \leq 1$*.*

6. *Let* G *be a symmetric binary graph. We present graphs model-theoretically, i.e. as sets of vertices on which certain edge relations hold. Throughout this chapter* $R(x, y)$ *is a binary edge relation, which will sometimes (we will clearly say when) be interpreted as* P_2*.*

7. *A graph is a simple graph: no loops and no multiple edges.* $\forall x(P_1(x) \rightarrow P_2(x, x))$*, but we will, by convention, not count loops when taking* P_2 *as* R*.*

8. *Given a graph* G*:*

   - $|G|$ *is the size of* G*, i.e. the number of vertices.*
   - $e(G)$ *is the number of edges of* G*.*
   - $\hat{e}(G)$ *is the number of edges omitted in* G*.*
   - *An* empty graph *is a graph with no edges.*

- *A* complete graph *is a graph with all edges, i.e. in which* $x, y \in G, x \neq y \implies R(x, y)$.
- *The* degree *of a vertex is the number of edges which contain it.*
- *The* dual graph G' *has the same vertices and inverted edges, i.e. for* $x \neq y$, $G' \models R(x, y) \iff G \models \neg R(x, y)$.

9. Write (X, Y) *to indicate a a bipartite graph. Then:*

- $e(X, Y)$ *is the number of edges between elements* $x \in X$ *and* $y \in Y$. *Note that if* $G = A \cup B$ *then possibly* $e(G) \neq e(A, B)$, *as the latter counts only edges between* A *and* B.
- $\hat{e}(X, Y)$ *is the number of edges omitted between elements* $x \in X$ *and* $y \in Y$.
- *The* density *of a finite bipartite graph* (X, Y) *is* $\delta(X, Y) := e(X, Y)/|X||Y|$ *when* $|X|, |Y| \neq 0$, *and* 0 *otherwise.*
- *An* infinite empty pair *is* (X, Y) *such that* $|X| = |Y| \geq \aleph_0$ *and for all* $x \in X$, $y \in Y$, *we have* $\neg R(x, y)$.
- *A* complete bipartite graph *is* (X, Y) *such that for all* $x \in X, y \in Y$, $R(x, y)$.
- *The dual* $(X, Y)'$ *of a bipartite graph inverts precisely the edges between the components* X *and* Y.

Note that many of the results in this chapter could be made stronger by replacing the assumption that φ is simple, φ is stable, etc. with "some localization φ^f is simple, stable..." in the sense of Definition 2.27.

3.2 Counting functions on simple φ

Throughout this section, we consider the binary edge relation P_2 from the characteristic sequence of φ.

Observation 3.2. *Suppose φ is stable and let A be a positive base set (possibly empty). Then after localization, for any two disjoint finite $X, Y \subset P_1$, $\delta(X, Y) = 1$.*

Proof. Theorem 2.55 says that when φ is stable, then after localization P_1 is a complete graph, so a fortiori there are no edges omitted between disjoint components. □

Definition 3.3. *Define $\alpha : \omega \to \omega$ to be*

$$\max \quad \{\hat{e}(X) : X \subset P_1, |X| = n\}$$

i.e. the largest number of P_2-edges omitted over an n-size subset of P_1.

Observation 3.4. *Suppose φ does not have the tree property. Then after localization $\alpha(n) < \frac{n(n-1)}{2}$.*

Proof. The maximum possible value $\frac{n(n-1)}{2}$ of any $\alpha(n)$ is attained on a P_2-empty graph, on which $x \neq y \implies \neg P_2(x, y)$. Apply Theorem 2.60 which says that when φ does not have the tree property then we have, after localization, a uniform finite bound k on the size of a P_2-empty graph $X \subset P_1$. So the function α is eventually strictly below the maximum. $\square$

These two Observations show that the function $\alpha(n)$ is meaningful, i.e.

$$\frac{n(n-1)}{2} > \quad \alpha(n) \quad > 0$$

precisely when φ is simple unstable. With some care we can easily restrict the range further. A famous theorem of Turán says that:

Theorem K. (Turán, [15]:Theorem 2.2) *If G_n is a graph with n vertices and*

$$e(G) > \left(1 - \frac{1}{k-1}\right)\frac{n^2}{2}$$

then G_n contains a complete subgraph on k vertices.

Recall also from Chapter 2, Claim 2.21 that if φ is simple unstable, φ has the independence property and so P_2 contains an $(\omega, 2)$-array, Definition 2.18.

Observation 3.5. *Suppose that P_2 contains an $(\omega, 2)$-array. Then $\alpha(n) \geq \lfloor \frac{n}{2} \rfloor$.*

Corollary 3.6. *When φ is simple unstable, then after localization*

$$\left(1 - \frac{1}{k-1}\right)\frac{n^2}{2} \geq \alpha(n) \geq \left\lfloor \frac{n}{2} \right\rfloor$$

Proof. The righthand side is Observation 3.5. For the lefthand side, let $k > 1$ be the uniform finite bound on the size of an empty graph, given by simplicity (see the proof of Observation 3.4), and apply Turán's theorem to the dual graph. $\square$

At the end of Section 3.3 we will give a proof of the following:

Proposition 3.7. *When φ is simple unstable either*

$$\left(1 - \frac{1}{1-k}\right)\frac{n^2}{2} \ \geq \ \alpha(n) \ \geq \ \frac{n^2}{4} \qquad \text{or} \qquad \mathcal{O}(n^2) > \alpha(n) \geq \left\lfloor \frac{n}{2} \right\rfloor$$

The proof will follow from Proposition 3.14 below, which will show more, namely that for φ simple unstable, either $\mathcal{O}(n^2) > \alpha(n)$ or there exists an infinite empty pair in P_1.

Our strategy is going to be to show that in the absence of such an "empty pair" we can partition sufficiently large graphs into many pieces of roughly equal size so that, asymptotically, almost no edges are omitted between pairs. Thus the bulk of the omitted edges must occur inside the (relatively much smaller) components. The main tool will be Theorem L below.

3.3 Szemerédi regularity

We begin with a review of Szemerédi's celebrated *regularity lemma*. Throughout this section ϵ, δ are real numbers, $0 < \epsilon < 1$ and $0 \leq \delta \leq 1$.

Definition 3.8. *[29], [15] The finite bipartite graph (X, Y) is ϵ-regular if for every $X' \subset X$, $Y' \subset Y$ with $|X'| \geq \epsilon|X|$, $|Y'| \geq \epsilon|Y|$, we have: $|\delta(X,Y) - \delta(X',Y')| < \epsilon$.*

The regularity lemma says that sufficiently large graphs can always be partitioned into k pieces X_i of approximately equal size so that almost all of the pairs (X_i, X_j) are ϵ-regular.

Theorem L. (Szemerédi [15], [29]) *For every ϵ, m_0 there exist $N = N(\epsilon, m_0)$, $m = m(\epsilon, m_0)$ such that for any graph X, $N \leq |X| < \aleph_0$, for some $m_0 \leq k \leq m$ there exists a partition $X = X_1 \cup \cdots \cup X_k$ satisfying:*

- $||X_i| - |X_j|| \leq 1$ *for* $i, j \leq k$
- *All but at most ϵk^2 of the pairs (X_i, X_j) are ϵ-regular.*

One important consequence is that we may, approximately, describe large graphs G as random graphs where the edge probability between x_i and x_j is the density $d_{i,j}$ between components X_i, X_j in some Szemerédi-regular decomposition. We include here two formulations of this idea from the literature, the first for intuition and the second for our applications.

Theorem M. (from Gowers [11]) *For every $\alpha > 0$ and every k there exists $\epsilon > 0$ with the following property. Let $V_1, \ldots V_k$ be sets of vertices in a graph G, and suppose that for each pair (i, j) the pair (V_i, V_j) is ϵ-regular with density δ_{ij}. Let H be a graph with vertex set $(x_1, \ldots x_k)$ and let $v_i \in V_i$ be chosen uniformly at random, the choices being independent. Then the probability that $v_i v_j$ is an edge of G iff $x_i x_j$ is an edge of H differs from $\Pi_{x_i x_j \in H} \delta_{ij} \Pi_{x_i x_j \notin H} (1 - \delta_{ij})$ by at most α.*

The formulation we will use, Theorem N, requires a preliminary definition.

Definition 3.9. *[15]* (The reduced graph)

1. *Let $G = X_1, \ldots X_k$ be a partition of the vertex set of G into disjoint pieces. Given parameters ϵ, δ, define the* reduced graph $R(G, \epsilon, \delta)$ *to be the graph with*

vertices x_i $(1 \leq i \leq k)$ *and an edge between* x_i, x_j *just in case the pair* (X_i, X_j) *is* ϵ*-regular of density* $\geq \delta$.

2. *Write* $R(t)$ *for a full graph of height* t *whose reduced graph is* R, *i.e.,* $R(t)$ *consists of* k *clusters* $X_1, \ldots X_k$, *each with* t *vertices, such that* $e(X_i) = 0$, *and* $\delta(X_i, X_j) = 1$ *iff there is an edge between* x_i *and* x_j *in* R.

The following lemma (called the "Key Lemma" in [15]) says that sufficiently small subgraphs of the reduced graph must actually occur in the original graph G.

Theorem N. (Key Lemma, [15]:Theorem 2.1) *Given* $\delta > \epsilon > 0$, *a graph* R, *and a positive integer* m, *let* G *be any graph whose reduced graph is* R, *and let* H *be a subgraph of* $R(t)$ *with* h *vertices and maximum degree* $\Delta > 0$. *Set* $d = \delta - \epsilon$ *and* $\epsilon_0 = d^{\Delta}/(2+\Delta)$. *Then if* $\epsilon \leq \epsilon_0$ *and* $t - 1 \leq \epsilon_0 m$, *then* $H \subset G$. *Moreover the number of copies of* H *in* G *is at least* $(\epsilon_0 m)^h$.

Remark 3.10. *In the statement of the Key Lemma, "*$H \subset G$*" means that there is a bijection* $f : H \to X \subset G$ *such that* $e(h_1, h_2)$ *implies* $e(f(h_1), f(h_2))$. *With some slight modifications (recording whether a missing edge in the reduced graph means the density is near 0 or the pair is not regular; and using the dual graphs when necessary) we may assume "*$H \subset G$*" has the usual meaning of isomorphic embedding, but this will not be an issue for the arguments in this section.*

Definition 3.11. *Let* G *be a graph and let* $G = X_1 \cup \cdots \cup X_n$ *be a decomposition into disjoint pieces, e.g. as given by Theorem L. Call the edges between vertices in* *different components* interstitial edges.

An easy application of the Key Lemma shows that

Observation 3.12. *Suppose that there exists δ, $0 < \delta < 1$ such that for all $0 < \epsilon < 1$ and all $N \in \mathbb{N}$ there exist disjoint subsets $X_N, Y_N \subset P_1$, $|X_N| = |Y_N| \geq N$ such that (X_N, Y_N) is ϵ-regular with density δ. Then P_1 contains an infinite empty pair.*

Proof. Apply the Key Lemma to each dual graph $(X_N, Y_N)'$, which is still regular, of density bounded away from 0 and 1. For each $t < \omega$, for all N sufficiently large, $(X_N, Y_N)'$ contains a complete bipartite graph on t vertices, as this occurs as a subset of $R(t)$. □

Corollary 3.13. *We can explictly reword this as:*

1. *If P_1 does not contain an infinite empty pair then we can define a function $f : (0,1) \times \omega \to (0,1)$ monotonic increasing as $\epsilon \to 0$ and $N \to \infty$ such that if (X, Y) is an ϵ-regular pair of size N then its density is at least $\delta := f(\epsilon, N)$.*

2. *Likewise we can define $g : (0,1) \times \omega \to (0,1)$ monotonic increasing as $\epsilon \to 0$ and $k \to \infty$ such that if X is a graph large enough to admit an ϵ-regular decomposition into k-many pieces then the density between any two regular components is at least $\delta := g(\epsilon, k)$.*

We are now prepared to prove:

Proposition 3.14. *When φ is simple unstable, if there does not exist an infinite empty pair $X, Y \subset P_1$, then $\alpha(n) < \mathcal{O}(n^2)$.*

Proof. By the analysis above, if there is no infinite empty pair then the density between ϵ-regular pairs of size n must tend to 1 as $n \to \infty$, $\epsilon \to 0$. Let us count the interstitial edges as we partition graphs X into ℓ pieces, most of which are pairwise ϵ-regular, by applying the Regularity Lemma (as ℓ grows). Recall that the goal is to show that for any positive constant c, eventually the count $\alpha(n)$ falls below cn^2.

Suppose that X, $|X| = n$ is large enough to admit an ϵ-regular decomposition into ℓ-many pieces of size m. Then in the notation of Corollary 3.13, setting $\delta := g(\epsilon, \ell)$, the contribution of the interstitial edges is at most:

$$\epsilon \ell^2 m^2 + (1 - \epsilon)\ell^2(1 - \delta)m^2$$

as the term on the left assumes the irregular pairs are missing all possible edges, and the term on the right counts the expected number missing from the regular pairs. But $m = |X|/\ell$, so given some such regular decomposition the count is simply

$$\epsilon n^2 + (1 - \epsilon)(1 - \delta)n^2$$

As $n \to \infty$, it continually passes the threshold for ϵ-regular decompositions into at least k pieces. Corollary 3.13 ensures that for any $\delta < 1$, as $\epsilon \to 0$, for cofinally many k there is a threshold size $N_{\epsilon,k}$ such that any in graph X of size $n > N_{\epsilon,k}$ there is a Szemerédi decomposition into k pieces such that the pairwise regular components have density at least δ. Thus as $n \to \infty$, $\epsilon \to 0$ and $\delta \to 1$. So there will eventually be less than cn^2 interstitial edges added to α, for any $c > 0$.

It remains to consider edges omitted entirely within components. The idea is that

$$\alpha(m\ell) \leq \ell\frac{m^2}{2}$$

which falls below the order of $(m\ell)^2$ as m stays fixed and $\ell \to \infty$. Let us be more careful, however, as these quantities are not entirely independent. Notice that the first part of the argument shows that for any constant c, for n greater than some N_c, there exist ϵ, k such that all graphs of size n admit a Szemerédi decomposition into k components, almost all of which are ϵ-regular and such that no more than cn^2 of the interstitial edges are missing. When $n/k > N_c$, each of the components in the decomposition has this property in turn. For any finite d, by choosing n sufficiently large, we can ensure that this decomposition continues to a depth of d steps. At each stage in the decomposition, the potentially large number of missing edges disappears inside the relatively much smaller components. As we can do this for any ϵ, for cofinally many k depending on ϵ, and for any d, we see that the total edge count must also fall below $c'n^2$ for any $c' > 0$. $\square$

Proof. (of Proposition 3.7) This is now an immediate corollary of Proposition 3.14, $\frac{n^2}{4}$ being the number of edges omitted in an empty pair. $\square$

Remark 3.15. *Proposition 3.14, and thus Proposition 3.7, are more natural than might appear. We know that when φ is simple unstable P_2 will contain a tuple of inconsistent elements; in fact, even the random graph, the "least complex" example of a simple theory, contains an infinite empty pair, for instance* $(\{(a, x) : x \in M, x \neq$

$a\}, \{(y, a) : y \in M, y \neq a\})$ when $\varphi = xRy \wedge \neg xRz$. Also, as Szemerédi regularity deals with density, it cannot (in this formulation) give precise information about edge counts below $\mathcal{O}(n^2)$. On the other hand, notice that the analysis of types in the random graph, Example 2.7, reduces to a definable $(\omega, 2)$-array, on which α is essentially linear.

3.4 Order and genericity

The main idea of the previous proof was that in the course of a regular decomposition, if a growing sequence of ϵ-regular pairs X_n, Y_n failed to have density sufficiently close to 1, we could in due course extract an infinite pair X', Y' with density 0. This was because, by the Blow-up Lemma, if the density of some increasing sequence of ϵ-regular pairs stayed bounded away from 1, we could obtain increasingly large empty graphs as subgraphs.

Examining this assumption more carefully, we can give an "excluded-middle" characterization for the associated P_2 in a class of theories which strictly includes simple theories: in some localization, the density of any sufficiently large ϵ-regular pair $X, Y \subset P_1$ must approach either 0 or 1 (as a function of ϵ and N), Corollary 3.20 below.

The next few results are given for any symmetric binary relation R.

Lemma 3.16. *Suppose that for some $0 < \delta < 1$ and for all ϵ, n with $0 < \epsilon < 1$, $n \in \mathbb{N}$*

we have a bipartite R-graph (X, Y), $|X| = |Y| \geq n$, *such that* (X, Y) *is* ϵ*-regular with density* d, *where* $|d - \delta| < \epsilon$. *Then* R *has the order property (Definition 0.9).*

Proof. In order to apply the reformulated Key Lemma, it suffices to show that for sufficiently small ϵ_0 and sufficiently large k_0 there is a Szemerédi-regular decomposition of X and of Y into k_0 pieces such that all but $k_0(\epsilon_0)^2$ of the pairs X_i, Y_i are ϵ_0-regular with density near δ.

Given ϵ_0, k, let k_0, N_0 be the number of components and threshold size, respectively, given by the regularity lemma. Choose ϵ so that $\frac{1}{k_0} > \epsilon$ and $n > N_0$. Let (X, Y) be the ϵ-regular pair of size at least n and density near δ, given by hypothesis.

By regularity, $n > N_0$ means that there is a decomposition $X = \cup_{i \leq k_0} X_i$, $Y = \cup_{i \leq k_0} Y_i$ into disjoint pieces of near equal size and that all but $\epsilon_0(k_0)^2$ of the pairs (X_i, Y_j) are ϵ_0-regular. However any one of these regular pairs (X_i, Y_j) will satisfy $|X_i|, |Y_j| = n/k_0 > \epsilon n$, so $|d(X_i, Y_j) - d(X, Y)| = |d(X_i, Y_j) - \delta \pm \epsilon| < \epsilon$ and $|d(X_i, Y_j) - \delta| < 2\epsilon$, as desired. □

This is a version of the graph-theoretic "Slicing Lemma":

Fact 3.17. (Slicing Lemma, [15]:Fact 1.5) *Let* (A, B) *be an* ϵ*-regular pair with density* d, *and, for some* $\alpha > \epsilon$, *let* $A' \subset A$, $|A'| \geq \alpha|A|$, $B' \subset B$, $|B'| \geq \alpha|B|$. *Then* (A', B') *is an* ϵ'*-regular pair with* $\epsilon' = max\{\epsilon/\alpha, 2\epsilon\}$, *and for its density* d' *we have* $|d' - d| < \epsilon$.

The remaining ingredient is a lemma from Chapter 4 which inverts a construction of Shelah:

[**Lemma 4.5**, quoted for completeness.] *Suppose the formula $R(x;y)$ has the order property. If T does not have the strict order property, then on some infinite (A,B), R is a bipartite random graph.*

Definition 3.18. *Fix a binary edge relation R. Call a density $0 \leq \delta \leq 1$* attainable *if for all ϵ there exists a sequence $\langle S^\delta_\epsilon = \langle (X_i, Y_i) : i < \omega \rangle$ of finite bipartite R-graphs such that for all $n < \omega, \epsilon > 0$ there is $N < \omega$ such that for all $i > N$,*

- $|X_i| = |Y_i| \geq n$,
- (X_i, Y_i) *is ϵ-regular with density d_i, where $|d_i - \delta| < \epsilon$.*

Theorem 3.19. *(NSOP) The following are equivalent for a binary relation $R(x,y)$:*

1. *For some $0 < \delta < 1$ and for all N, ϵ there exist disjoint X, Y with $|X| = |Y| \geq N$ such that (X,Y) is ϵ-regular with density d, $|d - \delta| < \epsilon$.*
2. *For any attainable $0 < \delta < 1$ such that for all N, ϵ there exist disjoint X, Y with $|X| = |Y| \geq N$ such that (X,Y) is ϵ-regular with density d, $|d - \delta| < \epsilon$.*
3. *R has the order property.*

Proof. (2) $\rightarrow$ (1) Attainable densities exist, e.g. $\frac{1}{2}$: consider subgraphs of an infinite random bipartite graph.

(1) $\rightarrow$ (3) Lemma 3.16.

(3) $\rightarrow$ (1) Chapter 4, Lemma 4.5, which says that from (3), assuming $NSOP$, we can construct an infinite random bipartite graph with edge relation R. $\square$

Corollary 3.20. *Suppose that $(T, \varphi) \mapsto \langle P_n \rangle$, and T is simple. Then after localization, the density of any sufficiently large P_2-regular pair (X, Y) must approach either 0 or 1. More precisely, for each such P_2 in some finite localization there exists $f : \mathbb{N} \times (0,1) \to [0, \frac{1}{2}]$ monotonic increasing as $n \to \infty$, $\epsilon \to 0$ such that if $X, Y \subset P_1$, $|X|, |Y| \geq n$ and (X, Y) is ϵ-regular, then either $d(X, Y) < f(n, \epsilon)$ or $d(X, Y) > 1 - f(n, \epsilon)$.*

Proof. The analysis of Chapter 2 showed that when T is simple, after localization the formulas P_n are stable (Conclusion 2.33), so in particular P_2 cannot have the order property. $\square$

Remark 3.21. *This is a class strictly containing the simple theories. For instance, there is a theory with NTP_1 page 61 whose P_2, after localization, contains an (ω, ω) array but does not contain the order property. In the next section, we consider the force of this dividing line.*

3.5 Two kinds of order property

The previous section gave some insight into the idea, which arose in Chapter 2, that having the order property in P_2 is a signal of complexity for φ (Observation 2.30). That proof simply showed that the order property in P_2 implies dividability of φ_2. Theorem 3.19 said much more, namely that in the absence of strict order, P_2 has the order property iff it is, in some sense, universal for finite bipartite graphs.

"In some sense" refers to the issue of edges within the components X, Y of a bipartite graph. This section considers two polar opposite order properties and their implications in P_2. In the language of Theorem 3.19, the first case arises when P_2 has nontrivial density between some pair of positive base sets for types; the second, when P_2 has nontrivial density between some pair of dividing sequences.

Definition 3.22. (Two kinds of order property) *Let $\langle P_n \rangle$ be the characteristic sequence of φ.*

1. *φ has the n-*compatible order property*, for some $n < \omega$ (or $n = \infty$) if there exist $\langle a_i, b_i : i < \omega \rangle$ such that for all $m \leq n$ (or $m < \omega$), $P_m((a_{i_1}, b_{j_1}), \ldots (a_{i_m}, b_{j_m}))$ iff* $\max\{i_1, \ldots i_m\} < \min\{j_1, \ldots j_m\}$.

1.′ *When the sequence has support 2 this becomes: there exist $\langle a_i, b_i : i < \omega \rangle$ such that $P_2(a_i, a_j)$, $P_2(b_i, b_j)$ for all i, j and $P_2(a_i, b_j)$ iff $i < j$.*

2. *φ has the n-*empty order property*, for some $n \in \omega$, if:*

   there exist $\langle a_i, b_i : i < \omega \rangle$ such that (i) $P_2(a_i; b_j)$ iff $i < j$ and (ii) $\neg P_n(a_{i_1}, \ldots a_{i_n})$, $\neg P_n(b_{i_1}, \ldots b_{i_n})$ hold for all $i_1, \ldots i_n < \omega$.

Let us briefly justify excluding a natural third possibility, the "semi-compatible order property:"

Observation 3.23. *There is a formula in the random graph which has the semi-compatible order property.*

Proof. Let us apply the construction of Example 2.2, page 59. Choose two distinguished elements $0, 1$ (this can be coded without parameters; see the example). Define $\psi(x; y, z)$ to be $x = y$ if $z = 0$, xRy otherwise. Then on any sequence of distinct elements $\langle a_i b_i : i < \omega \rangle \subset M$ which witness the order property ($a_i R b_j \iff i < j$), we have additionally that

$$\exists x \left(\psi(x; a_i, 0) \land \psi(x; b_j, 1)\right) \iff \exists x \left(x = a_i \land xRb_j\right) \iff i < j$$

so P_2 has the order property on the sequence $\langle (a_i, 0), (b_i, 1) : i < \omega \rangle$. On the other hand, $\exists x(x = a_i \land x = a_j) \iff i = j$, so the row of elements $(a_i, 0)$ is a P_2-empty graph. Finally, $\exists x(xRb_i \land xRb_j)$ always, by the axioms of the random graph; so the row of elements $(b_j, 1)$ is a P_∞-complete graph. $\square$

Remark 3.24. *Assuming $MA + 2^{\aleph_0} > \aleph_1$, Shelah has constructed an ultrafilter on ω which saturates certain models of the random graph, but not of theories with the tree property, thus a fortiori not maximal theories (Theorem H, page 38). This fact, coming after the analysis in this section, is a strong argument for the "semi-compatible order property" being less complex: it cannot, by itself, imply maximality.*

We return to the study of the compatible and empty order properties.

Convention 3.25. *When more than one characteristic sequence is being discussed, write $P_n(\varphi)$ to indicate the nth hypergraph associated to the formula φ. Recall that $\varphi_\ell(x; y_1, \ldots y_\ell) := \bigwedge_{i \leq \ell} \varphi(x; y_i)$.*

The following general principle will be useful.

Claim 3.26. *Suppose that we have a sequence $\langle c_i : i \in \mathbb{Z} \rangle$ and a formula $\rho(x; y, z)$ such that:*

- $\exists x \rho(x; c_i, c_j) \iff i < j$
- $\exists x \left(\bigwedge_{\ell \leq n} \rho(x; c_{i_\ell}, c_{j_\ell}) \right)$ *just in case* $\max\{i_1, \ldots i_n\} < \min\{j_1, \ldots j_n\}$

Then ρ has the ∞-compatible order property.

Proof. This is a fairly straightforward translation of the definition, but let us give a picture. Essentially, we can now describe intervals. For instance, setting $b_k = (c_{-2i}, c_{2i})$ and $a_k = (c_{2i-1}, c_{3i})$ for $1 \leq k \leq n$, $j = n\text{-}k$ is sufficient. In the following picture, matching parentheses are pairs:

$$\leftarrow [-[-[-[-(-]-(-]-(-]-(-]-)-)-)-)\rightarrow$$

□

Observation 3.27. *Suppose that φ has the strict order property, i.e. there is an infinite sequence $\langle c_i : i < \omega \rangle$ on which $\exists x(\neg\varphi(x; c_i) \wedge \varphi(x; c_j)) \iff i < j$. Then $\neg\varphi(x; y) \wedge \varphi(x; z)$ has the ∞-compatible order property.*

Proof. Writing $\rho(x; y, z) = \neg\varphi(x; y) \wedge \varphi(x; z)$,

- $\exists x \rho(x; c_i, c_j) \iff i < j$, by definition of strict order;
- $\exists x(\rho(x; c_i, c_j) \wedge \rho(x; c_k, c_\ell)) \iff i, k < j, \ell$

and the characteristic sequence $P_\infty(\rho)$ has support 2. Apply Claim 3.26. □

Recall from Chapter 2, Definition 2.5 that T_0-configurations are the P_n- (here, P_2-) graphs which can consistently occur in some characteristic sequence.

Observation 3.28. *(NSOP) Let $\langle P_n : n < \omega \rangle$ be the characteristic sequence of φ and $\langle P'_n : n < \omega \rangle$ that of φ_2. Suppose that φ has the compatible order property witnessed by $X \subset P_1$, and that the sequence $\langle P_n \rangle$ restricted to X has support 2. Then P'_2 is universal for finite T_0-configurations. Moreover, there is an infinite subset $C' \subset P'_1$ over which P_2 is a random graph, and the sequence P'_n restricted to C' has support 2.*

Proof. By Corollary 4.6 to Lemma 4.5, which assumes NSOP, since P_2 has the order property between P_∞-complete graphs we may find two disjoint infinite sets A, B which are P_∞-complete graphs so that (A, B) is an infinite bipartite random graph for P_2 and inherits support 2. Let $C = A \cup B$. Let us show by induction on ℓ that the graph with vertices in P'_1 and edge relation P'_2 is universal for finite symmetric binary graphs on ℓ vertices. For $\ell = 1$ it is true. Suppose $\ell = m + 1$, and we would like to embed a graph given by $\{x_1, \dots x_{m+1}\}$ and edge relation R. Suppose that, by inductive hypothesis, we have found $\{(a_i, b_i) : i \leq m\} \subset P'_1$ such that:

- $a_i \in A$, $b_i \in B$ for $i \leq m$
- $P_2(a_i, a_j)$ and $P_2(b_i, b_j)$ for $1 \leq i, j \leq m$, and therefore $P'_1((a_i, a_j))$, $P'_1((b_i, b_j))$
- $x_i \mapsto (a_i, b_i)$ $(i \leq m)$ is a graph isomorphism,

  i.e. if $i \neq j$ then $R(x_i, x_j)$ iff $P'_2((a_i, b_i), (a_j, b_j))$.

Now $R(x_{m+1}, x_i) \iff i \in \sigma$ for some $\sigma \subset m$. By the choice of A, B it is easy to find $a_{m+1} \in A$ such that $P_2(a_{m+1}, b_j) \iff j \in \sigma$, and b_{m+1} such that $P_2(b_{m+1}, a_j)$ for all $1 \leq j \leq m$. Then P_2' holds on the pair $(a_{m+1}, b_{m+1}), (a_j, b_j)$ iff

$$\exists x\, (\varphi_2(x; a_{m+1}, b_{m+1}) \wedge \varphi_2(x; a_j, b_j)) \iff P_4(a_{m+1}, b_{m+1}, a_j, b_j) \iff P_2(a_{m+1}, b_j)$$

as we built all other pairs to be consistent, and assumed that on C the characteristic sequence of φ depends on 2. As we have built C' directly from C, it is straightforward to check that in fact the entire sequence $\langle P_n' \rangle$, restricted to C', inherits support 2. $\square$

Remark 3.29. *Thus the compatible order property pushes forward to a random graph, whereas the incompatible order property would simply push forward to a $P_2(\varphi_2)$-empty graph.*

Example 3.1. *Let T be the theory of the triangle-free random graph with edge relation R. Consider $\varphi(x; y, z) = xRy \wedge xRz$. (The negative instances could be added but are not necessary.) Then:*

- $P_1((y, z)) \iff \neg yRz$.
- $P_2((y, z), (y', z'))$ *iff* $\{y, y', z, z'\}$ *is an empty graph.*
- *The sequence has support* 2*, as the only problems come from a single new edge:* $P_n((y_1, z_1), \ldots (y_n, z_n))$ *iff*

$$\exists x \left(\bigwedge_{i \leq n} xRy_i \wedge \bigwedge_{j \leq n} xRz_j \right) \quad \textit{that is, if } \bigcup_i y_i \cup \bigcup_j z_j \textit{ is a } P_2\textit{-empty graph.}$$

Let $\langle a_i, b_i : i < \omega\rangle$ be a sequence witnessing the incompatible *order property with respect to the edge relation R, say $a_i R b_j$ iff $j \leq i$. Then $\exists x(xRa_i \wedge xRb_j)$ iff $i < j$, i.e. $(a_i, b_j) \in P_1$ iff $i < j$. Also, $\exists x(xRa_i \wedge xRb_j \wedge xRa_k \wedge xRb_\ell)$ if, in addition, $i, k < j, \ell$. Apply Claim 3.26.*

SOP_3 will be discussed in Chapter 4, §4.3. Recall from Chapter 1 that SOP_3 implies maximality in the Keisler order.

Lemma 3.30. *Suppose that $\theta(x; y)$ has SOP_3, so $\ell(x) = \ell(y)$. Let $\varphi_r = \varphi$, $\psi_\ell = \psi$ be the formulas from Definition 4.8. Then $\rho(x; y, z) := \varphi_r(y, x) \wedge \psi_\ell(x, z)$ has the ∞-compatible order property on some $A' \subset P_1$. Moreover, we can choose A' so that the sequence restricted to A' has support 2.*

Remark 3.31. *This is an existential assertion, and it is straightforward to check that it remains true if we modify ρ to include the corresponding negative instances.*

Proof. (of Lemma) Let $A := \langle a_i : i < \mathbb{Q}\rangle$ be an infinite indiscernible sequence from Definition 4.8. Then

$$P_1((a_i, a_j)) \iff \exists x\, (\varphi_r(a_i, x) \wedge \psi_\ell(x, a_j)) \iff i < j$$

by the choice of φ, ψ. More generally,

$$P_n((a_{i_1}, a_{j_1}), \ldots (a_{i_n}, a_{j_n})) \iff \exists x \left(\bigwedge_{t \leq n} \varphi_r(x; a_{i_t}) \wedge \bigwedge_{t \leq n} \psi_\ell(x; a_{j_t}) \right)$$

which, again applying Definition 4.8, happens iff $\max\{i_1, \ldots i_n\} < \min\{j_1, \ldots j_n\}$, a condition which has support 2. We now apply Claim 3.26 to obtain $A' \subset A \times A$

witnessing the compatible order property. Note that while $\langle P_n \rangle$ need not depend on 2 elsewhere in P_1 (we know very little about ρ off A), it does depend on 2 on elements from the sequence A'. $\square$

Assuming $NSOP$, we thus obtain a simpler proof of Shelah's theorem that any theory with SOP_3 is maximal in the Keisler order ([27], [28]):

Theorem 3.32. *Suppose T has SOP_3 and not strict order. Then there exist a formula ρ of T with characteristic sequence $\langle P_n \rangle$ and an infinite set $C \subset P_1$ such that:*

- *P_2 is a random graph on C.*
- *The sequence $\langle P_n \rangle$ restricted to $C \subset P_1$ depends on 2.*

Thus T is maximal in the Keisler order.

Proof. Let ρ' be the formula defined in Lemma 3.30, and A' the compatible order property sequence from that proof. Apply Observation 3.28 to ρ' obtain the random graph with support 2 in the characteristic sequence of $\rho := \rho'_2$. In order to show that these conditions imply maximality by Lemma 5.14, recall that there exists (T, φ) whose sequence depends on 2 and which is maximal in the Keisler order: namely, $\varphi(x; y, z) = y > x > z$ in any theory of strict linear order. $\square$

Conclusion 3.33. *For any theory T without the strict order property, the following are equivalent. Each clause mentions a formula φ and its associated characteristic sequence $\langle P_n \rangle$. In the proof, we indicate by subscripting how the formulas relate.*

1. *For some formula φ of T, φ has the ∞-compatible order property (Definition 3.22).*

2. *For some formula φ of T, there exists an infinite $C \subset P_1(\varphi)$ witnessing the compatible order property such that $P_\infty(\varphi)$ restricted to C has support 2.*

3. *For some formula φ of T, there exist infinite positive base sets $A, B \subset P_1(\varphi)$ such that $P_\infty(\varphi)$ restricted to $A \cup B$ has support 2 and (A, B) is a bipartite random graph with edge relation $P_2(\varphi)$.*

4. *For some formula φ of T, there exists an infinite $C' \subset P_1(\varphi)$ on which the characteristic sequence $\langle P_n(\varphi) \rangle$ has support 2, and on which $P_2(\varphi)$ is a random graph.*

5. *For some formula φ of T, in $P_1(\varphi)$ there are arbitrarily large ϵ-regular pairs of positive base sets whose interstitial density stays bounded away from 0 and 1, where the edge relation is taken to be $P_2(\varphi)$. That is, there exist $\delta, 0 < \delta < 1$ such that for every $n < \omega$ and $0 < \epsilon$, there are positive base sets $A, B \subset P_1(\varphi)$, $n \leq |A| = |B| < \aleph_0$ so that (A, B) is ϵ-regular in the sense of Szemerédi, with density d such that $|d - \delta| < \epsilon$.*

6. *(Translating (5)) For some formula φ of T, there exists $0 < \delta < 1$ such that for every $n < \omega, 0 < \epsilon$ there exist parameter sets $A, B \subset P_1(\varphi)$, $n < |A| = |B| < \aleph_0$ such that*

- $\{\varphi(x;a) : a \in A\}$, $\{\varphi(x;b) : b \in B\}$ *are both consistent partial* φ*-types*

- *the likelihood that instances chosen from each are consistent is roughly* δ*,*

$$\left| \frac{|\{(a,b) : a \in A, b \in B, \{(\varphi(x;a), \varphi(x;b)\} \text{ consistent}\}|}{|A||B|} - \delta \right| \leq \epsilon$$

- *and in fact for every* $A' \subset A, B' \subset B, |A'| \geq \epsilon|A|, |B'| \geq \epsilon|B|$*,*

$$\left| \frac{|\{(a,b) : a \in A', b \in B', \{(\varphi(x;a), \varphi(x;b)\} \text{ consistent}\}|}{|A'||B'|} - \delta \right| < \epsilon$$

Proof. The only time the formula changes is from (3) to (4), when φ becomes φ_2.

$(1)_\varphi \implies (2)_\varphi$ by compactness.

$(2)_\varphi \implies (3)_\varphi$ is Corollary 4.6 to Lemma 4.5.

$(3)_\varphi \implies (4)_{\varphi_2}$ is Observation 3.28.

$(3)_\varphi \implies (5)_\varphi$ is clear.

$(1)_\varphi \iff (5)_\varphi$ is Theorem 3.19.

$(5)_\varphi \iff (6)_\varphi$ is a direct translation.

$(4)_\varphi \implies (1)_\varphi$ as this random graph will be universal for finite T_0-configurations.

$\square$

Chapter 4

Depth of independence

In this chapter, we observe and explicate a discrepancy between the model-theoretic notion of an infinite random k-partite graph and the finitary version given by Szemerédi regularity, showing essentially that a class of infinitary k-partite random graphs which do not admit reasonable finite approximations must have the strong order property SOP_3 (a model-theoretic notion of rigidity, Definition 4.7 below). This is structurally interesting, but also suggestive because SOP_3 is known to imply maximality in the Keisler order.

4.1 A seeming paradox

The formal definitions of "bipartite random graph," etc. will be given in the next section, but there are no real surprises.

Observation 4.1. *Let T be the theory of the triangle-free random graph, with edge relation R. Then it is consistent with T that there exist disjoint infinite sets X, Y, Z such that each pair $(X, Y), (Y, Z), (X, Z)$ is a bipartite random graph.*

Proof. The construction has countably many stages. At stage 0, let $X_0 = \{a\}, Y_0 = \{b\}, Z_0 = \{c\}$ where a, b, c have no R-edges between them. At stage $i + 1$, let X_{i+1} be X_i along with $2^{|Y_i|+|Z_i|}$-many new elements:

1. for each subset $\tau \subset Y_i$, a new element x_τ such that for $y \in Y$, $x_\tau R y \iff y \in \tau$, however $\neg x_\tau R x$ for any x previously added to X_{i+1}.

2. for each subset $\nu \subset Z_i$, a new element x_ν such that for $z \in Z$, $x_\nu R z \iff z \in \nu$, with x_ν likewise R-free from previous elements of X_{i+1}.

Y_{i+1}, Z_{i+1} are defined symmetrically. As we are working in the triangle-free random graph, in order that the the construction be able to continue, it is enough that the sets X_i, Y_i, Z_i are each empty graphs, i.e., at no point do we ask for a triangle.

To finish, set $X = \bigcup_i X_i$, $Y = \bigcup_i Y_i$, $Z = \bigcup_i Z_i$. Each pair is a bipartite random graph, as desired. $\square$

Theorem O. *(weak version of Key Lemma, Chapter 2) Fix $1 > \delta > 0$ and a binary edge relation R. Then there exist $\epsilon' = \epsilon'(\delta), N' = N'(\epsilon', \delta)$ such that:* if $\epsilon < \epsilon'$, *$N > N'$, X, Y, Z are disjoint finite sets of size at least N, and each of the pairs $(X, Y), (Y, Z), (X, Z)$ is ϵ-regular with density δ,* then *there exist $x \in X, y \in Y, z \in Z$ so that x, y, z is an R-triangle.*

Obviously, we cannot have an R-triangle in the triangle-free random graph. Nonetheless each of the pairs (X, Y) in Observation 4.1 manifestly has finite subgraphs of any attainable density.

The difficulty comes when we try to choose finite subgraphs $X' \subset X, Y' \subset Y, Z' \subset Z$ so that the densities of all three pairs are *simultaneously* near the same $\delta > 0$. If (X', Y') and (Y', Z') are reasonably dense, (X', Z') will be near 0. Put otherwise, we may choose elements of X independently over Y, and independently over Z, but not both at the same time.

The constructions in this chapter generalize this example, and give a way of measuring the "depth" of independence in a constellation of sets $X_1, \dots X_n$, where any pair (X_i, X_j) is a bipartite random graph. The example of the triangle-free random graph is paradigmatic: we shall see that a bound on the depth of independence will produce the 3-strong order property SOP_3.

4.2 Independence and order

In all of our applications the formula $R(x; y)$ will be symmetric. For the definitions, we just ask that $\ell(x) = \ell(y)$.

Definition 4.2. *(Constellations of order and independence properties)*

1. *The sets A, B* witness the order property for φ *if $|A| = |B| \geq \aleph_0$ and there exist enumerations $A = \langle a_i : i < \rho \rangle$, $B = \langle b_i : i < \rho \rangle$ such that either: $\varphi(a_i, b_j)$ iff*

$i < j$ for all $i, j < \rho$, or $\neg\varphi(a_i, b_j)$ iff $i < j$ for all $i, j < \rho$.

2. *Fix some formula $\varphi(x; y)$. Let A, B be disjoint sets of k- and n-tuples respectively, where $k = \ell(x), n = \ell(y)$. Then A is* independent over *B with respect to φ just in case for any two finite disjoint $\eta, \nu \subset B$, there exists $a \in A$ such that $b \in \eta \rightarrow \varphi(a; b)$ and $b \in \nu \rightarrow \neg\varphi(a; b)$.*

3. *Fix a formula $R(x; y)$. Let $A_1, \ldots A_k$ be disjoint sets (of m-tuples, where $m = \ell(x) = \ell(y)$). Then A_1 is* independent over *$A_2, \ldots A_k$ with respect to R just in case A_1 is independent over $B := \bigcup_{2 \leq i \leq k} A_i$ in the sense of (2).*

4. *The formula $R(x; y)$ is a* bipartite random graph *if there exist disjoint infinite sets A, B such that A and B are each independent over the other wrt R.*

5. *A formula R is X_2^m if there exist disjoint infinite sets $\langle A_i : i < m \rangle$ such that for any two distinct $i, j < m$, the sets A_i, A_j witness the order property for R.*

6. *A formula R is I_k^m, $m \geq k \geq 2$, if there exist disjoint infinite sets $\langle A_i : i < m \rangle$ such that for any distinct $i_1, \ldots i_k < \omega$, A_{i_1} is independent over $\bigcup_{2 \leq j \leq k} A_{i_j}$ w.r.t. R. Notice that k refers to the* scope *of the independence (i.e. the number of columns involved), not the size of the finite disjoint η, ν.*

Observation 4.3. *1. Let $R(x; y)$ be a symmetric formula. If R is I_ω^ω then there is an infinite subset of the monster model on which R is a random graph. (Certainly this need not be definable or interpretable in any way). The converse is also true.*

2. *If $R(x;y)$ has the strict order property then $R'(x;y,z) = R(x,y) \wedge \neg R(x,z)$ is X_2^ω. So this phenomenon occurs both in order and independence.*

3. *If $\varphi(x;y)$ (not necessarily symmetric) has the order property then $\theta(xy;zw) = \varphi(x;w)$ does too, and furthermore θ is X_2^ω and so, assuming $NSOP$, also I_2^ω (see below). But the asymmetry between the parameter and object variables is not erased; the result is really a bipartite random graph for φ, and not a random graph.*

We first sketch a classic proof of Shelah, as the details of its argument will be used in Lemma 4.5. Recall that:

Definition 4.4. *([26]) A formula $\varphi(x;y)$ has the* strict order property *if there exists an infinite sequence $\langle a_i : i < \omega \rangle$ such that: $\exists x(\varphi(x;a_i) \wedge \neg\varphi(x;a_j))$ iff $j < i$.*

Theorem P. *(Shelah, [26]:Theorem II.4.7)*

Suppose $\varphi(x;y)$ is unstable. Then either φ has the independence property, or for some $n < \omega, \eta \in 2^n$, $\bigwedge_{i<n} \varphi(x;y_i)^{\eta[i]}$ has the strict order property.

Proof. (Sketch) Let $A = \langle a_i : i < \omega \rangle$ be a sequence on which φ has the order property, i.e. for any $j \in \omega$, $\exists x(i \leq j \rightarrow \neg\varphi(x;a_i) \;\wedge\; i > j \rightarrow \varphi(x;a_i))$. By convention, $\varphi(x;y)^0 \equiv \neg\varphi(x;y)$.

Suppose that φ does not have the independence property on A, so there are $n < \omega$, $a_0, \ldots a_n$ and $\sigma \subset n$, $|\sigma| = k < n$ such that

$$\neg\exists x(\bigwedge_{i\leq n}\varphi(x;a_i)^{if\ \ i\in\sigma}) \tag{4.1}$$

Because of the order property, there is some $\tau \subset n$, $|\tau| = k$ such that

$$\exists x(\bigwedge_{i\leq n}\varphi(x;a_i)^{if\ \ i\in\tau}) \tag{4.2}$$

namely, $\tau = \{n-k+1, \ldots n\}$. Thus, for some finite m, we can find a sequence of m sets $\sigma_i \subset n$, such that $|\sigma_i| = k$, $\sigma_0 = \tau$, $\sigma_m = \sigma$, and for each i, σ_{i+1} is obtained from σ_i by swapping two consecutive elements. Let $t+1 \leq m$ be the first time when (2) yields to (1) and the witness vanishes. By indiscernibility of A, $\neg\exists x(\bigwedge_{i\leq n}\varphi(x;y_i)^{if\ \ i\in\sigma_{t+1}})$ is true of any increasing n-tuple from A of the same order-type. In particular, suppose that the inconsistency appeared when swapping elements $a_\ell, a_{\ell+1}$. Write $p(y_1, \ldots y_n)$ for the type $\bigwedge_{i\leq n}\varphi(x;y_i)^{if\ \ i\in\sigma_t}$, and $q(y_1, \ldots y_{\ell-1}, y_{\ell+2}, \ldots y_n)$ for p on all but the crucial pair $y_\ell, y_{\ell+1}$.

By choosing an n-tuple from A whose first $\ell-1$ elements A_0 and last $n-(\ell+2)$ elements A_1 are sufficiently far apart, we get arbitrarily large indiscernible sequences in between on which $\exists x(q(A_0, A_1) \wedge \varphi(x;a_i) \wedge \varphi(x;a_j))$ is true iff $i < j$. This gives strict order. $\square$

Let us turn the construction around.

Lemma 4.5. *Suppose the formula $R(x;y)$ is X_2^2. If T does not have the strict order property, then R is I_2^2.*

Proof. Let us begin with an indiscernible sequence $\langle a_i b_i : i < \mathbb{Q} \rangle$ which witnesses the order property for R. Let $A_0 = \{a_i : i < \mathbb{Q}\}$, $B = \{b_i : i < \mathbb{Q}\}$. Let $D_0 = \emptyset$. Fix some increasing ω-indexed subset of A_0; call this C_0, and set $A := A_0 \setminus C_0$.

We will construct two countable increasing sequences of sets C_n, D_n $(n < \omega)$ so that $(\bigcup_n C_n, \bigcup_n D_n)$ is a bipartite random graph. The idea is that at stage n, the "independent" sets C_n, D_n will be relatively sparse inside A, B respectively, which serve as a large scaffolding for the construction. More precisely:

Inductive hypothesis. A, B are fixed. At stage $n \geq 1$, suppose that:

1. C_n, D_n are countably infinite.
2. $(A \cup C_n, B)$ and $(A, B \cup D_n)$ are indiscernible sequences witnessing the order property for R.
3. D_n is independent over C_n.
4. C_n is independent over D_{n-1}.

We have defined A, B, C_0, D_0. We give a definition for any D_n $(n \geq 1)$. This will be essentially symmetric for C_n: the only important difference is that C_n in the first sentence becomes D_{n-1} rather than D_n.

Definition of D_n. Fix some enumeration $f : \omega \to \omega^{<\omega} \times \omega^{<\omega}$ of pairs of disjoint finite subsets of C_n. For each $i < \omega$, we want to choose an element d_i such that, writing D_n^i for $\{d_j : j \leq i\}$,

1. this instance of independence is satisfied: $d \models p_{\sigma,\tau}(x) := (y \in \sigma \rightarrow R(x,y),$ $z \in \tau \rightarrow \neg R(x,y))$

2. the order property between $(A, B \cup D^i_n)$ is preserved: d_i realizes some cut $q(x) \in S(A)$ over A, which is not $\pm\infty$, and not the same cut as any element previously added.

3. the *indiscernibility* of $(A, B \cup D^i_n)$ is preserved: d_i realizes the correct average type $r(x) \in S(B \cup D_{n-1})$ (if $i = 0$), or $\in S(B \cup D^j_n)$ for $j = i - 1$ (otherwise).

Let p, q', r' be finite fragments of the three types. $p(x) \cup q'(x)$ together can be thought of as an instance of independence. To show that it is consistent, we now threaten strict order by running through Shelah's argument from Theorem P, with the additional criterion r: that is, we ask after each swap whether there exists an x such that $r(x)$ and ... . This is a legitimate addition because at the first stage of the construction, when the elements are in the expected order, there are many witnesses from the sequence B. Almost any element from B will satisfy r, as the inductive hypothesis says that $(A, B \cup D^j_n)$ is indiscernible, for $j < i$. $\square$

The following Corollary will be useful for §3.5.

Corollary 4.6. *Suppose that we strengthen the hypotheses of Lemma 4.5 as follows. We are given a characteristic sequence $\langle P_n \rangle$ and an infinite indiscernible sequence $\langle a_i b_i : i < \omega \rangle$ of pairs of elements of P_1 on which:*

- $P_2(a_i, b_j) \iff i < j$

- *$P_2(a_i, a_j)$ and $P_2(b_i, b_j)$ for all $i, j < \omega$*

- *$P_n(a_{i_1}, \ldots a_{i_k}, b_{j_1}, \ldots b_{j_{n-k}})$ iff P_2 holds on all 2-element subsets of $\{a_{i_1}, \ldots a_{i_k}, b_{j_1}, \ldots b_{j_{n-k}}\}$.*

i.e. P_2 is X_2^2 on this sequence, but satisfies some additional compatibility conditions.

Then there exist sets $C, D \subset P_1$ such that:

(i) (C, D) is a bipartite random graph for P_2

(ii) $P_2(c, c')$ and $P_2(d, d')$ for all $c, c' \in C$ and $d, d' \in D$

(iii) $P_n(c_{i_1}, \ldots c_{i_k}, d_{j_1}, \ldots d_{j_{n-k}})$ iff P_2 holds on all 2-element subsets of $\{c_{i_1}, \ldots c_{i_k}, d_{j_1}, \ldots d_{j_{n-k}}\}$

In other words, in carrying out the construction of Lemma 4.5 we can without loss of generality retain the compatibility of each half as well as support 2.

Proof. At the inductive step, to the conditions (1)-(3) in the definition of d_i in the proof of Lemma 4.5 we add the following schemata:

(4) $d_i \cup B \bigcup_{j \leq n-1} D_j \cup D_n^i$ is a P_∞-complete graph

$(5)_k$ $(k < \omega)$ for all $X \subset A \cup \bigcup_{j \leq n-1} C_j \cup B \bigcup_{j \leq n-1} D_j \cup D_n^i$, $X^2 \subset P_2$ and $d_i \in X$ implies $X^k \subset P_k$

That is, we would like to show that we can add the condition of support 2 to our inductive hypothesis. At stage $t + 1$ in the construction (fixing, if necessary, an enumeration of the countably many but not ω-indexed stages in the proof of the previous Lemma), let S_t be the union of A, B and every element added thus far. The claim is that we can choose the element at stage $t+1$ so that S_{t+1} has support 2.

At stage 0, this is true by hypothesis. At stage $t+1$, without loss of generality suppose we are adding an element to D^i_n. What the argument of Lemma 4.5 shows is that any finite fragment of the type p, q', r' which our new element must satisfy is realized by many elements of the original sequence B, provided that p over σ, τ is swapped to its "initial position" in that argument, i.e. in linear order in the sense of the (A, B)-order property. A fortiori, any element b from B has the property that $S_t \cup \{b\}$ has support 2. So we in fact have a witness to the larger finite fragment including pieces of (4), (5). This is all we need to run the argument from Lemma 4.5, progressively swapping the elements in σ, τ, at each stage using the threat of *SOP* to continue. $\square$

4.3 Towards SOP_3

Definition 4.7. (Shelah, [27]:Definition 2.5) *For* $n \geq 3$, *the theory* T has SOP_n *if there is a formula* $\varphi(x; y)$, $\ell(x) = \ell(y) = k$, $M \models T$ *and a sequence* $\langle a_i : i < \omega \rangle$ *with each* $a_i \in M^k$ *such that:*

1. $M \models \varphi(a_i, a_j)$ *for* $i < j < \omega$

2. $M \models \neg \exists x_1, \ldots x_n (\bigwedge \{\varphi(x_m, x_k) : m < k < n \text{ and } k = m + 1 \text{ mod } n\})$

Theorem Q. (Shelah, [27]: (1) is Claim 2.6, (2) is Theorem 2.9)

1. *For a theory T, $SOP \implies SOP_{n+1} \implies SOP_n$, for $n \geq 3$ (not necessarily for the same formula).*

2. *If T is a complete theory with SOP_3, then T is maximal in the Keisler order.*

For an alternate proof of (2), see Chapter 3, Theorem 3.32.

We will derive SOP_3 from failures of randomness, using the following equivalent definition. Remember that, by convention, $a_i, x, \ldots$ need not be singletons.

Definition 4.8. ([28]:Fact 1.3) *T has SOP_3 iff there is an indiscernible sequence $\langle a_i : i < \omega \rangle$ and $\mathcal{L}$-formulas $\varphi(x;y), \psi(x;y)$ such that:*

1. *$\{\varphi(x;y), \psi(x;y)\}$ is contradictory.*

2. *there exists a sequence of elements $\langle c_j : j < \omega \rangle$ such that*

   - $i \leq j \implies \varphi(c_j; a_i)$

   - $i > j \implies \psi(c_j; a_i)$

3. *if $i < j$, then $\{\varphi(x; a_j), \psi(x; a_i)\}$ is contradictory.*

The idea of the construction is contained in the following straightforward example.

Example 4.1. *Let T be the triangle free random graph, with edge relation R. Then R is I^3_2 but not I^3_3, and T is SOP_3.*

Proof. Let us prove the final clause (for the rest see Observation 4.1 and the discussion following).

The theory by definition contains a forbidden configuration, a triangle. Suppose A, B, C are disjoint infinite sets witnessing I_2^3. Let us construct a sequence of triples $S = \langle a_i, b_i, c_i : i < \omega \rangle$ such that, for $i < \omega$,

- For all $j \leq i$, $b_i R a_j$.
- For all $j \leq i$, $c_i R b_j$.
- For all $j \leq i$, $a_{i+1} R c_j$.

Define a binary relation $<_\ell$ on triples by:

$$(x, y, z) \leq_\ell (x', y', z') \iff ((xRy' \wedge yRz' \wedge zRx'))$$

While $<_\ell$ need not be a partial order on the model, it does linearly order the sequence S by construction. Looking towards Definition 4.8, let us define two new formulas (the variables t stand for triples):

- $\varphi(t_0; t_1, t_2) = t_1 <_\ell t_2 <_\ell t_0$
- $\psi(t_0; t_1, t_2) = t_0 <_\ell t_1 <_\ell t_2$

Let us check that these formulas give SOP_3. For condition (1), $\varphi(t_0; t_1, t_2), \psi(t_0; t_1, t_2)$ means that $(x_0, y_0, z_0) <_\ell (x_1, y_1, z_1) <_\ell (x_2, y_2, z_2) <_\ell (x_0, y_0, z_0)$. Then $x_i R y_j$, $y_j R z_k$, $z_k R x_i$ which gives a triangle, contradiction.

It is straightforward to satisfy (2) by compactness (e.g. by choosing S codense in a larger indiscernible sequence).

Finally, for condition (3), suppose $i < j$ but $\varphi(t; \gamma_i), \psi(t; \gamma_j)$ is consistent, where $t = (x, y, z)$. This means that $(x, y, z) <_\ell (a_i, b_i, c_i) <_\ell (a_j, b_j, c_j) <_\ell (x, y, z)$ (where the middle $<_\ell$ comes from the behavior of $<_\ell$ on the sequence S). As in condition (1), this gives a triangle, contradiction. $\square$

We can in fact build a much larger engine for producing the rigidity of SOP_3 from a forbidden configuration.

Theorem 4.9. *Suppose that for some $2 \leq n < \omega$, the formula R of T is I^{n+1}_n but not I^{n+1}_{n+1}. Then T is SOP_3.*

Proof. The construction is arranged into four stages.

Step 1: Finding a universally forbidden configuration G.

By hypothesis, R is not I^{n+1}_{n+1}. This means that the infinitary type $p(X_0, \ldots X_n)$, which describes $n+1$ infinite sets X_i which are I^{n+1}_{n+1} in the sense of Definition 4.2, is not consistent. Let G be a finite inconsistent subset of height h, in the variables $V_G = \{x^i_j : 1 \leq i \leq h, 0 \leq j \leq n\}$, and described by the edge map $E_G : \{((i,j),(i',j')) : i, i' \leq h, j \neq j' \leq n\} \to \{0, 1\}$. As the inconsistency of p is a consequence of T, G will be a universally forbidden configuration:

$$T \vdash \neg(\exists x_0^1, \dots x_n^h) \left(\bigwedge_{i,i' \leq h,\ j \neq j' \leq n} R(x_j^i, x_{j'}^{i'}) \iff E((i,j),(i',j')) = 1 \right) \tag{4.3}$$

Note that the configuration remains agnostic on edges between elements in the same column, in keeping with the definition of I_ℓ^m.

In what follows G will appear as a template which we shall try to approximate using I_n^{n+1}. Here are the vertices of G arranged as they will be visually referenced (the edges are not drawn in):

$$\begin{array}{ccccc}
x_0^h & & x_k^h & & x_n^h \\
\vdots & & \vdots & & \vdots \\
x_0^\rho & & x_k^\rho & & x_n^\rho \\
\vdots & & \vdots & & \vdots \\
x_0^1 & \dots & x_k^1 & \dots & x_n^1
\end{array}$$

Figure 4.1: Vertices of the forbidden configuration G, arranged in columns. When comparing this configuration to an array whose rows are indexed modulo h, the superscript of the top column becomes 0.

Step 2: Building an array A of approximations to G.

Let $A_0, \dots A_n$ be disjoint infinite sets witnessing I_n^{n+1} for R. As in Example 4.1, we will use elements from these columns A_i to build an array $A = \langle a_i^\rho : 1 \leq \rho < \omega, 0 \leq i \leq n \rangle$. Fixing notation,

- $a_0^\rho, \dots a_n^\rho$ is called a *row*.
- $\mathrm{Col}(i) = \{j : j \neq i, i+1 \pmod{n+1}\}$ is the set of column indices associated to the column index i.

- Define an ordering on pairs of indices (β for "before"):

$$\beta((t', i'), (t, i)) \iff_{def}$$
$$\Big((t' < t \wedge i' \in \mathrm{Col}(i)) \vee (t' = t \wedge i' < i)\Big)$$

Claim 4.10. *We may build the array A to satisfy:*

1. *For all ρ, $a^\rho_k \in A_k$.*
2. *For any ρ', ρ, k, k' such that $\beta((\rho', k'), (\rho, k))$,*

$$a^\rho_k \ R \ a^{\rho'}_{k'} \iff E_G((r, k), (r', k')) = 1$$

where $r \equiv \rho$ (mod h), $r' \equiv \rho'$ (mod h).

Proof. We choose elements in a helix $(a^1_0, a^1_1, \ldots a^1_n, a^2_0, a^2_1, \ldots)$ so that $\beta((\rho', k'), (\rho, k))$ implies that $a^{\rho'}_{k'}$ is chosen before a^ρ_k.

When the time comes to choose a^ρ_k, we look for an element of A_k which satisfies Condition (2) of the Claim, that is, which, by Condition (1), realizes a given R-type over disjoint finite subsets of the columns A_i ($i \in \mathrm{Col}(k)$). As $(A_0, \ldots A_n)$ was chosen to be I^{n+1}_n and $|\operatorname{Col} k| = n - 1$, an appropriate a^ρ_k exists. $\square$

Step 3: Defining the relation $<_\ell$, which has no pseudo-$(n+1)$-loops.

$$
\begin{array}{cccccccc}
 & & & & \vdots & & & \\
 & & & & \vdots & & & \\
 & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- \\
 & a_0^{\ell_n h+h} & & \cdots & & a_k^{\ell_n h+h} & & \mathbf{a}_{\mathbf{n}}^{\ell_{\mathbf{n}}\mathbf{h}+\mathbf{h}} \\
 & \vdots & & & & \vdots & & \vdots \\
B_{\ell_n} = & a_0^{\ell_n h+\rho} & & & & a_k^{\ell_n h+\rho} & & \mathbf{a}_{\mathbf{n}}^{\ell_{\mathbf{n}}\mathbf{h}+\rho} \\
 & \vdots & & & & \vdots & & \vdots \\
 & a_0^{\ell_n h+1} & & & & a_k^{\ell_n h+1} & & \mathbf{a}_{\mathbf{n}}^{\ell_{\mathbf{n}}\mathbf{h}+\mathbf{1}} \\
 & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- \\
 & & & & \vdots & & & \\
 & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- \\
 & a_0^{\ell_k h+h} & & \cdots & & \mathbf{a}_{\mathbf{k}}^{\ell_{\mathbf{k}}\mathbf{h}+\mathbf{h}} & & a_n^{\ell_k h+h} \\
 & \vdots & & & & \vdots & & \vdots \\
B_{\ell_k} = & a_0^{\ell_k h+\rho} & & & & \mathbf{a}_{\mathbf{k}}^{\ell_{\mathbf{k}}\mathbf{h}+\rho} & & a_n^{\ell_k h+\rho} \\
 & \vdots & & & & \vdots & & \vdots \\
 & a_0^{\ell_k h+1} & & \cdots & & \mathbf{a}_{\mathbf{k}}^{\ell_{\mathbf{k}}\mathbf{h}+\mathbf{1}} & \cdots & a_n^{\ell_k h+1} \\
 & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- \\
 & & & & \vdots & & & \\
 & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- \\
 & a_0^{2h} & \mathbf{a}_{\mathbf{1}}^{\mathbf{2h}} & & & a_k^{2h} & & a_n^{2h} \\
 & \vdots & \vdots & & & \vdots & & \vdots \\
B_1 = & a_0^{h+\rho} & \mathbf{a}_{\mathbf{1}}^{\mathbf{h}+\rho} & & & a_k^{h+\rho} & & a_n^{h+\rho} \\
 & \vdots & \vdots & & & \vdots & & \vdots \\
 & a_0^{h+1} & \mathbf{a}_{\mathbf{1}}^{\mathbf{h}+\mathbf{1}} & & & a_k^{h+1} & \cdots & a_n^{h+1} \\
 & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- & -\,-\,- \\
 & \mathbf{a}_{\mathbf{0}}^{\mathbf{h}} & & & & a_k^{h} & & a_n^{h} \\
 & \vdots & & & & \vdots & & \vdots \\
B_0 = & \mathbf{a}_{\mathbf{0}}^{\rho} & & & & a_k^{\rho} & & a_n^{\rho} \\
 & \vdots & & & & \vdots & & \vdots \\
 & \mathbf{a}_{\mathbf{0}}^{\mathbf{1}} & & & \cdots & a_k^{1} & \cdots & a_n^{1}
\end{array}
$$

Figure 4.2: Elements of the array A, arranged in blocks of h rows. The boldface refers to Step 4 of the proof, when a proposed witness to G is assembled from the ith columns of blocks B_i in a pseudo-$(n+1)$-loop.

We now define a binary relation $<_\ell$ on m-tuples, where $m = h(n+1)$. Fix the enumeration of these tuples to agree with the natural interpretation as blocks B_ℓ of h consecutive rows in the array A (see Figure 4.3). That is, write the variables $Y := \langle y_i^t : 1 \leq t \leq h, 0 \leq i \leq n \rangle$, $Z := \langle z_{i'}^{t'} : 1 \leq t' \leq h, 0 \leq i' \leq n \rangle$. Define:

$$\mathbf{Y} <_\ell \mathbf{Z} \iff_{(def)}$$

$$\bigwedge_{1 \leq t', t \leq h,\ 0 \leq i, i' \leq n} (i' \in \mathrm{Col}(i)) \implies \left(z_i^t \; R \; y_{i'}^{t'} \iff E_G((t,i),(t',i')) = 1 \right)$$

Let B be a partition of the array A into blocks B_k $(k < \omega)$ each consisting of h consecutive rows, so $B_k := \langle a_t^r : 0 \leq t \leq n, kh+1 \leq r \leq (kh)+h \rangle$, for each $k < \omega$ (see Figure 4.3). By Claim 4.10, $i \lneq j \implies B_i <_\ell B_j$.

Definition 4.11. *A* pseudo-$(n+1)$-loop *is a sequence* W_i $(0 \leq i \leq n)$ *such that for some* m, $1 \leq m < n$:

$$\left(\bigwedge_{(0<j<i\leq n)} W_j <_\ell W_i \right) \wedge \left(\bigwedge_{1 \leq j \leq m} W_0 <_\ell W_j \right) \wedge \left(\bigwedge_{m < j \leq n} W_j <_\ell W_0 \right) \tag{4.4}$$

Suppose it were consistent with T to have blocks of variables $W_0 \dots W_n$ which form a pseudo-$(n+1)$-loop. Write $W_k(i) = \{w_i^{hk+1}, \dots w_i^{hk+h}\}$ for the ith column of block W_k. Figure 4.3 gives the picture, where the elements a are replaced by variables w and the blocks B_i become W_i. Set $W_G = W_0(0) \cup \dots \cup W_n(n)$ (which can be visualized as the boldface columns in Figure 4.3).

By definition of $<_\ell$, the pseudo-(n+1)-loop (4.4) implies that whenever

$$((j \in \mathrm{Col}(i)) \wedge ((0 < j < i \leq n) \vee (j = 0 \wedge i \leq m) \vee (m < j \wedge i = 0)))$$

we will have:

$$\left(\forall\ w^t_k \in W(i),\ w^{t'}_{k'} \in W(j)\right) \left(w^t_k\ R\ w^{t'}_{k'} \iff E_G((t,k),(t',k')) = 1\right)$$

In other words, $<_\ell$ says that on certain pairs of elements in our proposed instance W_G of G, namely those elements whose respective columns "fall into each other's scope" as given by the Col operator, W_G faithfully follows the template of G. It is easy to check that in a pseudo-$(n+1)$-loop every pair $j \neq i$ in $\{0, \ldots n\}$ has this property. Thus pseudo-$(n+1)$-loops in $<_\ell$ are inconsistent with T.

Step 4: Obtaining SOP_3.

Step 3 showed that our array A of approximations had a certain rigidity, which we can now identify as SOP_3. Following Definition 4.8, let us define $\varphi_r(x; y_1, \ldots y_n)$ and $\psi_\ell(x; y_1, \ldots y_n)$, where the the variables are blocks, and the subscripts "ℓ" and "r" are visual aids: the element x goes to the left of the elements y_i under ψ, and to their right under φ.

That is, we set:

- $\varphi_r(x; y_1, \ldots y_n) =$

$$\bigwedge_{1 \leq i \neq j \leq n} y_i <_\ell y_j \ \wedge \ \bigwedge_{1 \leq i \leq n} y_i <_\ell x$$

- $\psi_\ell(x; y_1, \dots y_n) =$

$$\bigwedge_{1 \le i \le n} x <_\ell y_i \ \wedge \bigwedge_{1 \le i \ne j \le n} y_i <_\ell y_j$$

Now let us verify that the conditions of Definition 4.8 hold. Let B be the sequence of blocks defined in Step 3, and assume without loss of generality that $B = \langle B_k : k < \omega \rangle$ is indiscernible and moreover is dense and codense in some indiscernible sequence B'. Let $A = \langle A_i : i < \omega \rangle$ be an indiscernible sequence of n-tuples of elements of B.

1. $\{\varphi_r(x; y_1, \dots y_n), \psi_\ell(x; y_1, \dots y_n)\}$ is contradictory because it gives rise to a pseudo-$(n+1)$-loop.

2. By construction, for any $k < \omega$, the type

$$\{\psi_\ell(x; A_j) : j \le k\} \cup \{\varphi_r(x; A_i) : k < i\}$$

   is consistent, because $<_\ell$ linearly orders B, thus also B'. Choose the desired sequence of witnesses to be elements in the indiscernible sequence B' which are interleaved with B.

3. Suppose we have $\{\varphi_r(x; A_j), \psi_\ell(x; A_i)\}$ for some $i < j$, or in other words:

$$\{\varphi_r(x; B_{j_1}, \dots B_{j_n}), \psi_\ell(x; B_{i_1}, \dots B_{i_n})\} \ \text{ where } \ \{i_1, \dots i_n\} < \{j_1, \dots j_n\}$$

   Then $x <_\ell B_{i_1} <_\ell \cdots <_\ell B_{i_n} <_\ell B_{j_1} <_\ell \cdots <_\ell B_{j_n} <_\ell x$ is a pseudo-$(2n+1)$-loop (remember that $<_\ell$ holds between any increasing pair of elements of B by construction). Thus a fortiori we have a pseudo-$(n+1)$-loop, contradicting the conclusion of Step 3.

We have shown that the theory T has SOP_3, so we finish. $\square$

Chapter 5

Characteristic sequences and ultrapowers

This chapter builds the bridge necessary to apply the techniques of previous chapters to the realization of types in ultrapowers. In Section 5.1 we consider static and dynamic arguments in the characteristic sequence, i.e. the attempt to fit the countably many optimized predicates together to realize a type versus the attempt to optimize some given P_n. In the second section, we collect some results on transferring structure between sequences.

5.1 Static and dynamic arguments

This chapter builds on Chapters 1-2, but we recall here:

Convention 5.1. *Throughout this chapter:*

- *T is a first-order theory, φ is a formula of the language of T, $\langle P_n \rangle$ is the characteristic sequence of (T, φ), in the notation of Chapter 2. T_0, T_1 are as in Conventions 2.5-2.6.*
- $\langle P_n : n < \omega \rangle$ has support k *if for all finite $B \subset P_1$, $B^k \subset P_k$ implies that $B^n \subset P_n$, for all n.*
- *$M \models T$, $\mathcal{D}$ is a regular ultrafilter on λ, $N = M^\lambda/\mathcal{D}$, Definition 1.1.* Small *means of size $\leq \lambda$.*
- *A set $X \subset M^\lambda/\mathcal{D}$ is* induced *if is the product of its projections to the index models, Definition 1.26.*
- *Recall from Definition 2.1 that a consistent partial φ-type corresponds in the characteristic sequence to a complete P_∞-graph called a* positive base set, *i.e. $A \subset P_1$ such that $A^n \subset P_n$ for all n. Any such A gives rise to a consistent partial type $\{\varphi(x;a) : a \in A\}$. We may refer simply to* the type A.
- *A* distribution *$d : \mathcal{P}_{\aleph_0}(p) \rightarrow \mathcal{D}$ is a monotonic map from finite subsets of a small, consistent partial type p into the filter which refines the Łoś map and whose image is a regularizing set, Definition 1.4. Occasionally p will be identified with the corresponding A, in which case the distribution will be described on the singletons and naturally extended to all finite subsets of A.*

- *k^*, m^* are nonstandard integers used to indicate particular sizes, Remark 5.3.*

The following observation highlights the two key steps in our project of using characteristic sequences to analyze φ-types.

Observation 5.2. *The following are equivalent for a complete P_∞-graph A.*

1. *The type p corresponding to A is realized.*
2. *There exists an induced set $X \subset P_1^N$ and a distribution $d : A \to \mathcal{D}$ such that:*
   (a) $N \models A \subset X$
   (b) $N \models X^n \subset P_n$, *for all n*
   (c) *for each index $t \in \lambda$, $\{a[t] : t \in d(a)\} \subset X[t]$ in the index model $M[t]$*
   (d) *for each index $t \in \lambda$, if*
   $m[t] := |\{a[t] : t \in d(a)\}|$ *and*
   $k[t] := max\{k \leq m[t] : (X[t])^k \subset P_k \text{ in } M[t]\}$
   then $m[t] \leq k[t]$.
3. *There exists a distribution $d : A \to \mathcal{D}$ such that almost everywhere, $A[t]$ is a $P_{m[t]}$-complete graph.*

Proof. (2) $\rightarrow$ (1): The conditions guarantee that the $m[t]$-tuple of elements represented in each index model is in $P_{m[t]}$, and thus that the corresponding instances of φ have a common witness $\alpha[t]$. Then $\alpha := \Pi_t \alpha[t]/\mathcal{D}$ realizes the type p, by definition of distribution.

(1) $\to$ (2): By Chapter 1, Observation 1.10, p has a distribution which is multiplicative. This says precisely that the $m[t]$ elements assigned to each index model form a $m[t]$-complete graph.

(2) $\leftrightarrow$ (3): (3) is a direct translation, which amounts to saying that the set assigned to each index model is a complete graph (where "complete" means in its own arity). $\square$

Remark 5.3. *The project is therefore to analyze the comparative complexity of types A, A' by comparing:*

- *The relative difficulty of finding, for each $n < \omega$, an induced complete n-graph X_n containing A: in other words, considering the complexity of P_n as a hypergraph in the neighborhood of the (P_n-complete, but almost always not definably so) set A.*

- *The relative difficulty of putting these countably many predicates together in order to produce a realization of the type: the nonstandard integer k^* counting the uniform degree of consistency may be much smaller than the nonstandard integer m^* counting the number of elements assigned to each index model.*

Notice that each of these endeavors captures a different quality of limit argument. The search for an appropriate X_n involves *dynamic* arguments of the general form: If a particular P_n-configuration persists under any finite localization of P_1 around A, then something must be true of φ (φ is unstable, φ is not simple...). Whereas the slippage of putting the countably many predicates together involves *static* arguments of the general form: If something is true for all P_n, thus for P_∞, then something must be true of φ (φ has the finite cover property, φ is not low).

Remark 5.4. *The following are equivalent, for $(T, \varphi) \mapsto \langle P_n \rangle$:*

1. *There is $k < \omega$ such that the sequence $\langle P_n \rangle$ has support k.*

2. *φ does not have the finite cover property (Definition 0.9).*

When φ is unstable, some fixed finite conjunction of instances of φ has the finite cover property (Theorem C, page 16). Nonetheless, it may happen that there is a set $\Sigma \subset \mathcal{L}$ of formulas without the fcp such that $M \models T$ is λ^+-saturated iff M realizes all φ_0-types over sets of size λ for all $\varphi_0 \in \Sigma$. This is true, for instance, of $\Sigma = \{\psi(x; y, z) := xRy \land \neg xRz\}$ in the random graph, and of $\Sigma = \{\psi(x; y, z) := y < x < z\}$ in $(\mathbb{Q}, <)$.

When φ is a formula (without parameters) which does not have the finite cover property, issue (2) of Remark 5.3 is irrelevant:

Claim 5.5. *Suppose that φ does not have the finite cover property in T. Then any small φ-type is realized in any regular ultrapower of a model of T.*

Proof. By the previous observation $\langle P_n \rangle$ has support k. As φ nfcp implies φ stable, by Chapter 2 we can find an induced $X \subset P_1$, $A \subset X$ which is a complete P_k-graph; as this is a first-order property, it will be true almost everywhere. Distribute the elements of A so that $t \in d(a)$ implies that (i) $M[t] \models a[t] \in X$ and (ii) $X^{M[t]}$ is a complete P_k-graph. Thus, in *every* index model $M[t]$, the set $\{a[t] : t \in d(a)\}$ is a complete P_∞-graph, because "k-consistency implies n-consistency" is, for each n, a consequence of T and thus true in every index model. Notice that the X adds no essential structure in this argument: it is in fact a P_1-definable set, with parameters (Chapter 2), and serves simply as a template for refining the distribution. $\square$

We will not derive here Shelah's proof that countable stable theories fall into two equivalence classes in the Keisler order (or that this saturation depends on φ-types). However, to indicate the subtlety of fcp:

Example 5.1. *(fcp and parameters) Let $M \models T$, $N = M^\lambda/\mathcal{D}$. Suppose that $\varphi(x; y, z)$ has the finite cover property modulo T but for some parameter $a \in N$ $\varphi(x; y, a)$ does not have the fcp. Then the proof of Claim 5.5 may fail. For instance, let T be the theory of a single equivalence relation $E(x, y)$ with a class of size n for every n, and let $a \in N$ be an element in some nonstandard equivalence class. Then $\varphi(x; y, a) = E(x, a) \wedge x \neq a$ does not have the fcp but it is still possible for clause 2(d) of Observation 5.2 to fail, index model by index model, in the finite equivalence class $E(x, a[t])$.*

Our focus here will be the unstable case, where we can always find limit predicates, as we may assume that the lower cofinality $lcf(\omega, \mathcal{D}) \geq \lambda^+$ (see Definition 1.23, Theorem G).

Observation 5.6. *Suppose that $lcf(\omega, \mathcal{D}) \geq \lambda^+$, $N = M^\lambda/\mathcal{D}$ and $A \subset N$ is small. Let X_i $(i < \omega)$ be a sequence of induced predicates in N satisfying $X_i \supset X_{i+1} \supset A$, for all $i < \omega$. Then there exists an induced predicate P_∞ such that for each i, $X_i \supset X_\infty \supset A$.*

Proof. This is Chapter 1, Lemma 1.28. Notice, however, an important difference in the way it is used. In Chapter 1, the limit predicates X_∞ sit inside a concentric sequence X_i of definable complete graphs extending A. In other words, they are consistent types all the way down; the refinements are to control the interaction of these X_i with partial types in other formulas. Thus the fact that X_∞ is interpreted to be X_k in the index model $M[t]$ still allows us to realize the type X_k in $M[t]$.

A new problem, however, is highlighted by Observation 5.2.2(d), if we think of the predicate X as X_∞ which is interpreted as a P_k-complete graph X_k in the index model $M[t]$. Here k is not simply the amount of the sequence which the index model is able to code but *also* a degree of consistency. In Obs. 5.2.2(d) we cannot find a witness for more than k elements of X_k at a time, so the difference in size of the nonstandard integers k^* and m^* measuring consistency and size of the distribution, respectively, becomes important. □

Definition 5.7. *(Flexibility) Let $\mathcal{D}$ be a regular ultrafilter.*

1. *Let $X = \{X_i : i < \lambda\}$ be a regularizing set. For each index $t < \lambda$, set $n_t := |\{i : t \in X_i\}|$. The* size *of X is the nonstandard integer $n^* = \Pi_t n_t/\mathcal{D}$.*

2. *Say that $\mathcal{D}$ is* flexible *if for every nonstandard integer n^*, $\mathcal{D}$ contains a regularizing set of size $\leq n^*$.*

3. *The theory T is* low *if for every formula φ there exists $k < \omega$ such that if φ divides it $\leq k$-divides.*

Observation 5.8. *If φ is stable then φ is low.*

Proof. To show that φ has the independence property, it suffices to establish the consistency of the following schema. For $k < \omega$, Ψ_k says that there there exist $y_1, \dots y_{2k}$ such that for every $\sigma \subset 2k$, $|\sigma| = k$,

$$\exists x \left(\varphi(x; y_i) \iff i \in \sigma \right)$$

But Ψ_k will be true on any subset of size $2k$ of an indiscernible sequence on which φ is k-consistent but $(k+1)$-inconsistent, and such sequences exist for arbitrarily large k by hypothesis of non-lowness. $\square$

Remark 5.9. *It was shown in Chapter 1, Lemma 1.21 that saturation of a non low theory requires the filter to be flexible. This establishes flexibility as a property which is "seen" by theories, meaning that its absence can be detected by failures of saturation.*

Thus it is a non-trivial property, i.e. one which any ultrafilter which saturates all countable theories must be able to handle:

Corollary 5.10. *(of Chapter 1, Lemma 1.21)*

1. *Not all regular filters are flexible. In particular, a filter with $\mu < \lambda$ will not be flexible.*

2. *Good filters must be flexible.*

The utility of flexibility is that it can reconcile the sizes of any two nonstandard integers.

Claim 5.11. *Suppose $\mathcal{D}$ is flexible, $N = M^\lambda/\mathcal{D}$. Let $A \subset N$ be small, and suppose that for each $n < \omega$ there is an induced X_n, $A \subset X_n \subset P_1$ such that $N \models (X_n)^n \subset P_n$. Then the type A is realized in N.*

Proof. We simply need to show that we can refine the distribution of A to satisfy the conditions of Observation 5.2.2, in particular clause (d). In the notation of that Observation, let $k^* = \Pi_t k[t]/\mathcal{D}$ be the product of the degrees of consistency $k[t]$ (in index model $M[t]$), modulo $\mathcal{D}$; the worry is that $k^* << m^*$. Let $\{Y_i : i < \lambda\}$ be any regularizing set below k^*. Refine the distribution of A by intersecting $d(a_i)$ with Y_i; the resulting $m^* \leq k^*$ modulo $\mathcal{D}$, so we are done. $\square$

Remark 5.12. *Flexibility is sufficient for the conclusion of Claim 5.11, but it may not be necessary, provided we add some conditions on the theory. Many theories will not be able to code all possible discrepancies between k^* and m^*, with nfcp theories as an obvious example.*

As an aside, non lowness cannot be *necessary* for flexibilty as we know that there is a maximal theory, namely $(\mathbb{Q}, <)$, which is low.

Conclusion 5.13. *Whenever we may assume the filter $\mathcal{D}$ is flexible (e.g. if T is non low), the distance between k^* and m^* in any distribution is immaterial, which means we may focus on finding predicates for each of the complete P_n-graphs X_n, without worrying about how to put them together. Otherwise, the distances $k^* << m^*$ which T is able to represent will play a role.*

5.2 Comparing sequences

Let us write down an explicit way in which sequences are comparable. This will be useful in applying arguments which appeal to particular configurations, e.g. those of Chapter 3. In the next lemma, T_1-configurations are consistent sub(hyper)graphs of the characteristic sequence of T, 2.6.

Lemma 5.14. *Suppose that we have two characteristic sequences, $(T, \varphi) \mapsto \langle P_n \rangle$ and $(T', \varphi') \mapsto \langle P'_n \rangle$. Suppose that every finite T'_1-configuration of $\langle P'_n \rangle$ can be isomorphically embedded into $\langle P_n \rangle$. Then any regular ultrafilter which realizes all small φ-types must also realize all small φ'-types.*

Proof. Let $A' \subset P'_1$ be a positive base set for φ', given with a distribution $d' : A' \to \mathcal{D}$. Essentially, the hypothesis allows us to transfer the blueprint of A' (index model by index model) over into the characteristic sequence of φ, where it will push forward by Łoś' theorem to a positive base set A for φ. A realization of this φ-type solves the combinatorial problem, and this solution can then be transferred back to solve A'.

More precisely, write $A'[t]$ for the projection to each index model given by d'. Define for each index $t < \lambda$ an isomorphic embedding $f[t]$ of $A'[t]$ into P_1. Then $f := \prod_{t<\lambda} f[t]/\mathcal{D}$ is an isomorphic embedding of A' into P_1; call its image A. Then by Łoś' theorem, A will be a positive base set (for φ), and as we chose each $f[t]$ to be an isomorphism, we can define a distribution $d : A \to \mathcal{D}$ by $d'(a') = d(f(a))$, for each $a' \in A'$.

Suppose the type correponding to the positive base set A is realized. Then by Observation 5.2 there is a distribution $d_1 : A \to \mathcal{D}$ which refines d and is a complete graph almost everywhere. Use d_1 and f to define d'_1 refining d', by setting $d'_1(a') = d_1(f(a'))$. As f was everywhere a graph isomorphism, we have successfully found a distribution of A' which is a.e. a complete graph; so the corresponding type must be realized. $\square$

Corollary 5.15. *If $(T, \varphi) \mapsto \langle P_n \rangle$ and this sequence is universal for finite T_0-configurations, then T is maximal in the Keisler order.*

Proof. By Lemma 5.14 in light of Theorem 1.33, which says that the Keisler order depends on an analysis of types in a finite language. $\square$

Bibliography

[1] Ax and Kochen, "Diophantine problems over local fields. I" American Journal of Mathematics 87 (1965), pp. 605-630

[2] Ax and Kochen, "Diophantine problems over local fields. II" American Journal of Mathematics 87 (1965), pp. 631-648

[3] Ax and Kochen, "Diophantine problems over local fields. III" Annals of Mathematics, Ser. 2 83 (1966), pp. 437-456

[4] Baldwin and Shelah, "Randomness and Semigenericity," Transactions AMS, 349 (1997) 1359-1376.

[5] Buechler, "Lascar Strong Types in Some Simple Theories," Journal of Symbolic Logic 64(2) (1999), 817-824.

[6] Casanovas and Kim, "A supersimple nonlow theory." Notre Dame J. Formal Logic 39 (1998), no. 4, 507–518.

[7] Chang and Keisler, *Model theory.* Third edition. Studies in Logic and the Foundations of Mathematics, 73. North-Holland Publishing Co., Amsterdam, 1990.

[8] Comfort and Negrepontis, *The theory of ultrafilters.* Die Grundlehren der mathematischen Wissenschaften, Band 211. Springer-Verlag, New York-Heidelberg, 1974.

[9] Džamonja and Shelah, "On $\triangleleft^*$-maximality," Annals of Pure and Applied Logic 125 (2004) 119-158.

[10] Elek and Szegedy, "Limits of Hypergraphs, Removal and Regularity Lemmas. A Non-standard Approach," (2007) arXiv:0705.2179.

[11] Gowers, "Hypergraph Regularity and the multidimensional Szemerédi Theorem." Ann. of Math. (2) 166 (2007), no. 3, 897–946.

[12] Keisler, "Ultraproducts which are not saturated." Journal of Symbolic Logic, 32 (1967) 23-46.

[13] Kim, *Simple first order theories*, Ph.D. thesis, University of Notre Dame, 1996.

[14] Kim and Pillay, “From stability to simplicity.” Bull. Symbolic Logic 4 (1998), no. 1, 17–36

[15] Komlós and Simonovits, “Szemerédi’s Regularity Lemma and its Applications in Graph Theory,” Combinatorics, Paul Erdös is Eighty, vol. 2, Budapest (1996) pps. 295-352.

[16] Kunen, “Ultrafilters and independent sets,” Trans. A.M.S. 172 (1972) 299-306.

[17] Hrushovski, “A stable $\aleph_0$-categorical pseudoplane,” preprint (1988).

[18] Macintyre, “Model theory: Geometrical and set-theoretic aspects and prospects,” Bulletin of Symbolic Logic, vol. 9, no 2 (2003), pps 197-212.

[19] Malliaris, “Realization of φ-types and Keisler’s order,” to appear, Annals of Pure and Applied Logic.

[20] Makowsky and Zilber, “Polynomial invariants of graphs and totally categorical theories,” MODNET Preprint, 2006.

[21] Rowbottom, “The Łoś conjecture for uncountable theories,” Notices. A. M. S. 11 (1964) 248.

[22] Shelah, “Stable theories.” Israel J Math 7 (1969) 187-202.

[23] Shelah, “Simple unstable theories.” Annals Math Logic 19 (1980) 177-203.

[24] Shelah, “On the cardinality of ultraproduct of finite sets.” J Symbolic Logic 35 (1970) 83-84.

[25] Shelah, “Saturation of ultrapowers and Keisler’s order,” Annals Math. Logic 4 (1972) 75-114.

[26] Shelah, *Classification Theory and the number of non-isomorphic models*. Studies in Logic and the Foundations of Mathematics, 92. First edition 1978, revised 1990. North-Holland Publishing Co., Amsterdam.

[27] Shelah, “Toward classifying unstable theories,” Annals of Pure and Applied Logic 80 (1996) 229-255.

[28] Shelah and Usvyatsov, “More on SOP_1 and SOP_2,” (2007), to appear, Annals of Pure and Applied Logic.

[29] Szemerédi, ”On Sets of Integers Containing No k Elements in Arithmetic Progression.” Acta Arith. 27, 199-245, 1975a.