ELLIPTIC CURVES AND COMPLEX MULTIPLICATION

JAKE ZWEIFLER

Abstract. In this paper, we will be discussing elliptic curves and complex multiplication. Namely, we will focus on the curve $y^2 = f(x) = x^3 + x$. Such curve has complex multiplication and thus, we can use this to show that $\mathbb{Q}(C[n])(i))/\mathbb{Q}(i)$ is an abelian Galois extension.

Contents

1. An Introduction to Elliptic Curves

An elementary problem in number theory is to find solutions, (x, y), to the quadratic equation $ax^2 + bxy + cy^2 + dx + ey + f = 0$, where the coefficients are rational numbers. Usually, we focus on rational or integer solutions, although we can find solutions over any field. Solving this problem is fairly trivial because we can project this curve onto a line. This is due to the fact that an arbitrary line usually crosses a conic two times.

Now, consider the problem of finding solutions to the cubic equation $f(x, y) = ay^3+bxy^2+cx^2y+dx^3+ex^2+fxy+gy^2+hx+iy+j = 0$, where the coefficients are rational numbers. We cannot solve this problem using the method above because an arbitrary line will usually cross a cubic three times. However, this allows us to view the set of points on the cubic as a group. This fact is discussed in great detail by Silverman and Tate [3, Section 1.2]. For the rest of this paper, we will use ideas motivated by these insights. We start with the following definition:

Definition 1.1. Let F be some field. For our purposes, an elliptic curve, C, is defined to be the set of points satisfying the equation

$$f(x, y) = ay^3 + bxy^2 + cx^2y + dx^3 + ex^2 + fxy + gy^2 + hx + iy + j = 0.$$

More specifically, we let $C(F)$ denote the set of points $(x, y) \in F^2$ satisfying $f(x, y) = 0$, along with the point $\mathcal{O}$, which we will define in the next section.

As will be discussed in more detail in the next section, we can define a group law on $C(F)$. As with any group, we ensure the existence of a zero element which we denote $\mathcal{O}$. We can then begin to look at points in $C(F)$ with finite order, i.e, points P such that $\underbrace{P+P+...+P}_{n \text{ times}} = nP = \mathcal{O}$. This will lead us to our penultimate goal of studying fields generated by the coordinates of these finite ordered points. Finally, we will discuss the field $\mathbb{Q}(C[n])(i)$, proving that it is Galois over $\mathbb{Q}(i)$. All these ideas will be defined and discussed in more detail in the following sections.

2. Developing A Group Law for $C(F)$

First, we define the following operation that, while useful, ultimately fails to satisfy the group axioms:

Definition 2.1. Let C be some elliptic curve and F a field. For $P, Q \in C(F)$ let $P * Q$ denote the third intersection point of C and the line connecting P and Q. If $P = Q$, then we say that $P * Q$ is the second point on the line tangent to P.

Immediately, problems arise with this definition since it is not always the case that the line though P and Q will intersect the curve a third time. However, viewing the curve in the projective plane, we can show that the line through P and Q indeed does always intersect C exactly once more so long as we include "extra points" at infinity [3, Appendix].

Having established that this operation is well defined, one can check via simple algebra and coordinate geometry that $P * Q$ is indeed in $C(F)$ when P and Q are in $C(F)$. Clearly, this operation is commutative since the order of P and Q does not affect the line through them. However, one can also find that this operation is not associative. It is also unclear what the zero element would be. Thus, we need to do a bit more to formulate a proper group law. For brevity purposes, we give the following definition:

Definition 2.2. We say a curve C is in Wierstrass normal form if

$$C : y^2 = f(x) = x^3 + ax^2 + bx + c$$

for $a, b, c, \in \mathbb{Q}$ and $f(x)$ separable, i.e, $f(x)$ has three distinct roots.

Although we can formulate the group law for any cubic curve, it is much easier to invoke more projective geometry to show that every elliptic curve can be transformed to be in Weierstrass normal form [3, Section 1.3]. Thus, from now on, we assume that C is in Wierstrass normal form.

The final thing we need is to define the zero element, denoted $\mathcal{O}$. For the sake of intuition, we can think of $\mathcal{O}$ as being the point $(0, \pm\infty)$, the point at infinity directly above the origin. While this definition is sufficient for our discussion, this notion is rigorized by projective geometry in Chapter 1 [3, Section 1.2]. With the addition of $\mathcal{O}$, we can now define the complete set of points on the curve by $C(F) = \{(x, y) \in F^2 \mid y^2 = f(x) = x^3 + ax^2 + bx + c\} \cup \{\mathcal{O}\}$. To provide more intuition of how $\mathcal{O}$ works, consider the following example:

Example 2.3. For $P, Q \in C$ with $P = (x, y)$ and $Q = (x, -y)$, it is clear that the line through P and Q is a vertical line. Since C is of the form $y^2 = f(x)$, we know that C is symmetric about the x-axis. Thus, the line intersecting P and Q only intersects the third point $\mathcal{O}$ at infinity. This is seen in Figure 2.1 below.

Furthermore, for general points $P, Q \in C$ the line intersecting P and Q will by definition intersect $P * Q$. Thus, the line through P and $P * Q$ intersects Q, so we can conclude that $P * (P * Q) = Q$. Thus, since $(x, y) * (x, -y) = \mathcal{O}$, we have that $(x, y) * \mathcal{O} = (x, -y)$. This also makes sense since from Figure 2.1 , it we draw a line through P and a point infinitely high up, the line will also go through Q.

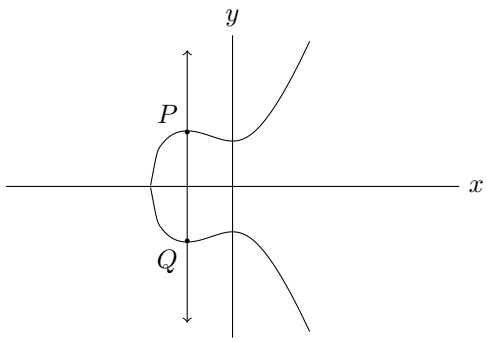

Figure 2.1

The final thing to note is that we define $\mathcal{O} * \mathcal{O} = \mathcal{O}$. Again, this can be more rigorously derived using projective geometry. We can now define the group law for C:

Definition 2.4. For $P, Q \in C(F)$, define the binary operation $+$ on the set $C(F)$ by $P + Q = (P * Q) * \mathcal{O}$.

We now must show that $C(F)$ with this operation is a group. We start with the following lemma:

Lemma 2.5 ([3, Section 1.2])**.** *Let C be an elliptic curve. The set $C(F)$ with the binary operation, $+$, defined in Definition 2.4, contains a zero element and additive inverses for each $P \in C(F)$.*

Proof. We know that $\mathcal{O} \in C(F)$ by how we defined $C(F)$ above. We must now actually show that $\mathcal{O}$ always acts as the zero element. If $P = \mathcal{O}$, we know that $\mathcal{O} + \mathcal{O} = (\mathcal{O} * \mathcal{O}) * \mathcal{O} = \mathcal{O} * \mathcal{O} = \mathcal{O}$ where we use the fact that $\mathcal{O} * \mathcal{O} = \mathcal{O}$. Thus, $\mathcal{O} + \mathcal{O} = \mathcal{O}$. For nonzero P we know that $P = (x, y)$. Then, we have that $P + \mathcal{O} = (P * \mathcal{O}) * \mathcal{O}$. As discussed in Example 2.3, $P * \mathcal{O} = (x, -y)$. Using this fact again, $(x, -y) * \mathcal{O} = (x, y)$. Thus, $P + \mathcal{O} = (x, y) = P$. So $\mathcal{O}$ acts as a zero element.

Now, we will show inverses exist. We just established that $\mathcal{O} + \mathcal{O} = \mathcal{O}$, so $-\mathcal{O} = \mathcal{O}$. As discussed in Example 2.3, when $P = (x, y)$ and $Q = (x, -y)$, we have that $P * Q = \mathcal{O}$. Thus, $P + Q = (P * Q) * \mathcal{O} = \mathcal{O} * \mathcal{O} = \mathcal{O}$. Therefore, $Q = -P$. Thus, when $P = (x, y)$, we have that $-P = (x, -y)$. Note that since $y \in F$ implies that $-y \in F$, we know that $-P \in C(F)$. Thus, for all $P \in C(F)$, P has an additive inverse also in $C(F)$. □

It would be helpful to have an explicit formula for $P + Q$ when both are nonzero. Say that $P = (x_1, y_1)$ and $Q = (x_2, y_2) \neq \pm P$. Then, let $\lambda = \frac{y_2 - y_1}{x_2 - x_1}$ and $\nu = y_2 - \lambda x_2 = y_1 - \lambda x_1$, noting that $x_2 - x_1 \neq 0$ since $Q \neq \pm P$. The line through P and Q is then $y = \lambda x + \nu$. We want to first find $P * Q = (x_3, y_3)$. Plugging the line into $y^2 = f(x) = x^3 + ax^2 + bx + c$, we get that $(\lambda x + \nu)^2 = x^3 + ax^2 + bx + c$. We

know that x_1, x_2 and x_3 are the three solutions to this equation. Thus, we should have that $x^3+ax^2+bx+c-(\lambda x+\nu)^2 = (x-x_1)(x-x_2)(x-x_3)$. Simplifying, we see that the coefficient of the x^2 term is $a-\lambda^2$ on the left-hand side and $-(x_1+x_2+x_3)$ on the right-hand side. So it must be the case that $x_1 + x_2 + x_3 = \lambda^2 - a$. Thus, $x_3 = \lambda^2 - a - x_1 - x_2$. We can plug this into the $y = \lambda x + \nu$ to find that $y_3 = \lambda(\lambda^2-a-x_1-x_2)+\nu$. For notation purposes, let $x(P)$ and $y(P)$ denote the x and y coordinates of the point P, respectively. So we have that $x(P+Q) = \lambda^2-a-x_1-x_2$ and $y(P + Q) = \lambda(\lambda^2 - a - x_1 - x_2) + \nu$.

This does not cover the case when $P = Q$. Say that $P = (x, y)$. Then, line tangent to C at P will have slope $\lambda = y' = \frac{f'(x)}{2y}$. Thus, using the equation above with $x_1 = x_2 = x$, we have that the $x(P + P) = \frac{f'(x)^2}{4y^2} - a - 2x = \frac{f'(x)^2}{4f(x)} - a - 2x$. Then, $y(P + P) = \lambda(\frac{f'(x)^2}{4f(x)} - a - 2x) + y - \lambda x$. Finally, when $P = -Q$, we know that $P + Q = \mathcal{O}$, so we don't require a formula.

With these formulae, we can now prove that $C(F)$ is indeed a group.

Theorem 2.6 ([3, Section 1.2])**.** *Let F be a field such that* $\mathrm{char}(F) \neq 2$*. Then set $C(F)$ with the binary operation defined in Definition 2.4 is an abelian group.*

Proof. By Lemma 2.5, we already know that $C(F)$ has a zero element, namely $\mathcal{O}$, and that every $P \in C(F)$ has an inverse.

Then, let $P = (x_1, y_1)$ and $Q = (x_2, y_2)$ be points in $C(F)$. So $x_1, x_2, y_1, y_2 \in F$. By the formulae developed above, we know that $x(P+Q)$ and $y(P+Q)$ are given by rational functions of x_1, x_2, y_1, y_2 and therefore must also be in F. One can check that if λ is ever undefined since the denominator is zero, it is necessarily the case that $P + Q = \mathcal{O}$ by previous definitions and observations. Thus, in either case, we have that $P + Q \in C(F)$. Therefore, $C(F)$ is closed under the addition operation.

Finally, it is clear from the formulae that $P + Q$ is always equal to $Q + P$. Although it takes a lot of algebraic manipulation, we can also use the formulae to show that $(P + Q) + R = P + (Q + R)$. One can also take a geometric approach to show this fact [3, Section 1.2]. So the addition operation is both associative and commutative, showing that $C(F)$ is an abelian group.

Note that we require $\mathrm{char}(F) \neq 2$ since we divide for 4 in the equation for $x(P + P)$.

□

Now that we have established that $C(F)$ is a group, we can begin to examine $C(F)$ for specific fields F. Also, note that our formulae for $P + Q$ are completely algebraic, despite the fact that $P + Q$ is based off a geometric concept. This is important because we can ensure that $P + Q$ is still well defined even when $F = \mathbb{C}$ or when $F = \mathbb{F}_q$, the finite field of q elements.

3. Points of Finite order

Definition 3.1. If C is the curve $y^2 = f(x) = x^3 + ax^2 + bx + c$, let $C[n]$ denote the points in C such that $\underbrace{P + P + ... + P}_{n \text{ times}} = nP = \mathcal{O}$. If $n \in \mathbb{N}$ is the smallest such number such that $nP = \mathcal{O}$ then we say that P has order n.

As with any group, the elements of order dividing n form a subgroup. Thus, $C[n]$ is a group in its own right. However, we can show more. First, let us examine $C[n]$ for small values of n:

Example 3.2. Let $n = 2$. Thus, $C[2]$ contains points P such that $P + P = \mathcal{O}$. Subtracting P from both sides, this is equivalent to $P = -P$. Clearly, $\mathcal{O} \in C[2]$ since $\mathcal{O} + \mathcal{O} = \mathcal{O}$. Now suppose that $P \neq \mathcal{O}$. So, we have that $P = (x, y) = -P = (x, -y)$. So $y(P)$ must equal 0. Therefore, x must be a solution to $f(x) = 0$. Since we have assumed that $f(x)$ is separable, there will be three distinct roots of $f(x)$ in $\mathbb{C}$. Call these roots α_1, α_2 and α_3. Therefore, over $\mathbb{C}$ we have that $C[2] = \{\mathcal{O}, (\alpha_1, 0), (\alpha_2, 0), (\alpha_3, 0)\}$.

With much complex analysis [2, Section 9.2], we can show that there is a group isomorphism between $C(\mathbb{C})$ and $\mathbb{C}/L$ where $L = \{a_1\omega_1 + a_2\omega_2 \mid a_1, a_2 \in \mathbb{Z}\}$, a 2-dimensional lattice in $\mathbb{C}$. Note that the *periods* ω_1 and ω_2 are determined by the coefficients of $f(x)$. With this, we can see that a point of order dividing n in $C(\mathbb{C})$ will correspond to $\frac{a_1\omega_1}{n} + \frac{a_2\omega_2}{n}$ in $\mathbb{C}/L$ for $a_1, a_2 \in \mathbb{Z}/n\mathbb{Z}$. Thus, this is a group of size n^2 and is clearly generated by $\frac{\omega_1}{n}$ and $\frac{\omega_2}{n}$. Therefore, looking at $C[n]$ over $\mathbb{C}$, there is a group isomorphism $C[n] \cong (\mathbb{Z}/n\mathbb{Z}) \times (\mathbb{Z}/n\mathbb{Z})$. It is quite miraculous that such a complicated group is isomorphic to the well-understood group, $(\mathbb{Z}/n\mathbb{Z}) \times (\mathbb{Z}/n\mathbb{Z})$. Importantly, we now know that $C[n]$ is a finite group with size n^2

4. Preliminary Galois Theory

We now digress to develop the Galois Theory we plan on using later. For more preliminary details about field theory and Galois theory, please refer to Sections 13 and 14 of Dummit and Foote [1]. We start be defining field extensions contained in $\mathbb{C}$:

Definition 4.1. We call K a field extension of F, written K/F, when K is a field that can be viewed as a vector space over F, a subfield. We then say that K/F is finite if the dimension of K over F is finite, i.e, there is a finite basis. Let $[K : F]$ denote the size of this basis, i.e, the dimension of K over F.

As with every finite dimensional vector space, we can find some finite basis, $b_1, ..., b_n$. We can then say that $K = F[b_1, ..., b_n]$. For any field F and any algebraic element α over K, Bézout's identity for coprime polynomials in $F[x]$ tells us that $F(\alpha) = F[\alpha]$. Using other techniques, we can prove the Primitive Element Theorem [4], that $F(\alpha_1, \alpha_2, ..., \alpha_n) = F(\gamma)$ for some $\gamma \in F(\alpha_1, \alpha_2, ..., \alpha_n)$. Thus, with these two tools, we know that for every finite extension K/F, there exists some $\alpha \in K$ such that $F(\alpha) = K$. Then, if m_α is the minimal polynomial of α over F, it is clear that $[K(\alpha) : F] = \deg(m_\alpha)$. These facts will be important when we discuss field embeddings.

From now on, we will narrow our focus to fields that are contained in $\mathbb{C}$.

Definition 4.2. Let K be a subfield of $\mathbb{C}$. We call $\sigma : K \hookrightarrow \mathbb{C}$ a field embedding of K into $\mathbb{C}$ if σ is a nontrivial ring homomorphism.

Notice that since K and $\mathbb{C}$ are fields, σ is necessarily injective as long as σ is not the trivial homomorphism. For reasons that will soon be clear, we would also like to further restrict these embeddings such that they fix F, i.e, $\sigma(a) = a$ for all $a \in F$. We can then say something about the number of these embeddings:

Theorem 4.3. *Let K/F be a finite extension contained in $\mathbb{C}$. Then, the number of distinct field embeddings of K into $\mathbb{C}$ that fix F is exactly $[K : F]$.*

Proof. As mentioned above, since K/F is a finite extension, we know that $K = F(\alpha)$ for some $\alpha \in K$. Thus, every $\mu \in K = F(\alpha)$ is equal to $a_0 + a_1\alpha + ... + a_{n-1}\alpha^{n-1}$ for unique $a_i \in F$ where $n = \deg(m_\alpha)$. We then know that $\sigma(\mu) = \sigma(a_0 + a_1\alpha + ... + a_{n-1}\alpha^{n-1}) = a_0 + a_1\sigma(\alpha) + ... + a_{n-1}\sigma(\alpha)^{n-1}$ since σ fixes F. Thus, σ is entirely determined by $\sigma(\alpha)$. For $f(x) \in F[x]$, it is easy to check that $\sigma(f(\beta)) = f(\sigma(\beta))$ as long as σ is a homomorphism that fixes F. Thus, $m_\alpha(\sigma(\alpha)) = \sigma(m_\alpha(\alpha)) = \sigma(0) = 0$, implying that $\sigma(\alpha)$ is also a zero of m_α. Since m_α is the minimal polynomial, it is necessarily irreducible. It is a general fact that irreducible polynomials in fields with characteristic zero have no double roots[1, Cor 13.34]. Therefore, we can conclude that there are exactly $n = \deg(m_\alpha)$ choices for $\sigma(\alpha)$, namely the n unique roots of m_α in $\mathbb{C}$. Thus, there are exactly $n = \deg(m_\alpha) = [K : F]$ field embeddings that fix F. □

It is sometimes the case that the codomain of σ is K, itself. This occurs when $\sigma(K) \subset K$. With this idea, we can now introduce the automorphism group, further restricting our choices for σ:

Definition 4.4. Let $\mathrm{Aut}(K/F) := \{\sigma \mid \sigma \text{ is a field embedding that fixes } F \text{ such that } \sigma(K) \subset K\}$. Then, $\mathrm{Aut}(K/\mathbb{Q})$ is a group where the group operation is composition, i.e $\sigma\tau := \sigma \circ \tau$.

First, note that $\sigma(K) \subset K$ really just implies that $\sigma(K) = K$. We can see that this is true since σ is injective and thus, $\sigma(K)$ must be a vector space over F with the same dimension as K/F, implying that $\sigma(K) = K$, i.e, σ is surjective. Thus, these homomorphisms are really automorphisms from K to itself, which is why we call the group an *automorphism* group. Since each σ is an isomorphism, we know that σ^{-1} also exists. Furthermore, it is easy to check that $\sigma \circ \tau$ is an automorphism when σ and τ are automorphisms. Thus, $\mathrm{Aut}(K/\mathbb{Q})$ is indeed a group. We can say the following about the size of this group:

Corollary 4.5. *For any finite extension K/F, we have that $|Aut(K/F)| \leq [K : F]$.*

Proof. It is clear than any element of $\mathrm{Aut}(K/F)$ is also a field embedding of K that fixes F. By Theorem 4.3, there are $[K : F]$ such embeddings. Thus, $|\mathrm{Aut}(K/F)|$ is at most $[K : F]$. □

More specifically, $|\mathrm{Aut}(K/F)|$ is equal to the number of roots of m_α that are in K. Something very nice occurs when all of the roots of m_α are in K, implying that every field embedding is in the automorphism group. This allows us to make the following definition:

Definition 4.6. We say that the field extension K/F is Galois when $|\mathrm{Aut}(K/F)| = [K : F]$. In this case, we write $\mathrm{Gal}(K/F)$ instead of $\mathrm{Aut}(K/F)$.

To summarize, based on our discussions above, we know that a Galois extension K/F will also have that $\sigma(K) = K$ for every field embedding σ. Since $\sigma(\alpha)$ must be a root of m_α, this also means that m_α will split completely in K, implying that $|\mathrm{Aut}(K/F)| = [K : F] = \deg(m_\alpha)$ which is equal to the total number of field embeddings. Thus, if we show that the size of the automorphism group is equal to any of these values or that $\sigma(K) = K$ for each embedding, we will have shown that K/F is Galois.

Example 4.7. The field $K = \mathbb{Q}(\sqrt{d})$ is Galois over $\mathbb{Q}$ for $d \in \mathbb{Z}$, d nonsquare. Note that the minimal polynomial of $\sqrt{d}$ is $m(x) = x^2 - d$. Since this polynomial splits in K (with roots $\pm\sqrt{d}$), we would expect the extension to be Galois. The two automorphisms are the identity, and the map that sends $\sqrt{d}$ to its conjugate root, $-\sqrt{d}$. Thus, $2 = \deg(x^2 - d) = |\operatorname{Aut}(\mathbb{Q}(\sqrt{d})/\mathbb{Q})| = [\mathbb{Q}(\sqrt{d}) : \mathbb{Q})]$. So $K/\mathbb{Q}$ is Galois.

Example 4.8. The field $K = \mathbb{Q}(\sqrt[3]{2})$ is not Galois over $\mathbb{Q}$ since there are elements missing in K. The three embeddings of K into $\mathbb{C}$ send $\sqrt[3]{2}$ to $\sqrt[3]{2}$, $\sqrt[3]{2}\omega$ and $\sqrt[3]{2}\omega^2$ in $\mathbb{C}$ where ω is the cube root of unity. However, $\omega \notin K \subset \mathbb{R}$. Thus, $\operatorname{Aut}(K/\mathbb{Q}) = \{\mathrm{id}\}$.

5. Looking at $C(K)$ when K is generated by points $P \in C$ of finite order

First, we will prove some results about $C(K)$ when K is any Galois extension of $\mathbb{Q}$. More specifically, we can study how $\mathrm{Gal}(K/\mathbb{Q})$ acts on $C(K)$ for some curve $C : y^2 = f(x) = x^3 + ax^2 + bx + c$. For $\sigma \in \mathrm{Gal}(K/\mathbb{Q})$ and $P \in C(K)$, define $\sigma(P)$ as follows:

$$\sigma(P) = \begin{cases} (\sigma(x), \sigma(y)) & \text{if } P = (x, y) \\ \sigma(P) = \mathcal{O} & \text{if } P = \mathcal{O} \end{cases}$$

With this definition, $\mathrm{Gal}(K/\mathbb{Q})$ acts very nicely on $C(K)$:

Theorem 5.1 ([3, Proposition 6.3])**.** *For all $\sigma \in \mathrm{Gal}(K/\mathbb{Q})$ and $P \in C(K)$, $\sigma(P) \in C(K)$. Additionally, $\sigma(P + Q) = \sigma(P) + \sigma(Q)$ and $\sigma(-P) = -\sigma(P)$.*

Proof. First, notice that when $P = \mathcal{O}$, $\sigma(P) = \mathcal{O} \in C(K)$. When $P = (x, y) \in C(K)$, we have that $y^2 = x^3 + ax^2 + bx + c$. Applying σ to both sides, we get that $\sigma(y^2) = \sigma(x^3 + ax^2 + bx + c)$. Since σ is a homomorphism that fixes $\mathbb{Q}$, we have $\sigma(y)^2 = \sigma(x)^3 + a\sigma(x)^2 + b\sigma(x) + c$. Since $K/\mathbb{Q}$ is Galois and $x, y \in K$, we know that $\sigma(x), \sigma(y) \in K$ as well. Therefore, since $\sigma(P)$ satisfies $y^2 = f(x)$ and has coordinates in K, we can conclude that $\sigma(P) \in C(K)$.

To check that $\sigma(P + Q) = \sigma(P) + \sigma(Q)$, we would have to check many cases, depending on the relationship between P and Q. We will check the most general case, when $P = (x_1, y_1)$ and $Q = (x_2, y_2) \neq \pm P$. From Section 2, we know that $x(P+Q) = (\frac{y_2 - y_1}{x_2 - x_1})^2 - a - x_1 - x_2$. As before, since σ is a homomorphism that fixes $\mathbb{Q}$, $\sigma(x(P + Q)) = \sigma((\frac{y_2 - y_1}{x_2 - x_1})^2 - a - x_1 - x_2) = (\frac{\sigma(y_2) - \sigma(y_1)}{\sigma x_2 - \sigma(x_1)})^2 - a - \sigma(x_1) - \sigma(x_2)$. Note that this is the same as $x(\sigma(P) + \sigma(Q))$. Similar logic works to show that the same is true for the y coordinates, implying that $\sigma(P + Q) = \sigma(P) + \sigma(Q)$. We can do something similar for the cases where P and Q are more related, for example when $P = \pm Q$.

Finally, since $-P = (x, -y)$, we have that $\sigma(-P) = (\sigma(x), \sigma(-y))$. Since $\sigma(-y) = -\sigma(y)$, we can conclude that $\sigma(-P) = (\sigma(x), -\sigma(y)) = -\sigma(P)$, thus completing the proof. □

Corollary 5.2 ([3, Proposition 6.3])**.** *For $\sigma \in \mathrm{Gal}(K/\mathbb{Q})$ and $P \in C(K)$, if $n \in \mathbb{Z}$ then $\sigma(nP) = n\sigma(P)$.*

Proof. First, notice that when $n = 0$, both sides are zero so the statement is true. For $n \geq 1$, we can iteratively apply the fact that $\sigma(P + P) = \sigma(P) + \sigma(P)$ from

Theorem 5.1 until we have that $\sigma(nP) = n\sigma(P)$. Then, since $\sigma(-P) = -\sigma(P)$, we can show that the statement is true for negative n as well. □

Thus, with Theorem 5.1 and Corollary 5.2, we can conclude that σ is a group homomorphism on $C(K)$. We will also want to prove the following fact about the order of $\sigma(P)$:

Theorem 5.3 ([3, Proposition 6.3])**.** *Let $\sigma \in \mathrm{Gal}(K/\mathbb{Q})$. If the order of P is n, then the order of $\sigma(P)$ is also n.*

Proof. Suppose that the order of P is n. This implies that $nP = \mathcal{O}$. Then, $\sigma(nP) = \sigma(\mathcal{O}) = \mathcal{O}$. By Corollary 4.2, $\sigma(nP) = n\sigma(P)$. So $n\sigma(P) = \mathcal{O}$. Let m be the order of $\sigma(P)$. Thus, $m \mid n$. Since $\mathrm{Gal}(K/\mathbb{Q})$ is a group, we can consider σ^{-1}. We know that $m\sigma(P) = \mathcal{O}$. Using Theorem 4.1 again, we have that $\sigma(mP) = \mathcal{O}$. Applying σ^{-1} to both sides, $mP = \sigma^{-1}(\mathcal{O}) = \mathcal{O}$. Thus, we must have that $n \mid m$. Therefore, $m = n$ and P and $\sigma(P)$ have the same order. □

Now, we want to construct some field that is related to $C[n]$. In Section 3, we proved that $C[n] \cong (\mathbb{Z}/n\mathbb{Z}) \times (\mathbb{Z}/n\mathbb{Z})$. This tells us that $|C[n]| = n^2$. Explicitly, we have that $C[n] = \{\mathcal{O}, (x_1, y_1), ..., (x_{n^2-1}, y_{n^2-1})\}$ for $x_i, y_i \in \mathbb{C}$. This motivates the following:

Definition 5.4. For $C[n] = \{\mathcal{O}, (x_1, y_1), ..., (x_{n^2-1}, y_{n^2-1})\}$, define $\mathbb{Q}(C[n]) := \mathbb{Q}(x_1, y_1, ..., x_{n^2-1}, y_{n^2-1})$, a field extension generated by the coordinates of the nonzero points in $C[n]$.

Theorem 5.5 ([3, Proposition 6.5])**.** *If (x_i, y_i) are the nonzero points in $C[n]$ for $1 \leq i \leq n^2 - 1$, then the field $\mathbb{Q}(C[n]) = \mathbb{Q}(x_1, y_1, ..., x_{n-1}, y_{n-1})$ has only finitely many embeddings into $\mathbb{C}$. Furthermore, this field is Galois over $\mathbb{Q}$*

Proof. Let $\sigma : \mathbb{Q}(C[n]) \to \mathbb{C}$ be an embedding that fixes $\mathbb{Q}$. Every element of $\mathbb{Q}(C[n])$ is a linear combination of products of the x_i, y_i so σ is completely determined by where it sends x_i and y_i. Let $P_i = (x_i, y_i)$ be some nonzero point in $C[n]$ with order $m \mid n$. From the first part of our proof of Theorem 5.3, we know that the order of $\sigma(P_i)$ divides m and thus, divides n. Therefore, $n\sigma(P_i) = \mathcal{O}$ and so $\sigma(P_i)$ is just one of the points in $C[n]$. So we know that $\sigma(x_i) = x_j$ and $\sigma(y_i) = y_j$ for some $1 \leq j \leq n^2 - 1$. This is true for all $1 \leq i \leq n^2 - 1$, so we know that there are at most $(n^2 - 1)^2$ possibilities for where σ sends the x_i and y_i. Therefore, there are only finitely many embeddings of $\mathbb{Q}(C[n])$ in $\mathbb{C}$.

As mentioned, every element $\mu \in \mathbb{Q}(C[n])$ is a linear combination of products of x_i and y_i. Since $\sigma(x_i) = x_j \in \mathbb{Q}(C[n])$ and $\sigma(y_i) = y_j \in \mathbb{Q}(C[n])$ for all $1 \leq i \leq n^2 - 1$, we know that $\sigma(\mu)$ will also be in $\mathbb{Q}(C[n])$. Thus, $\sigma(\mathbb{Q}(C[n])) \subset \mathbb{Q}(C[n])$. As mentioned in Section 4, since σ is necessarily injective, this implies that $\sigma(\mathbb{Q}(C[n])) = \mathbb{Q}(C[n])$ for every embedding σ. Thus, $\mathbb{Q}(C[n])$ is Galois over $\mathbb{Q}$. □

It will be helpful to describe the Galois group of $\mathbb{Q}(C[n])/\mathbb{Q}$ and we can do so using the fact that $C[n] \cong (\mathbb{Z}/n\mathbb{Z}) \times (\mathbb{Z}/n\mathbb{Z})$. Because of this group isomorphism, we can view $C[n]$ as a $\mathbb{Z}/n\mathbb{Z}$-module with rank 2. While not every R-module has a basis, $C[n]$ certainly does. An example of a basis would be the elements that get mapped to $\frac{\omega_1}{n}$ and $\frac{\omega_2}{n}$ in $\mathbb{C}/L$. Thus, for every $P \in C[n]$, there exist $a, b \in \mathbb{Z}/n\mathbb{Z}$ such that $P = aP_1 + bP_2$ where $\{P_1, P_2\}$ is some basis over $\mathbb{Z}/n\mathbb{Z}$. We can then use Theorem 5.1 and Corollary 5.2 with $K = \mathbb{Q}(C[n])$ to say that

$\sigma(P) = \sigma(aP_1 + bP_2) = a\sigma(P_1) + b\sigma(P_2)$. Note that we have used the fact that $C(K)$ contains $C[n]$.

So we have shown that σ as an action on $C[n]$ only depends on $\sigma(P_1)$ and $\sigma(P_2)$. As mentioned in the proof of Theorem 5.5, for $P \in C[n]$ and $\sigma \in \mathrm{Gal}(\mathbb{Q}(C[n])/\mathbb{Q})$, we know that $\sigma(P) \in C[n]$. Thus, we know that $\sigma(P_1) = \alpha_\sigma P_1 + \gamma_\sigma P_2$ and $\sigma(P_2) = \beta_\sigma P_1 + \delta_\sigma P_2$ where $\alpha_\sigma, \beta_\sigma, \gamma_\sigma, \delta_\sigma \in \mathbb{Z}/n\mathbb{Z}$. It is easy to see the connection to linear algebra: we can write $M_\sigma = \begin{pmatrix} \alpha_\sigma & \beta_\sigma \\ \gamma_\sigma & \delta_\sigma \end{pmatrix}$ and then notice that $\sigma(P) = (P_1, P_2) \cdot M_\sigma \cdot \begin{pmatrix} a \\ b \end{pmatrix}$ when $P = aP_1 + bP_2$.

We can then check that $M_\sigma M_\tau = M_{\sigma\tau}$. Also, since σ has an inverse, we know that M_σ must be invertible and therefore in $\mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$, the multiplicative group of invertible 2-by-2 matrices with entries in $\mathbb{Z}/n\mathbb{Z}$. Finally, suppose that M_σ is the identity matrix. This implies that $\sigma(P) = P$ for all $P \in C[n]$. So σ fixes every generating element of $\mathbb{Q}(C[n])$, implying that σ fixes the entire field. Thus, σ is the identity of $\mathrm{Gal}(\mathbb{Q}(C[n])/\mathbb{Q})$. With this fact, we can show that $M_\sigma = M_\tau$ implies that $\sigma = \tau$. Therefore, there is a one-to-one correspondence between elements of $\mathrm{Gal}(\mathbb{Q}(C[n])/\mathbb{Q})$ and elements of $\mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$. Since $M_\sigma M_\tau = M_{\sigma\tau}$, we can conclude that $\mathrm{Gal}(\mathbb{Q}(C[n])/\mathbb{Q})$ is isomorphic to a subgroup of $\mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$.

6. Complex Multiplication

The last topic we must introduce is that of complex multiplication. For the curve $C(\mathbb{C})$, the multiplication-by-n map is an endomorphism. This is clearly the case since $C(\mathbb{C})$ is abelian so $nP + nQ = n(P + Q)$. Note that the kernel of this map is $C[n]$. In most cases, the only endomorphisms that exist for $C(\mathbb{C})$ are the multiplication-by-n maps. However, for some curves, there are other examples of endomorphisms.

Definition 6.1. Let C be an elliptic curve. If there exists some endomorphism $\phi : C(\mathbb{C}) \to C(\mathbb{C})$ that is not the multiplication-by-n map, we say that the curve C has complex multiplication.

We can view the set of endomorphisms as a ring where addition is defined such that $(\phi_1 + \phi_2)(P) = \phi_1(P) + \phi_2(P)$ and multiplication is defined as composition, i.e, $\phi_1\phi_2(P) = \phi_1(\phi_2(P))$. It is then clear that if C lacks complex multiplication, this ring is isomorphic to $\mathbb{Z}$.

If C does have complex multiplication, we can prove some things about the extra endomorphisms. First, recall that $C(\mathbb{C}) \cong \mathbb{C}/L$. Let f be an endomorphism from $\mathbb{C}/L$ to $\mathbb{C}/L$. Using more complex analysis, we can show that $f(z) = cz$ for some $c \in \mathbb{C}$ [3, Prop 6.17]. Furthermore, we can show that if c is real, then $c \in \mathbb{Z}$. Thus, in order for ϕ to not be the multiplication-by-n map, $c \in \mathbb{C} \setminus \mathbb{R}$ [3, Prop 6.18]. Thus, the ring of isogenies is isomorphic to some subring of $\mathbb{C}$, hence we say C has "complex" multiplication.

Remark 6.2. *For any homomorphism $h : C(\mathbb{C}) \to C(\mathbb{C})$, we can use similar arguments as in Theorem 5.1 and Corollary 5.2 to show that $h(P+Q) = h(P)+h(Q)$ and that $h(nP) = nh(P)$. Additionally, if h has an inverse, Theorem 4.3 is also true for h, that is, if P has order n, then $h(P)$ has order n. Thus, these facts apply to all endomorphisms as well.*

Example 6.3. Consider the curve $C : y^2 = x^3 + x$. Then, the map $\phi(x, y) = (-x, iy)$ is an endomorphism. We can check that $\phi(P) \in C$ when $P \in C$: $(iy)^2 = -y^2 = -(x^3 + x) = (-x)^3 + (-x)$. With casework and algebra, we can also check that $\phi(P + Q) = \phi(P) + \phi(Q)$, thus verifying that ϕ is indeed a homomorphism when considered as a function from $C(\mathbb{C})$ to $C(\mathbb{C})$. Also note that $\phi \circ \phi(P) = -P$. Thus, $\phi^2 = -1$, so ϕ is the element that is analogous to i. Thus, the ring of endomorphisms is isomorphic to $\mathbb{Z}[i]$, the Gaussian integers.

7. Abelian Galois Extensions of $\mathbb{Q}(i)$

Continuing from Example 6.3, we consider the curve $C : y^2 = x^3 + x$ with complex multiplication $\phi(x, y) = (-x, iy)$. Let $K/\mathbb{Q}$ be any Galois extension with $i \in K$. Then, for $P \in C(K)$ and $\sigma \in \mathrm{Gal}(K/\mathbb{Q})$, we have that $\sigma(P) \in C(K)$. We also have that $\phi(P) \in C(K)$. For reasons that will be clear soon, we want σ and ϕ to commute. Thus, we should have that $\sigma(\phi(P)) = \phi(\sigma(P))$ for all $P \in C(K)$. More specifically, $\sigma(\phi(P)) = \sigma(-x, iy) = (-\sigma(x), \sigma(i)\sigma(y))$ and $\phi(\sigma(P)) = \phi(\sigma(x), \sigma(y)) = (-\sigma(x), i\sigma(y))$. So ϕ and σ commute when $\sigma(i) = i$. Thus, we want to consider all homomorphisms σ that fix $\mathbb{Q}(i)$ not just $\mathbb{Q}$. Therefore, we must also adjoin i to $\mathbb{Q}(C[n])$, motivating the following definition:

Definition 7.1. For $n \geq 1$ and the curve $C : y^2 = x^3 + x$, define $K_n := \mathbb{Q}(C[n])(i)$.

We can then show that the extension $K_n/\mathbb{Q}(i)$ is Galois. Even more spectacular, we can show that the Galois group of $K_n/\mathbb{Q}(i)$ is abelian. We start by proving the first claim:

Theorem 7.2 ([3, Theorem 6.19])**.** *For K_n, as defined above, we have that K_n is a Galois extension of $\mathbb{Q}(i)$.*

Proof. In Theorem 5.5, we showed that $\mathbb{Q}(C[n])$ is Galois over $\mathbb{Q}$. Also, as discussed in Example 4.7 with $d = -1$, $\mathbb{Q}(i)$ is Galois over $\mathbb{Q}$. Thus, by general Galois theory [1, Prop 14.21], the compositum of these two fields, $\mathbb{Q}(C[n])(i) = K_n$ is Galois over $\mathbb{Q}$. Consider an arbitrary embedding $\sigma : K_n \to \mathbb{C}$ that fixes $\mathbb{Q}(i)$. This embedding automatically fixes $\mathbb{Q}$. Thus, $\sigma \in \mathrm{Gal}(K_n/\mathbb{Q})$, implying that $\sigma(K_n) = K_n$. Thus, we know that K_n is Galois over $\mathbb{Q}(i)$, as well. □

Now, we want to show that $\mathrm{Gal}(K_n/\mathbb{Q}(i))$ is an abelian group. As we did in Section 5, if we fix generators $P_1, P_2 \in C[n]$, we get that $\sigma(P_1) = \alpha_\sigma P_1 + \gamma_\sigma P_2$ and $\sigma(P_2) = \beta_\sigma P_1 + \delta_\sigma P_2$ where $\alpha_\sigma, \beta_\sigma, \gamma_\sigma, \delta_\sigma \in \mathbb{Z}/n\mathbb{Z}$. Thus, we can associate the matrix $M_\sigma = \begin{pmatrix} \alpha_\sigma & \beta_\sigma \\ \gamma_\sigma & \delta_\sigma \end{pmatrix}$ to σ. Because of points made in Remark 6.2, this actually works for any endomorphism, not just the Galois elements. Thus, there is also a matrix $M_\phi = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ that corresponds to ϕ. Furthermore, the whole point of focusing on $K_n/\mathbb{Q}(i)$ was to ensure that ϕ and σ commute for all $\sigma \in \mathrm{Gal}(K_n/\mathbb{Q}(i))$. Thus, $M_\sigma M_\phi = M_\phi M_\sigma$. Now comes the crucial part of the proof that the Galois group is abelian. We will first show that M_ϕ is a special type of matrix. We will then show that all invertible matrices which commute with this special type of matrix commute with each other, proving that $\mathrm{Gal}(K_n/\mathbb{Q}(i))$ is abelian. We show these facts in the following lemmas:

Lemma 7.3 ([3, Lemma 6.20])**.** *The matrix M_ϕ is invertible. Furthermore, when reduced modulo q for any prime q dividing n, the matrix is not a scalar matrix, i.e.,*

$M_\phi \not\equiv cI \pmod q$ *for any* $c \in \mathbb{Z}/q\mathbb{Z}$ *so long as* $q \mid n$. *Note that* I *is the identity matrix.*

Proof. First, note that $\phi(\phi(P)) = -P$. Thus, we have that $M_\phi^2 = -I$, implying that $\det(M_\phi^2) = \det(-I) = 1$. Thus, $\det(M_\phi)^2 = 1$, implying that $\det(M_\phi)$ is a unit in $\mathbb{Z}/n\mathbb{Z}$. By general linear algebra, this tells us that M_ϕ is invertible. Explicitly, this inverse is M_{ϕ^3} where $\phi^3 = \phi \circ \phi \circ \phi$, another endomorphism.

Let q be a prime dividing n and suppose that $M_\phi \equiv \begin{pmatrix} m & 0 \\ 0 & m \end{pmatrix} \equiv mI \pmod q$ for some $m \in \mathbb{Z}/q\mathbb{Z}$. We can examine how ϕ acts on the points in $C[q] \subset C[n]$. Since $M_\phi \equiv mI \pmod q$, we can see that for all $P \in C[q]$, we have that $\phi(P) = mP$. We will show that this leads to a contradiction.

Consider $P \neq \mathcal{O}$ in $C[q]$. So the order of P is q. Let $\tau : K_n \to K_n$ be complex conjugation restricted to K_n and let $P' = \tau(P)$. Note that $\tau \in \mathrm{Gal}(K_n/\mathbb{Q})$. Then, by Theorem 5.3, the order of P' is also q. Furthermore, by Corollary 5.2, we know that $\tau(mP') = m\tau(P')$. Thus, $\tau(\phi(P')) = \phi(\tau(P'))$. We also know that $\tau(\phi(P')) = \tau(-x, iy) = (-\tau(x), -i\tau(y)) = -(-\tau(x), i\tau(y)) = -\phi(\tau(P'))$. Thus, we have that $\phi(\tau(P')) = -\phi(\tau(P'))$, implying that $2\phi(\tau(P')) = \mathcal{O}$. Since $\tau^2 = \mathrm{id}$, it is also clear that $\tau(P')$ equals P. Thus, $2\phi(P) = 2mP = \mathcal{O}$. Since P was arbitrary, we know that every point of order q satisfies this property.

Then, by the properties of orders, we know that $q \mid 2m$. Suppose that $q \mid m$. This would imply that $\phi(P) = mP = \mathcal{O}$ for all $P \in C[q]$. However, we know that $\phi^2(P) = -P$, so this clearly is not the case when $P \neq \mathcal{O}$. Therefore, we know that $q \nmid m$. So we must have that $q \mid 2$, implying that $q = 2$. However, we can explicitly show that this is not the case by calculating M_ϕ for $n = 2$.

From Example 3.2, we know that the points of order 2 are those satisfying $y^2 = f(x) = x^3 + x = 0$. Thus, for this curve, $C[2] = \{\mathcal{O}, (0,0), (i,0), (-i,0)\}$. We can let $P_1 = (0,0)$ and $P_2 = (i,0)$. Using the addition formula from Section 1, we can calculate that $P_1 + P_2 = (-i, 0)$. Thus, $\phi(P_1) = (-i,0) = P_1 + P_2$ and $\phi(P_2) = (0,0) = P_2$. Thus, we can conclude that $M_\phi = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$. This is clearly not a scalar matrix, so we know that M_ϕ is never equal to a scalar matrix when reduced modulo any prime $q \mid n$. □

We will now prove a series of claims about matrices in $\mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$:

Lemma 7.4 ([3, Lemma 6.21])**.** *Let* $A \in \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$ *be a matrix that is not a scalar matrix modulo* q *for any prime* q *dividing* n. *Then, there exists some* $T \in \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$ *such that* $T^{-1}AT = \begin{pmatrix} 0 & b' \\ 1 & d' \end{pmatrix}$ *for some* $b', d' \in \mathbb{Z}/n\mathbb{Z}$.

Proof. We prove that such a T exists by first showing that for every prime power q^k dividing n, there exists some $T_q \in \mathrm{GL}_2(\mathbb{Z}/q^k\mathbb{Z})$ such that $T_q^{-1}AT_q \equiv \begin{pmatrix} 0 & b' \\ 1 & d' \end{pmatrix} \pmod{q^k}$.

So let q^k be some prime power dividing n. Let $\overline{A} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{GL}_2(\mathbb{Z}/q^k\mathbb{Z})$ such that $A \equiv \overline{A} \pmod{q^k}$. We now will show that we can find some change of basis matrix T_q such that $T_q^{-1}\overline{A}T_q = \begin{pmatrix} 0 & b' \\ 1 & d' \end{pmatrix}$. This is known as putting $\overline{A}$ into

rational normal form. To do this, consider a basis of the form $\{\mathbf{v}, \overline{A}\mathbf{v}\}$. If we let the columns of T_q be these column vectors, it is clear that $\overline{A}T_q = T_q \begin{pmatrix} 0 & b' \\ 1 & d' \end{pmatrix}$ for proper b', d'. We still need to show that T_q has an inverse by choosing a correct $\mathbf{v}$.

By our supposition, we know that $\overline{A}$ is not a scalar matrix modulo q. Thus, one of the following is true:

(1) $b \not\equiv 0 \pmod{q}$
(2) $c \not\equiv 0 \pmod{q}$
(3) $a \not\equiv d \pmod{q}$

Each of these possibilities matches up with the following cases:

(1) If $b \not\equiv 0 \pmod{q}$, let $\mathbf{v} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$. Then $T_q = \begin{pmatrix} 0 & b \\ 1 & d \end{pmatrix}$
(2) If $c \not\equiv 0 \pmod{q}$, let $\mathbf{v} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$. Then $T_q = \begin{pmatrix} 1 & a \\ 0 & c \end{pmatrix}$
(3) If $b \equiv c \equiv 0$ and $a \not\equiv d \pmod{q}$, let $\mathbf{v} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$. Then $T_q = \begin{pmatrix} 1 & a+c \\ 1 & b+d \end{pmatrix}$

In each of these cases, once can check that the determinant of T_q is nonzero modulo q, and thus is a unit in $\mathbb{Z}/q^k\mathbb{Z}$, implying that T_q^{-1} exists, so we can say that $T_q^{-1}AT_q \equiv \begin{pmatrix} 0 & b' \\ 1 & d' \end{pmatrix} \pmod{q^k}$.

Now, by the Chinese Remainder Theorem, we can find T such that $T \equiv T_q \pmod{q^k}$ for each q dividing n. Similarly, let $T' \equiv T_q^{-1} \pmod{q^k}$ for each q dividing n. Then, it is clear that $TT' \equiv I \pmod{q^k}$ for each q. By Chinese Remainder Theorem, this implies that $TT' = I \pmod{n}$. So $T' = T^{-1}$, implying that $T \in \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$. Then, since $T_q^{-1}AT_q \equiv \begin{pmatrix} 0 & b' \\ 1 & d' \end{pmatrix} \pmod{q^k}$ for each q, we can use similar logic with the Chinese Remainder Theorem to say that $T^{-1}AT \equiv \begin{pmatrix} 0 & b' \\ 1 & d' \end{pmatrix} \pmod{n}$. Thus, when looking at A as a matrix in $\mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$, we can say that there exists some T such that $T^{-1}AT = \begin{pmatrix} 0 & b' \\ 1 & d' \end{pmatrix}$.

□

Lemma 7.5 ([3, Lemma 6.20])**.** *If B and B' both commute with $A = \begin{pmatrix} 0 & b \\ 1 & d \end{pmatrix}$, then B and B' commute with each other, i.e, $BB' = B'B$.*

Proof. Let $B = \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix}$ be some element such that $AB = BA$. This means we must have that $\begin{pmatrix} b\gamma & b\delta \\ \alpha + d\gamma & \beta + d\delta \end{pmatrix} = \begin{pmatrix} \beta & b\alpha + d\beta \\ \delta & b\gamma + d\delta \end{pmatrix}$. By simple algebra, it is clear that this is equivalent to $\beta = b\gamma$ and $\delta = \alpha + d\gamma$. Therefore, every B that commutes with A will be of the form $B = \begin{pmatrix} \alpha & b\gamma \\ \gamma & \alpha + d\gamma \end{pmatrix}$ for fixed b and d. Thus, it only neccesary to check that $BB' = B'B$ for B, B' of this form. In other words, we want to show that the following always holds:

$$\begin{pmatrix} \alpha & b\gamma \\ \gamma & \alpha + d\gamma \end{pmatrix} \begin{pmatrix} \alpha' & b\gamma' \\ \gamma' & \alpha' + d\gamma' \end{pmatrix} = \begin{pmatrix} \alpha' & b\gamma' \\ \gamma' & \alpha' + d\gamma' \end{pmatrix} \begin{pmatrix} \alpha & b\gamma \\ \gamma & \alpha + d\gamma \end{pmatrix}$$

By multiplying out these two products, we can see algebraically that the two sides are indeed equal and that $BB' = B'B$ so long as they both commute with A. □

Corollary 7.6 ([3, Lemma 6.20]). *For $A \in \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$, suppose that A is similar to $\begin{pmatrix} 0 & b \\ 1 & d \end{pmatrix}$, i.e, there exists some T such that $T^{-1}AT = \begin{pmatrix} 0 & b \\ 1 & d \end{pmatrix}$. Then, the set $\{B \in \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z}) \mid AB = BA\}$ is an abelian subgroup of $\mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$.*

Proof. First, note that $AB = BA$ implies that $B^{-1}(AB)B^{-1} = B^{-1}(BA)B^{-1}$, implying that $B^{-1}A = AB^{-1}$. Thus, we know that B^{-1} is in the set. We can do similar manipulations to show that the above set indeed closed. Clearly, I is in the set. Thus, it is a subgroup of $\mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z})$. Now, we must show that this subgroup is abelian.

Consider $B, B' \in \{B \in \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z}) \mid AB = BA\}$. So $AB = BA$. Since $TT^{-1} = I$, this is the same as $(T^{-1}AT)(T^{-1}BT) = (T^{-1}BT)(T^{-1}AT)$. Thus, $(T^{-1}BT)$ commutes with $T^{-1}AT = \begin{pmatrix} 0 & b \\ 1 & d \end{pmatrix}$. Similarly, so does $T^{-1}B'T$. Thus, by Lemma 7.5, $T^{-1}BT$ and $T^{-1}B'T$ commute with each other. So we have that $(T^{-1}BT)(T^{-1}B'T) = (T^{-1}B'T)(T^{-1}BT)$. Simplifying, we get that $BB' = B'B$, showing that the subgroup $\{B \in \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z}) \mid AB = BA\}$ is indeed abelian. □

We are now ready to prove our ultimate claim:

Theorem 7.7 ([3, Theorem 6.19]). *The Galois group $\mathrm{Gal}(K_n/\mathbb{Q}(i))$ is abelian.*

Proof. In Lemma 7.3, we showed that M_ϕ is not a scalar matrix modulo q for any prime q dividing n, so by Lemma 7.4, M_ϕ is able to be put into rational normal form. Therefore, by Corollary 7.6, the set $\{B \in \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z}) \mid M_\phi B = BM_\phi\}$ is an abelian group.

As we have discussed, by focusing on K_n over $\mathbb{Q}(i)$, we have ensured that $M_\phi M_\sigma = M_\sigma M_\phi$ for all $\sigma \in \mathrm{Gal}(K_n/\mathbb{Q}(i))$. Therefore, $\{M_\sigma \mid \sigma \in \mathrm{Gal}(K_n/\mathbb{Q}(i))\}$ can be identified as a subgroup of $\{B \in \mathrm{GL}_2(\mathbb{Z}/n\mathbb{Z}) \mid M_\phi B = BM_\phi\}$. This means that the elements of $\{M_\sigma \mid \sigma \in \mathrm{Gal}(K_n/\mathbb{Q}(i))\}$ commute with each other. We also know that $\mathrm{Gal}(K_n/\mathbb{Q}(i)) \cong \{M_\sigma \mid \sigma \in \mathrm{Gal}(K_n/\mathbb{Q}(i))\}$, by sending every σ to its corresponding matrix M_σ. Thus, we know that the elements of $\mathrm{Gal}(K_n/\mathbb{Q}(i))$ must commute, implying that $\mathrm{Gal}(K_n/\mathbb{Q}(i))$ is an abelian group. □

It can actually be shown that when $\mathrm{Gal}(K/\mathbb{Q}(i))$ is abelian, K must be contained in K_n. In this way, the proofs of Theorem 7.3 and Theorem 7.7 are the starting point to studying every abelian Galois extension over $\mathbb{Q}(i)$. This is certainly an area of study which would follow the discussions of this paper.

Acknowledgments

It is with extreme admiration and gratitude that I thank my mentor, Seraphina Lee. Her patience, advice, and ability to convey the many nuances of elliptic curves resulted in her being a crucial part of my REU experience. I would also like to thank Tarika Mane and Logan Quick for proving to be great group mates and for continually being excited and eager to discuss the fresh mathematical concepts we were studying. I would also like to thank Yulia Kotelnikova for educating me on so much Galois theory. Finally, I would like to thank Professor May for running the REU program. This program has certainly inspired me to continue pursuing math research and in addition provided me with an amazing summer.

References

[1] D.S. Dummit, R.M. Foote. Abstract Algebra, 3rd edn. Wiley. 2004.
[2] E.M. Stein, R. Shakarchi. Complex analysis. Princeton University Press. 2003.
[3] J.H. Silverman, J. Tate. Rational Points on Elliptic Curves. Springer, New York, NY. 1992.
[4] K. Brown. The Primitive Element Theorem. 2010. http://pi.math.cornell.edu/~kbrown/6310/primitive.pdf