

ABSTRACT

Title of dissertation: CALCULATIONS FOR THE ABELIANIZED
ÉTALE FUNDAMENTAL GROUP OF CURVES
OF GENUS ONE OR TWO

Adam Lizzi
Doctor of Philosophy, 2018

Dissertation directed by: Professor Niranjan Ramachandran
Department of Mathematics

The étale fundamental group $\pi_{1,\text{ét}}(C)$ is a profinite group which classifies covers of algebraic-geometric objects. Its abelianization classifies torsors over C with a finite, abelian structure group. When C is a curve over a number field K , $\pi_{1,\text{ét}}^{\text{ab}}(C)$ possesses a finite subgroup $\text{Ker}(C/K)$ which records which abelian étale covers of C come from geometry. In this thesis multiple algorithms for calculating $\text{Ker}(E/K)$ are presented when E is an elliptic curve over a number field. We indicate how to generalize one of these algorithms to the Jacobian of a genus two curve C , allowing for progress to be made on calculating $\text{Ker}(C/K)$.

CALCULATIONS FOR THE ABELIANIZED
ÉTALE FUNDAMENTAL GROUP
OF CURVES OF GENUS ONE OR TWO

by

Adam Joseph Lizzi

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2018

Advisory Committee:

Professor Niranjana Ramachandran, Chair/Advisor

Professor Lawrence Washington

Professor Patrick Brosnan

Professor Thomas Haines

Professor William Gasarch

© Copyright by
Adam Joseph Lizzi
2018

Acknowledgments

Thanks are first due to my advisor, Dr. Niranjan Ramachandran, for your infinite patience and unrelenting positive attitude despite numerous setbacks. Your advice to pursue whatever interests me was never wrong.

Thank you, Rebecca, for your support, sanity checks, and occasional-to-frequent sanity restoration. It's not clear to me I could have succeeded without you. I love you and am so proud of you for making it to the end with me.

Thank you, Dr. Washington—every time we talked I came away with a better sense of direction than I did at the outset. Given how many times this occurred, this was no small feat.

Thank you to my classmates Sean, Steve, Sam, Nathaniel, and Rebecca (again)—you all selflessly put up with many one-sided conversations. I hope I was even a fraction as useful to you as you were to me.

Lastly, thank you to my parents, whose heartfelt desire to see me finish was infectious.

Table of Contents

Acknowledgements	ii
0 Introduction	1
1 Preliminaries	5
1.1 The étale fundamental group	5
1.2 Elliptic curves	9
1.2.1 Torsion	11
1.2.2 Changes of coordinates, invariants	13
1.2.3 Reduction mod p	14
1.3 Isogenies	16
1.3.1 Vélu's formulas	19
1.4 The Tate module and the Galois representation	21
1.5 Elliptic curves over \mathbb{C}	22
1.6 Étale covers of abelian varieties	25
2 $g = 1$, first way: Division polynomials and Vélu's formulas	29
2.1 Division polynomials	29
2.2 Detecting μ -type subgroups	31
2.3 Algorithm	38
3 $g = 1$, second way: modular polynomials	44
3.1 Modular polynomials	44
3.2 Twisting theory of elliptic curves	45
3.3 Algorithm	49
3.4 Alternative method to find $j(E')$: via \mathbb{C}	53
4 Middlebrow arithmetic of curves of genus two	59
4.1 Genus two curves and their Jacobians	59
4.2 The Riemann-Roch formula; Jacobians as groups of divisors	65
4.3 Étale covers of genus two curves	67
4.4 Jacobians over \mathbb{C}	69

5	Calculating a cover in genus two	74
5.1	Igusa invariants	74
5.2	Theta constants	76
5.3	Twisting theory of curves of genus two	79
5.4	(p, p) -isogenies	82
5.5	Tentative algorithm	87
5.6	Example	89
6	Future directions	98
6.1	Proof of correctness for the method of §5.5	98
6.2	Equations for covering curves	100
6.3	What (ℓ, ℓ) -isogenies are possible over \mathbb{Q} ?	100
6.4	Isogenies besides (ℓ, ℓ) -isogenies	100
6.5	Reducible Jacobians	101
6.6	Torsion of Jacobians over \mathbb{Q}^{ab}	103
6.7	Determining $\text{Ker}(E/\mathbb{Z})$ or $\text{Ker}(E/\mathcal{O}_K)$	104
6.8	Relationship to bigger theories	104
A	Cardinality of $\text{Jac}(C)(\mathbb{F}_p)$	106
	Bibliography	112

Chapter 0: Introduction

The main goal in this thesis is to compute certain abelian covers of curves that “come from geometry,” in the sense that no part of them can be explained by a base change $C_L \rightarrow C$ along an abelian extension L/K of the field of definition of C . The group that measures such covers is a subgroup of the group $\pi_1(C)^{\text{ab}}$ classifying all étale covers. Due to an exact sequence in which it sits, we follow Katz and Lang [32] in calling it $\text{Ker}(C/K)$. The main theorem of [32] (Theorem 1.1.1) states that $\text{Ker}(C/K)$ is always finite for varieties over number fields. Our goal will be to calculate what we can regarding $\text{Ker}(C/K)$ when C is a curve of genus one or two.

Genus one curves with a rational point are elliptic, which opens up the door to applying techniques from arithmetic to attack the problem. The preliminaries are discussed in §1. The key result we need to get us off the ground is that of Katz [31] (Theorem 2.2.6) which tells us how to detect whether an elliptic curve could have a p -cover by looking at its reductions modulo primes ℓ of good reduction.

The new results of the thesis are concentrated in §2 and §3, where we give three (two and a half?) algorithms to calculate $\text{Ker}(E/\mathbb{Q})$, indicating how to adapt them to higher-degree ground fields. In some sense these algorithms are dual to Vélú’s

formulas [55], in that Velu’s formulas provide an algorithm for an isogeny with known domain, while our algorithms provide a procedure for creating an isogeny with known target.

As far as we are aware these are the first methods that compute $\text{Ker}(E/\mathbb{Q})$ for an general elliptic curve. In Katz and Lang [32, p. 302] the group $\text{Ker}(X_0(N)/\mathbb{Q})$ is calculated, where here $X_0(N)$ represents the usual compactified modular curve $\text{SL}(2, \mathbb{Z})/\Gamma_0(N)$. (This result is ultimately due to work of Mazur.) For elliptic curves, we’ll see in §1.6 that $\text{Ker}(E/\mathbb{Q})$ can be identified with a certain subset of $\text{Tors}(E(\mathbb{Q}^{\text{ab}}))$, a group which was conjectured finite by Katz and Lang, proven finite by Ribet [48], and calculated by Chou [10]. His characterization is Theorem 1.6.2. Note that Chou determines $\text{Tors}(E(\mathbb{Q}^{\text{ab}}))$ only as an abelian group, while $\text{Ker}(E/\mathbb{Q})$ cares about the Galois action on abelian torison points. See the discussion following Example 2.2.5 for information about the difference between the two. For example, Chou shows that 163 can divide $\# \text{Tors}(E(\mathbb{Q}^{\text{ab}}))$, but it cannot divide $\# \text{Ker}(E/\mathbb{Q})$.

The first of our procedures, Algorithm 2.3.1, works by constructing so-called μ -type subgroups of E directly; this involves studying the factors of the *division polynomials* of E . This algorithm involves constructing points on E defined over extensions of its field of definition. Attempting to stay within the ground field leads us to Algorithm 3.3.1, which works by studying the *modular equation*. This object will tell us the j -invariant of the relevant cover. Then we’ll have to solve a “twisting problem” to turn that one number into the curve we desired. In §3.4 we indicate how to avoid the modular polynomial by working with complex numbers—this latter method is similar in spirit to the procedures discussed in Cremona [16, §3.8].

Throughout multiple examples are given, in line with the over-arching philosophy of this thesis that anything worth doing is worth doing concretely. For instance, Example 2.3.3 uses Algorithm 2.3.1 to demonstrate a nontrivial element of $\text{Ker}(E/\mathbb{Z}[\sqrt{29}])$ for a certain elliptic curve E .

We were motivated to search for multiple algorithms for genus one because we were hoping for a way to expand the scope to genus two. All abelian covers of a genus two curve come via pullbacks of isogenies targeting its Jacobian, so we can again approach the problem from arithmetic. The third approach from §3 does admit a generalization to the Jacobians of genus two curves. We give the relevant arithmetic facts in §4 and explain how to port the algorithm to this larger setting in §5. At the moment we can only control one kind of isogeny, those which go between two Jacobians—these are called (n, n) -isogenies. To do this we need to find points in the Siegel upper half-space corresponding to lattices (p, p) -isogenous to a given one. These are cataloged by the coset representatives of the subgroup $\Gamma_0^{(2)}(p) \subset \text{Sp}(4, \mathbb{Z})$. A correct set of coset representatives does not appear to have existed in the literature; the accurate list is in Proposition 5.4.3. This is the key piece needed to complete Algorithm 5.5.1, and an example of it in action is given in §5.6. The genus two twisting problem is much more complicated than in genus one, so we limit ourselves to the easiest case (where $\text{Aut}(C) \cong \mathbb{Z}/2$).

In §6 we indicate future trails which are invited by our lines of inquiry.

One fact from arithmetic needed in §4, regarding the cardinality of a Jacobian mod p , is proven in Appendix A. Its generalizations to higher genus are discussed as well; for instance, Formula A.1 is the correct combination of $\#C(\mathbb{F}_p)$, $\#C(\mathbb{F}_{p^2})$, and

$\#C(F_{p^3})$ needed to calculate the cardinality of a Jacobian of a curve of genus three.

Chapter 1: Preliminaries

1.1 The étale fundamental group

In this section only we'll approach the subject with “highbrow” language, that is, the Grothendieck language of algebraic geometry. Since this is a relatively isolated phenomenon, we refer you to Hartshorne [23] for relevant background.

The material in this section largely comes from Katz and Lang [32].

Let X and Y be varieties defined over a field K .¹ A morphism $f : X \rightarrow Y$ is called **étale** if it is smooth of relative dimension zero. (This is equivalent to being flat and unramified.) In your head you should picture a topological covering space: the preimage of each point in Y needs to be a dimension zero subset of X , i.e., a finite number of points. Taken in aggregate X should look like a “stack of pancakes” over Y . Note that if $Y = \operatorname{Spec} k$ is the spectrum of a field, étale maps to Y are the spectra of appropriately-named étale algebras, that is, finite products of field extensions of k .

If étale maps are the analogue of covering spaces, then there should be an object like the fundamental group that catalogs them, and indeed this will be our

¹The theory works in more generality, but ultimately we're going to study abelian varieties, so let's not strain ourselves with the preliminaries.

object of study. To avoid technicalities regarding base points, we'll proceed as one does in topology and consider only the abelianization of $(\pi_1)_{\text{ét}}$, henceforth referred to as π_1^{ab} . The object $\pi_1^{\text{ab}}(Y)$ is a profinite group which classifies (fppf) torsors over Y with finite abelian structure group. By “classifies” I mean that for any finite abelian G there is a canonical isomorphism

$$\text{Hom}_{\text{ab}}(\pi_1^{\text{ab}}(Y), G) \cong H_{\text{ét}}^1(Y, G),$$

with the object on the right being étale cohomology. The étale cohomology group $H_{\text{ét}}^1(Y, G)$ classifies “ G -torsors,” that is, morphisms $Z \rightarrow Y$ whose structure group is isomorphic to G . The structure group of a torsor $\pi : X \rightarrow Y$ is the collection of maps $f : X \rightarrow X$ that form a commutative triangle, i.e., $\pi = \pi \circ f$. Note that π_1^{ab} packages together information about all torsors simultaneously with every possible abelian structure group. This makes π_1^{ab} quite large. π_1^{ab} , and indeed π_1 itself, is covariant functorial in its input.

The most basic example of an étale map is a group homomorphism between finite groups $\phi : G \rightarrow H$. The fibers of a group homomorphism are always of equal cardinality, and the inverse image of the various points of H are torsors for the kernel $K = \ker \phi$. The “structure group” in this case is K itself, since a map of the form $g \mapsto g + k$ will wash out upon applying ϕ . If K is abelian, this is exactly the situation we wish to analyze.

If Y is a variety defined over K , the map $Y \rightarrow \text{Spec } K$ induces by functoriality a homomorphism $\pi_1^{\text{ab}}(Y) \rightarrow \pi_1^{\text{ab}}(\text{Spec } K)$. Since Hom is contravariant, we then get

for any finite abelian group G a map

$$\mathrm{Hom}(\pi_1^{\mathrm{ab}}(\mathrm{Spec} K), G) \rightarrow \mathrm{Hom}(\pi_1^{\mathrm{ab}}(Y, G)).$$

If we interpret these Hom groups as H^1 -s, then this map is the “inverse image” function that pulls back a G -torsor over $\mathrm{Spec} K$ along the structure morphism.

The kernel of the map $\pi_1^{\mathrm{ab}}(Y) \rightarrow \pi_1^{\mathrm{ab}}(\mathrm{Spec} K)$ is our main object of study. We’ll denote it by $\mathrm{Ker}(Y/K)$. By definition it sits in an exact sequence

$$0 \rightarrow \mathrm{Ker}(Y/K) \rightarrow \pi_1^{\mathrm{ab}}(Y) \rightarrow \pi_1^{\mathrm{ab}}(\mathrm{Spec} K).$$

In the situation where the structure morphism $Y \rightarrow \mathrm{Spec} K$ has a section ϵ (i.e., Y has a K -rational point y_0), the kernel group classifies those torsors on Y whose inverse image via ϵ is trivial on $\mathrm{Spec} K$. The trivial G -torsor is $\#G$ many copies of $\mathrm{Spec} K$, permuted by the group G . Concretely a G -torsor classified by $\mathrm{Ker}(Y/K)$ is an étale map $f : X \rightarrow Y$ for which $f^{-1}(y_0)$ is a finite number of K -rational points on X .

If L/K is a finite abelian extension, the base change $Y_L \rightarrow Y$ is étale, and therefore is classified by $\pi_1^{\mathrm{ab}}(Y)$. Note how the inverse image of a K -rational point of Y is naturally going to be an L -rational point of Y_L . That suggests that these are the kinds of maps that we wish to avoid; they are the image of $\pi_1^{\mathrm{ab}}(Y)$ in $\pi_1^{\mathrm{ab}}(\mathrm{Spec} K)$. In fact there is an exact sequence

$$0 \rightarrow \mathrm{Ker}(Y/K) \rightarrow \pi_1^{\mathrm{ab}}(Y) \rightarrow \mathrm{Gal}(\overline{K} / K)^{\mathrm{ab}} \rightarrow 0$$

assuming that $Y(K) \neq \emptyset$. In this sequence $\pi_1^{\mathrm{ab}}(K)$ shows up as the group $\mathrm{Gal}(\overline{K} / K)^{\mathrm{ab}}$, since étale covers of $\mathrm{Spec} K$ are abelian extensions of K . (The structure group is the

Galois group!) In the situation where $Y(K) = \emptyset$, things are a little more complicated; there is an action of $\text{Gal}(\overline{K}/K)$ on $\pi_1^{\text{ab}}(Y_{\overline{K}})$ coming from the action of Galois on the equations defining étale covers of Y . The coinvariants of this action are the object that sit in the exact sequence above:

$$0 \rightarrow \pi_1^{\text{ab}}(Y_{\overline{K}})_{\text{Gal}(\overline{K}/K)} \rightarrow \pi_1^{\text{ab}}(Y) \rightarrow \text{Gal}(\overline{K}/K)^{\text{ab}} \rightarrow 0.$$

There is a surjection from these coinvariants to $\text{Ker}(Y/K)$ which is only assuredly an isomorphism if Y has a K -rational point.

The main theorem of Katz and Lang is that the kernel group is *finite*.

Theorem 1.1.1 (Katz-Lang [32], 1981). *Let K be a field finitely generated over its prime field, and let $f : Y \rightarrow \text{Spec } K$ be a variety. The group $\text{Ker}(Y/K)$ is finite if K has characteristic zero, and it is the product of a finite group with a pro- p group if K has characteristic p .*

The basic outline of the proof is to relate the group of coinvariants $\pi_1^{\text{ab}}(Y_{\overline{K}})_{\text{Gal}(\overline{K}/K)}$ to the coinvariants of Tate module of an extension of the Jacobian of Y . We then enter the realm of abelian varieties. The next step is to reduce the question of the finiteness of the coinvariants of the Tate module of the generalized Jacobian to the finiteness of the coinvariants of the Jacobian itself (as well as that of \mathbb{G}_m). We then can dominate this group with the group of $\text{Gal}(\overline{\mathbb{F}_q}/\mathbb{F}_q)$ -coinvariants for the reductions of these abelian varieties, and those quantities turn out to be the cardinality of \mathbb{F}_q -rational points, which is plainly finite.

$\text{Ker}(Y/K)$ is a group defined through abstract nonsense, and it is known to be finite. What can we say about its cardinality? If we cannot state its cardinality

exactly, can we at least say whether a given prime p divides it? This will occupy us for the rest of this thesis.

The next several chapters are concerned with calculating $\text{Ker}(E/\mathbb{Q})$ for E an elliptic curve. Later I'll also present calculations that calculate representatives for nontrivial classes in $\text{Ker}(J/\mathbb{Q})$ for J the Jacobian of a curve of genus two. Since we won't need the genus two theory for some time I'll postpone it until it's needed in Chapter 4.

1.2 Elliptic curves

Many things stated without proof in this section are in Silverman [51].

An **elliptic curve** over a field k is a curve of genus one, together with a marked point P defined over k . The Riemann-Roch theorem implies that any elliptic curve can be presented in so-called “Weierstrass form,” that, is, as the zero locus in \mathbb{A}_k^2 of an equation which has the shape²

$$y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$$

with the $a_i \in k$, together with one additional k -rational point called the point at infinity.³ This point at infinity lives at the top of the vertical axis and is the

²The subscript numbering is chosen so that the polynomial is homogenous of degree 6 should y be given weight 3 and x weight 2. This will be cosmetic for our purposes.

³Perhaps it's better to say that the elliptic curve is the projective closure of this zero locus in \mathbb{P}^2 —the closure will always contain the lone additional point $[0 : 1 : 0]$. Or perhaps a homogeneous equation would be more honest. But in what follows we'll typically work with this affine version.

common point of intersection of all vertical lines. If the characteristic of k is not 2, then completing the square on the left side of the equation converts it into the form

$$y^2 = x^3 + b_2x^2 + b_4x + b_6,$$

and if the characteristic is not 3 then depressing the cubic on the right side gives us

$$y^2 = x^3 + c_4x + c_6.$$


See [51, §III.1] for more details on this process. This will be referred to as “short Weierstrass form.” We’ll endeavor to avoid characteristics 2 and 3 whenever possible and will try to present curves in short Weierstrass form.

Conversely, a choice of two elements $A, B \in k$ will create an elliptic curve over k with equation $y^2 = x^3 + Ax + B$ precisely when the cubic equation $f(x) = x^3 + Ax + B$ has simple roots. This is controlled by the **discriminant** $-4A^3 - 27B^2$. This quantity will be called $\Delta(E)$, though be forewarned that it is a number attached to the equation we’ve chosen for E —it’s not intrinsic to the curve itself. See §1.2.2 regarding changes of coordinates for a fuller treatment.

For any $K \supseteq k$, the points of E with coordinates in K (henceforth denoted $E(K)$) form an abelian group. The group law is most succinctly stated as follows: three points in $E(K)$ sum to zero if and only if they are collinear, and the point at infinity is the neutral element.⁴ If E is in short Weierstrass form then it’s easy to see that a rational point $(x_0, y_0) \in E(K)$ will come with a “mate” $(x_0, -y_0)$. As these two points are connected by the vertical line $x = x_0$, the point at infinity is

⁴Technically any rational point can serve as the neutral element, but we’ll stick to it being the point at infinity, as it is a natural choice when the curve is presented in Weierstrass form.

collinear with these two points. It follows from the description of the group law that (x, y) and $(x, -y)$ are negatives of each other. This explains the placement of the point $P + Q$ in the schematic below.



The structure of the points on elliptic curve over a number field is explained by the following fundamental theorem.

Theorem 1.2.1 (Mordell-Weil). *Let E be an elliptic curve defined over a number field k . The group of points $E(k)$ is finitely generated.*

1.2.1 Torsion

The Mordell-Weil theorem tells us that there is a decomposition of the form

$$E(k) = \mathbb{Z}^r \oplus T$$

for a nonnegative integer r (the “rank”) and a finite abelian group T (the “torsion”). While the rank is mysterious, the torsion group is relatively well-understood. For instance, the Lutz-Nagell theorem gives an succinct algorithm for calculating the torsion points of an elliptic curve over \mathbb{Q} :

Theorem 1.2.1.1 (Lutz-Nagell). *Suppose E/\mathbb{Q} is an elliptic curve in short Weierstrass form, $y^2 = x^3 + ax + b$, with $a, b \in \mathbb{Z}$. Torsion points defined over \mathbb{Q} have integer coordinates. If (x, y) is a torsion point defined over \mathbb{Q} , then $y = 0$ or $y^2 \mid 4a^3 + 27b^2$.*

Points with $y = 0$ are 2-torsion, since the discussion above tells us that those points are equal to their negatives. Assuming we can factor the quantity $4a^3 + 27b^2$, the other torsion points will be laid bare among the factors.

Mazur showed that there's a finite list of options for the torsion group over \mathbb{Q} .

Theorem 1.2.1.2 (Mazur). *Let E/\mathbb{Q} be an elliptic curve. The torsion subgroup of $E(\mathbb{Q})$ is on the following list of sixteen groups:*

$$\left\{ \begin{array}{ll} \mathbb{Z}/n & \text{for } n = 1, 2, \dots, 10, 12, \text{ or} \\ \mathbb{Z}/2 \times \mathbb{Z}/n & \text{for } n = 2, 4, 6, 8. \end{array} \right.$$

Similar theorems are true over number fields of a fixed degree: Kenku and Momose [34] give the full characterization of possible torsion structures over quadratic fields, and various progress has been made on the cubic case [28] [44].

The n -torsion of an elliptic curve is endowed with a nondegenerate, Galois-invariant, alternating pairing called the **Weil pairing**. It is a map $E[n] \times E[n] \rightarrow \mu_n$. The knowledge that the Weil pairing exists is powerful: one basic consequence of its existence is that a field containing all of $E[n]$ must contain μ_n as well. Formulas for calculating it are given in [51, III §8].

1.2.2 Changes of coordinates, invariants

An **isomorphism** of elliptic curves $f : E_1 \rightarrow E_2$ is a regular (polynomial) function from one curve to another with a rational inverse. In the situation where both E_1 and E_2 are presented in Weierstrass form, such a map must look like

$$x \mapsto u^2x + r, \quad y \mapsto u^3y + u^2sx + t,$$

with $u \neq 0$, since we would prefer to fix the common point at infinity. Once we are in short Weierstrass form, linear shifts to either x or y will regenerate some of the missing terms like x^2 or xy . As such we become limited to the centered equations with $r = s = t = 0$; you can only scale x and then scale y by an appropriate amount to keep the cubic polynomial monic.

A change of variables of the form $x \mapsto u^2x, y \mapsto u^3y$ will affect the discriminant by the twelfth power of u . This further emphasizes the fact that the discriminant is only a statistic for the particular equation we have and does not necessarily record truths about the underlying geometric object. The quantity that is shared by all members of an isomorphism class is the **j -invariant**, defined by

$$j = \frac{c_4^3}{\Delta}.$$

If E came to us in short Weierstrass form $y^2 = x^3 + Ax + B$, then c_4 is a fancy way to say A ; the j -invariant in this case is

$$j(E) = 1728 \frac{4A^3}{4A^3 + 27B^2}.$$

In general c_4 is a combination of a_1, \dots, a_6 generated from the process described earlier that takes long Weierstrass equations to short ones.

1.2.3 Reduction mod p

If we are given an equation of the form $y^2 = f(x)$ defining an elliptic curve E over \mathbb{Q} , provided the polynomial $f(x)$ does not have repeated roots modulo p the same equation will define an elliptic curve modulo p . This is the **reduction** of E mod p . The condition on repeated roots is asking for p not to divide the discriminant of f . In this case we say E has **good reduction** mod p .⁵

Note that if p^{12} divides the discriminant of $f(x)$ it is possible that a change of variables could remove p from the discriminant. (See §3.2 where we investigate something similar.) We call an equation for E **minimal at p** if no change of variables can reduce the power of p in its discriminant any further. An equation is **globally minimal** if it is simultaneously minimal for all primes. A global minimal equation for E is only guaranteed to exist when the class number of the field of definition is one; see [51, VIII, Prop. 8.2] for a precise statement. Note that the minimal equation need not be in short form. The behavior of an equation minimal at p is the “true” (geometric) behavior of the curve E at p .

When the reduction is not good the curve will have a singularity modulo p . Since the degree of f is so small there are only two options for what the singularity could be: a node or a cusp. The reduction is called **multiplicative** if it has a node modulo p and **additive** if it has a cusp. Heuristically a node looks like $y^2 = x^3 + Ax^2$

⁵Things are a little more complicated modulo 2 and 3 since not every elliptic curve admits an equation in short Weierstrass form modulo those primes. This is another reason why we’re doing our best to avoid $p = 2$ and $p = 3$.

and a cusp looks like $y^2 = x^3$. Specific conditions for the reduction type are given by [51, VII, Prop. 5.1]:

Proposition 1.2.3.1. *Let E/K be an elliptic curve with a long Weierstrass equation that is minimal at p . Let Δ be the discriminant of this equation and c_4 the constant constructed above.*

1. *E has good reduction at p if $p \nmid \Delta$.*
2. *E has multiplicative reduction at p if $p \mid \Delta$ and $p \nmid c_4$.*
3. *E has additive reduction at p if $p \mid \Delta$ and $p \mid c_4$ (in this case p necessarily divides c_6 as well).*

The reduction types are often packaged into an invariant called the **conductor** of E , written N_E . When E is defined over \mathbb{Q} , the conductor is a product over all primes

$$N_E = \prod_p p^{a_p}$$

with exponents $a_p = 0$ if E has good reduction at p , $a_p = 1$ if E has additive reduction at p , and $a_p = 2$ if E has multiplicative reduction at p .⁶ (For larger number fields a more conceptual definition of N_E is needed, but we'll largely stay over \mathbb{Q} in what follows.)

One way in which good reduction is appropriately named is that the global torsion of E injects into the finite group $E(\mathbb{F}_p)$ for any prime of good reduction.

⁶The exponents on 2 and 3 could be higher if there is additive reduction at those primes.

1.3 Isogenies

We have groups (elliptic curves) to study. An obvious next step is to discuss the homomorphisms between those groups.

Definition 1.3.1. Let E_1 and E_2 be elliptic curves over a field k . An **isogeny** between them is a morphism of varieties sending the point at infinity on E_1 to the point at infinity on E_2 .

It may seem strange not to ask for our maps to be homomorphisms in the definition, but this turns out to be unnecessary: every isogeny is automatically a homomorphism. One could instead say define an isogeny as a homomorphism $E_1 \rightarrow E_2$ with finite kernel. (We need to exclude the zero map.) The **degree** of an isogeny is the cardinality of its kernel. Some caution is needed in finite characteristic because p -isogenies in characteristic p can be “inseparable”, like the **Frobenius endomorphism** $(x, y) \mapsto (x^p, y^p)$. This has only the point at infinity in its kernel, yet is degree p . Since étale covers need to have preimages that consist entirely of rational points, inseparable isogenies will not rear their heads in what follows.

An isogeny from a curve E to itself is called an **endomorphism**. A typical endomorphism is the “multiplication-by- n ” map, $[n] : E \rightarrow E$, defined by

$$[n](P) = \underbrace{P + \cdots + P}_{n \text{ times}} = nP.$$

The degree of multiplication by n is n^2 ; the kernel is the points of order n , which we had denoted $E[n]$. Over the algebraic closure, $E[n] \cong \mathbb{Z}/n \times \mathbb{Z}/n$. We’ll sketch a proof of this in §1.5. Under most circumstances, these are the only endomorphisms

that an elliptic curve E defined over \mathbb{Q} has—as a ring, $\text{End}(E) \cong \mathbb{Z}$. A special subset of curves over \mathbb{Q} have extra endomorphisms, called **complex multiplication** since $\text{End}(E) \otimes \mathbb{Q}$ is always an imaginary quadratic field when this occurs.⁷

If $\phi : E_1 \rightarrow E_2$ is an isogeny, there is always an isogeny $\hat{\phi} : E_2 \rightarrow E_1$ going the other way called the **dual isogeny**. This stems from the fact that an elliptic curve is isomorphic to its own Jacobian variety; see [51, III §6]. The composition of an isogeny with its dual is the endomorphism $[\deg \phi] : E_1 \rightarrow E_1$, where $\deg \phi$ is the degree of the map as defined above. Consequently we can impose an equivalence condition on elliptic curves by saying $E_1 \sim E_2$ if there is an isogeny between them. The conductor is an isogeny invariant: isogenous primes share the same reduction type at every prime [51, §16].

Let $E_1 : y^2 = f_1(x)$ and $E_2 : y^2 = f_2(x)$. In more down-to-earth terms an isogeny will be a pair of rational maps $X(x, y)$, $Y(x, y)$ which have the property that if (x, y) is a point satisfying $y^2 = f_1(x)$, then (X, Y) satisfies the equation $Y^2 = f_2(X)$. For instance, an isogeny of degree two between the curves $y^2 = x^3 + 4x$ and $y^2 = x^2 - 16x$ is given by the rational maps

$$(x, y) \mapsto \left(\frac{x^2 + 4}{x}, \frac{y(x^2 - 4)}{x^2} \right).$$

An isogeny is **k-rational**, and the two curves are called **isogenous over k**, if the equations defining the isogeny have coefficients in **k**. The isogeny is defined over the same field as its kernel, but beware: the field of definition of the kernel may be

⁷Modulo p there is a possibility that $\text{End}(E)$ could be an order in a quaternion algebra! These are called *supersingular* curves. If E is defined over \mathbb{Q} , the set of primes at which the reduction modulo p is supersingular is infinite. This is a result of Elkies [18].

different from the field of definition of the points in the kernel. An isogeny is defined over the same field as its dual.

If E/\mathbb{Q} is an elliptic curve, there are limits to how many isogenies E can have that are defined over \mathbb{Q} . A large research program culminated in a characterization of which degrees are possible. The precise statements here come from Chou [10].

Theorem 1.3.2 (Kenku, Mazur, Ogg, \dots). *Let E/\mathbb{Q} be an elliptic curve and $\phi : E \rightarrow E'$ a cyclic isogeny defined over \mathbb{Q} . Then $\deg \phi \leq 19$ or $\deg \phi \in \{21, 25, 27, 37, 43, 67, 163\}$.*

In fact, if you have an isogeny, that “crowds out” others.

Theorem 1.3.3 (Kenku [33]). *Up to isomorphism over \mathbb{Q} , at most eight other curves are isogenous to a given curve.*

There are specific rules about how many isogenies are possible if one is known; for example, if a curve possesses a \mathbb{Q} -rational 7-isogeny, it cannot possess any other isogenies except possibly one 2-isogeny or one 3-isogeny. The full set of stipulations are in Kenku’s paper.

For elliptic curves there is a pleasing dictionary between finite subgroups of E and isogenies out of E : every finite subgroup H determines an isogeny, the “mod out by H ” isogeny, to an elliptic curve whose points are abstractly isomorphic to E/H . See [51, III Prop. 4.12]. We’ll discuss this isogeny in detail in the next section.

1.3.1 Vélu's formulas

These formulas give in full detail the equations for the isogeny between two elliptic curves when the kernel is specified. As mentioned above, there is a dictionary between isogenies out of E and finite subgroups $H \subset E(\overline{K})$ which could serve as the kernel. We're only going to use the situation where H is cyclic, which will simplify matters somewhat. These formulas are, as the name implies, due to Vélu [55]; the treatment here is adapted from Washington's book [56, §12.3].

Formulas 1.3.1.1. *Let E be an elliptic curve defined over a field K of characteristic zero with a short Weierstrass equation*

$$y^2 = x^3 + a_4x + a_6$$

for constants $a_4, a_6 \in K$. Choose $C \subset E(\overline{K})$ a finite cyclic subgroup generated by a point P of order n . Let α be the isogeny out of E whose kernel is C and let E_2 be the codomain of α .

For a point $Q = (x_Q, y_Q) \neq \infty$ in C , define the quantities

$$g_Q^x = 3x_Q^2 + a_4$$

$$g_Q^y = -2y_Q$$

$$v_Q = \begin{cases} g_Q^x & \text{if } 2Q = \infty \\ 2g_Q^x & \text{if } 2Q \neq \infty \end{cases}$$

$$u_Q = (g_Q^y)^2.$$

Let S be the set $\{P, 2P, 3P, \dots, \lfloor \frac{n}{2} \rfloor P\}$. Set

$$v = \sum_{Q \in S} v_Q, \quad w = \sum_{Q \in S} (u_Q + x_Q v_Q).$$

The isogenous curve E_2 has equation

$$Y^2 = X^3 + (a_4 - 5v)X + (a_6 - 7w),$$

and the isogeny $E \rightarrow E_2$ is given by

$$X = x + \sum_{Q \in S} \left(\frac{v_Q}{x - x_Q} + \frac{u_Q}{(x - x_Q)^2} \right)$$

$$Y = y - \sum_{Q \in S} \left(u_Q \frac{2y}{(x - x_Q)^3} + v_Q \frac{y - y_Q}{(x - x_Q)^2} - \frac{g_Q^x g_Q^y}{(x - x_Q)^2} \right).$$

These equations work in positive characteristic as well, though the isogeny α is only unique if we ask it for it to be separable. They can also be adapted to deal with the case where the kernel is not cyclic.

Idea of proof. Showing that the formulas work amount to a calculation in the formal group of the elliptic curve. Develop x and y as Laurent series in the uniformizer $t = \frac{x}{y}$ at infinity. One then checks that the equations given for X and Y , as power series in t , satisfy the Weierstrass equation suggested for E_2 . It follows that the assignment $x \mapsto X, y \mapsto Y$ determines a function from E to E_2 . The fact that the denominators of X and Y are zero at the points of C implies that the points of C on E map to infinity on E_2 , i.e., they are in the kernel of the isogeny. \square

1.4 The Tate module and the Galois representation

Let p be a prime and E an elliptic curve defined over a field k . Repeated multiplication by p creates a projective system of groups

$$\cdots \rightarrow E[p^n] \rightarrow E[p^{n-1}] \rightarrow \cdots \rightarrow E[p^2] \rightarrow E[p] \rightarrow 0.$$

The limit $\varprojlim_{n \rightarrow \infty} E[p^n]$ of this sequence is called the **p -adic Tate module**, denoted $T_p E$. Since $E[p^n] \cong \mathbb{Z}/p^n \times \mathbb{Z}/p^n$, the Tate module is abstractly isomorphic to $\mathbb{Z}_p \times \mathbb{Z}_p$, with \mathbb{Z}_p the ring of p -adic integers.

There is an interesting Galois action on the Tate module stemming from the action of the absolute Galois group on the p -power torsion points. Form the tensor product $V_p = T_p E \otimes \mathbb{Q}_p$; this is a two-dimensional vector space which carries that action of Galois, extended linearly. Thus we have a representation $\rho_p : \text{Gal}(\bar{k}/k) \rightarrow \text{GL}_2(\mathbb{Q}_l)$ called the **p -adic Galois representation**. When k is a finite field or a number field, this representation is an isogeny invariant by the proven cases of the Tate conjecture [51, III.7.7].

The Galois representation can be “built up” projectively much as the Tate module was. The action of Galois on the p^n -torsion points defines a map

$$\rho_{p,n} : \text{Gal}(\bar{k}/k) \rightarrow \text{GL}_2(\mathbb{Z}/p^n).$$

Explicitly, suppose $E[p^n] = \langle P, Q \rangle$. An element $\sigma \in \text{Gal}(\bar{k}/k)$ must move P and Q to other p^n -torsion points, since multiplication by p^n is an isogeny from E to itself and therefore the equation “ $[p^n](P) = 0$ ” is a collection of rational expressions that the coefficients of P satisfy with k -rational coefficients. (In fact, the coefficients

of $[p^n]$ are in \mathbb{Z} .) As P and Q generate the p^n -torsion, $\sigma(P) = aP + cQ$ and $\sigma(Q) = bP + dQ$ for some values a, b, c, d determined modulo p^n . The effect that σ has on p^n -torsion is the same as the effect the matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ has on the two-dimensional \mathbb{Z}/p^n -module $E[n]$. We therefore define $\rho_{p,n}(\sigma) = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$. Assuming we have chosen compatible bases for $E[p^n]$ for all n , the Galois representation $\rho_p(\sigma)$ is the inverse limit of these matrices in $\mathrm{GL}_2(\mathbb{Q}_p)$.

1.5 Elliptic curves over \mathbb{C}

When $k = \mathbb{C}$, there is another way to describe elliptic curves: as the quotient of \mathbb{C} by a lattice $\Lambda \subset \mathbb{C}$. Here a **lattice** is the \mathbb{Z} -span of two complex numbers $\{\omega_1, \omega_2\}$ that are linearly independent over \mathbb{R} ; this makes the set $\Lambda = \{n_1\omega_1 + n_2\omega_2 : n_1, n_2 \in \mathbb{Z}\}$ a discrete subgroup of \mathbb{C} . The quotient \mathbb{C}/Λ is topologically a torus; if

$$\mathcal{F} = \{r_1\omega_1 + r_2\omega_2 : r_1, r_2 \in \mathbb{R}, 0 \leq r_i < 1\}$$

is a fundamental domain for the quotient, \mathcal{F} is a parallelogram with opposite sides identified. Topologically, \mathcal{F} is a surface of genus one which is a group under addition of complex numbers. Every elliptic curve over \mathbb{C} is isomorphic to a quotient of \mathbb{C} by some lattice—we'll worry about how to find a lattice representing a given elliptic curve in §3. Given a lattice Λ , one can create an equation for the elliptic curve \mathbb{C}/Λ using differential equations satisfied by the Weierstrass \wp -function whose set of poles is Λ . See Silverman [51, § VI.3] for more information.

Two lattices determine isomorphic elliptic curves if you can scale one to create

the other. I'll phrase this in a slightly different way with an eye to genus two.

Definition 1.5.1. Two lattices $\Lambda_1, \Lambda_2 \subset \mathbb{C}$ are called **homothetic** if there is a linear transformation $f : \mathbb{C} \rightarrow \mathbb{C}$ whereby $f(\Lambda_1) = \Lambda_2$. The two lattices are **isogenous** if there is a linear transformation $f : \mathbb{C} \rightarrow \mathbb{C}$ whereby $f(\Lambda_1) \subset \Lambda_2$. The **degree** of an isogeny is the index $[\Lambda_2 : f(\Lambda_1)]$.

Since the only linear transformations from \mathbb{C} to \mathbb{C} are scalar multiplications, this is the same as asking that there be a nonzero complex number α whereby $\alpha\Lambda_1 \subseteq \Lambda_2$. Homothety and isogeny of lattices correspond to isomorphism and isogeny of elliptic curves—if two lattices are isogenous then there will be a homomorphism $\mathbb{C}/\Lambda_1 \rightarrow \mathbb{C}/\Lambda_2$. A homothety is an isogeny of degree one.

The j -invariant of the elliptic curve \mathbb{C}/Λ can be calculated from any representative in the homothety class of Λ . Define the quantities

$$g_2(\Lambda) = 60 \sum_{\substack{\omega \in \Lambda \\ \omega \neq 0}} \omega^{-4} \quad \text{and} \quad g_3(\Lambda) = 140 \sum_{\substack{\omega \in \Lambda \\ \omega \neq 0}} \omega^{-6}.$$

Then

$$j(\Lambda) = 1728 \frac{g_2(\Lambda)^3}{g_2(\Lambda)^3 - 27g_3(\Lambda)^2}.$$

By using $\alpha = \frac{\omega_1}{\omega_2}$ or its reciprocal, every lattice is homothetic to one spanned by $\{1, \tau\}$ for a complex number τ whose imaginary part is positive. The entire upper half-plane is much larger than the space of isomorphism classes of elliptic curves, since lattices have multiple bases. The lattice spanned by $\{1, \tau\}$ is the same as the lattice spanned by $\{a\tau + b, c\tau + d\}$ whenever a, b, c, d are integers satisfying $\det \begin{bmatrix} a & b \\ c & d \end{bmatrix} = 1$. Again applying a homothety, this shows that the elliptic curve

$\mathbb{C}/\langle 1, \tau \rangle$ is isomorphic to the elliptic curve $\mathbb{C}/\langle 1, \frac{a\tau+b}{c\tau+d} \rangle$. If we define an action of $\mathrm{SL}(2, \mathbb{Z})$ on the upper half-plane by

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \cdot \tau = \frac{a\tau + b}{c\tau + d},$$

the orbit space is exactly the set of isomorphism classes of elliptic curves over \mathbb{C} .

This construction gives a pithy explanation for the shape of the n -torsion subgroup of an elliptic curve defined over a field of characteristic zero. The points in a fundamental parallelogram for $\mathbb{C}/\langle \omega_1, \omega_2 \rangle$ that are moved into Λ by multiplication by n are exactly the points of the form $\left\{ \frac{a\omega_1}{n} + \frac{b\omega_2}{n} \right\}$. As a group, this set is abstractly isomorphic to $\mathbb{Z}/n \oplus \mathbb{Z}/n$. If E is defined over a number field, its field of definition can be embedded into \mathbb{C} , implying that $E[n]$ is isomorphic to $\mathbb{Z}/n \oplus \mathbb{Z}/n$ as well. Note that this is only true over $\overline{\mathbb{Q}}$; the points of $E[n]$ need not be defined over the field of definition of E . (This is good for us—we want $E[n]$ to have an interesting Galois action!)

For later reference, we will need the set of lattices p -isogenous to a given lattice Λ . If Λ' is such an isogenous lattice with isogeny f , $\Lambda/f(\Lambda')$ is a group of order p . This implies $f(\Lambda')$ lives between Λ and $p\Lambda$. There are only $p+1$ such lattices, corresponding to the $p+1$ subgroups of $\Lambda/p\Lambda \cong \mathbb{Z}/p \oplus \mathbb{Z}/p$ of order p . Suppose that Λ_1 and Λ_2 are both p -isogenous to Λ via multiplication by α_1 and α_2 respectively, and further suppose $\alpha_1\Lambda_1$ and $\alpha_2\Lambda_2$ land on the same sublattice of Λ of index p . Then $\frac{\alpha_1}{\alpha_2}\Lambda_1 = \Lambda_2$, so Λ_1 and Λ_2 are homothetic. So if we can present one isogenous lattice corresponding to each sublattice of Λ of index p , we'll have our

set of representatives. As mentioned these are only determined up to homothety, but that means there will be a representative of each homothety class whereby the isogeny is multiplication by p . This representative will be a superlattice of Λ which possesses Λ as a sublattice of index p . Here are those representatives, together with which subgroup they correspond to in $\Lambda/p\Lambda$.

Proposition 1.5.2. *Let $\Lambda = \langle \omega_1, \omega_2 \rangle$. The following table lists representatives for the $p + 1$ homothety classes of lattices that are p -isogenous to Λ . Let $\Lambda/p\Lambda \cong \mathbb{Z}/p\mathbb{Z} \oplus \mathbb{Z}/p\mathbb{Z}$ via the isomorphism $(\omega_1, 0) + p\Lambda \mapsto (1, 0)$ and $(0, \omega_2) + p\Lambda \mapsto (0, 1)$.*

<i>basis for superlattice</i>	<i>generator for subgroup of $\Lambda/p\Lambda$</i>
$\left\langle \frac{\omega_1 + j\omega_2}{p}, \omega_2 \right\rangle_{j=0}^{p-1}$	$\langle (1, j) \rangle$
$\left\langle \frac{\omega_2}{p}, -\omega_1 \right\rangle$	$\langle (0, 1) \rangle$

For information on how this works when p is not prime, see Cox [15, § 11.C].

1.6 Étale covers of abelian varieties

We return in this section to the question at hand: computing $\text{Ker}(E/\mathbb{Q})$ for E an elliptic curve. We've seen that all morphisms between elliptic curves are necessarily group homomorphisms (isogenies). Let's interpret this in terms of the fundamental group. We know that $\text{Ker}(E/\mathbb{Q})$ should classify maps $\phi : E' \rightarrow E$ for which the inverse image of a rational point on E is a collection of rational points on E' . The Riemann-Hurwitz formula [23, IV §2 2.4] together with the fact that ϕ is unramified implies that the genus of E' must also be one. Now the fact that both curves have rational points means that we can apply translations to arrange it so

that the rational points in question on both curves are their origins. By the theory above, ϕ is a homomorphism, and the inverse image of the identity element will be the kernel of this homomorphism. Thus we have shown the following.

Proposition 1.6.1. *Let E/\mathbb{Q} be an elliptic curve. The group $\text{Ker}(E/\mathbb{Q})$ classifies elliptic curves E' which possess an isogeny into E for which the kernel consists entirely of rational points.*

Note that the ground field could be any number field without changing the argument. Similar logic applies to abelian varieties in any dimension, since the only covers of abelian varieties are other abelian varieties of the same dimension via isogenies [43, §18].

When A is an abelian variety over a number field K , Katz and Lang point out the following upper bound for the size of $\text{Ker}(A/K)$ on page 301 of [32]. The Weil pairings for abelian varieties are defined as maps $e_n : A[n] \times A^\vee[n] \rightarrow \mu_n$, where A^\vee is the dual variety to A . (Non-cognoscenti can think $A = A^\vee = E$, since elliptic curves are self-dual.) Now taking limits and colimits as appropriate leads to a Galois-equivariant pairing

$$T(A(\overline{K})) \times \text{Tors } A^\vee(\overline{K}) \rightarrow \mathbb{Q}/\mathbb{Z}(1).$$

Here $T(\cdot)$ is the product of all the Tate modules: the inverse limit along multiplications by n , for all n . Twisting by the inverse of the cyclotomic character will scrub out the action of Galois on the right.

$$T(A(\overline{K})) \times \text{Tors } A^\vee(\overline{K})(-1) \rightarrow \mathbb{Q}/\mathbb{Z}.$$

This duality, upon applying Galois invariants and coinvariants, becomes a perfect pairing of the shape

$$T(A(\overline{K}))_{\text{Gal}(\overline{K}/K)} \times \text{Tors } A^\vee(\overline{K})(-1)^{\text{Gal}(\overline{K}/K)} \rightarrow \mathbb{Q}/\mathbb{Z}.$$

The former object is isomorphic to $\text{Ker}(A/K)$ under the reductions from the proof of the main theorem. It follows that this last pairing is one between two finite groups, and that the size of $\text{Ker}(A/K)$ is that of the group $\text{Tors } A^\vee(\overline{K})^\chi$ of rational points of A^\vee which are transformed under Galois by the cyclotomic character.

In the case $A = A^\vee$ we are asking to find torsion points $P \in \text{Tors}(A(\overline{K}))$ for which the Galois action on P matches that of the cyclotomic character (hence the action of Galois on the roots of unity). Such points are necessarily defined over an abelian extension of K ; $\langle P \rangle \subset A$ is a sub-Galois module isomorphic to a group of roots of unity. A group admitting such an action is called **μ -type**. Thus we are hunting for the maximal μ -type subgroup of A , in the sense that all of its cyclic subgroups should be μ -type. This group can be expressed as a limit, as on page 302 of [32]; namely, $\varprojlim_N \text{Hom}_{K\text{-gpsch}}(\mu_N, A)$. Again, is the Pontrjagin dual to $\text{Ker}(A/K)$.

It follows from the above discussion that $\text{Tors}(A(\overline{K}))^\chi \subseteq \text{Tors}(A(K(\mu)))$, where $K(\mu)$ is the field obtained by adjoining all roots of unity to K . This latter group was proven to be finite for any abelian variety A over a number field by Ribet in the appendix to Katz-Lang [48]. In the case $K = \mathbb{Q}$, note $\mathbb{Q}(\mu) = \mathbb{Q}^{\text{ab}}$. In the case A is an elliptic curve, the classification of possible torsion structures over \mathbb{Q}^{ab} was recently completed by Chou [10]:

Theorem 1.6.2 (Chou). *If E/\mathbb{Q} is an elliptic curve, $\text{Tors}(E(\mathbb{Q}^{\text{ab}}))$ is one of the following groups:*

$$\mathbb{Z}/N \quad N \in \{1, 3, 5, \dots, 19, 21, 25, 27, 37, 43, 67, 163\}$$

$$\mathbb{Z}/2 \times \mathbb{Z}/2N \quad N \in \{1, 2, \dots, 9\}$$

$$\mathbb{Z}/3 \times \mathbb{Z}/3N \quad N \in \{1, 3\}$$

$$\mathbb{Z}/4 \times \mathbb{Z}/4N \quad N \in \{1, 2, 3, 4\}$$

$$\mathbb{Z}/5 \times \mathbb{Z}/5$$

$$\mathbb{Z}/6 \times \mathbb{Z}/6$$

$$\mathbb{Z}/8 \times \mathbb{Z}/8$$

We'll discuss the question of whether equality holds in the containment $\text{Ker}(E/\mathbb{Q}) \subseteq \text{Tors}(E(\mathbb{Q}^{\text{ab}}))$ in the next chapter.

Chapter 2: $g = 1$, first way: Division polynomials and

Vélu's formulas

In this chapter we'll characterize under what circumstances an elliptic curve E/K possesses an étale cover of degree p . We'll largely work under the assumption $K = \mathbb{Q}$, but when facts hold in more generality they will be pointed out.

2.1 Division polynomials

In the first chapter we considered the important multiply-by- n endomorphisms $[n] : E \rightarrow E$. The kernel of this map consisted of torsion points on E whose order divided n . As these isogenies consist of rational functions with coefficients in the ground field, their kernels will be defined over the ground field. This means the coordinates of these torsion points must satisfy a polynomial with coefficients in the ground field—we'll give a name to this polynomial.

Definition 2.1.1. The n -th **division polynomial** for E , written $\psi_{n,E}(t) = \psi_n(t)$, is the polynomial whose roots are the x -coordinates of the n -torsion points on E .¹

For example, if E has a short Weierstrass equation $y^2 = x^3 + Ax + B$ we saw

¹That is, the *unique* x -coordinates. If $P = (x, y) \in E[n]$ and P is not of order 2, then $-P = (x, -y) \in E[n]$ as well. Nonetheless $(t - x)$ divides ψ_n to the first power.

the 2-torsion points are the points of the form $(x, 0)$, so $\psi_2(t) = f(t)$. The division polynomials quickly get more complicated. $\psi_3(t) = 3t^4 + 6At^2 + 12Bt - A^2$ and so on. The degree of ψ_m is $\frac{m^2-1}{2}$. There is a recursive procedure that calculates the division polynomials in terms of previous ones which is implemented in any computer algebra system that has routines for elliptic curves.

Note that the minimal field of definition for the m -torsion subgroup of E is not exactly $K_m = K[x]/\psi_m(x)$. The issue is that we didn't try to include the y -coordinates in our definition of division polynomial. The minimal field of definition turns out to be either K_m or a quadratic extension of it.

In a similar manner we can define the **kernel polynomial** for any isogeny leaving E : its roots are the x -coordinates of the points in the kernel of the isogeny. If the isogeny $\phi : E \rightarrow E'$ has degree m , then all the points in the kernel must be contained in $E[m]$. Following the logic as above, the points in the kernel of the isogeny defined over the ground field must have x -coordinates which are the roots of a polynomial with coefficients in the ground field, and this polynomial must divide $\psi_m(x)$. So the division polynomials indirectly know about isogenies leaving E .

If $\deg \phi$ is odd, the degree of the kernel polynomial is $\frac{\deg \phi - 1}{2}$, as the points in the kernel of ϕ will pair up into packets of the form $\{P, -P\}$. Each packet contains two points with the same x -coordinate. If $\deg \phi$ is even and the kernel of ϕ contains s points of order two, the same logic implies the kernel polynomial will have degree $\frac{\deg \phi - s}{2}$.

2.2 Detecting μ -type subgroups

Let E/\mathbb{Q} be an elliptic curve. Recall some of the concepts from §1: For $n > 1$, the N -torsion subgroup of E is the group $E[n] = E(\overline{\mathbb{Q}})[n]$. It is abstractly isomorphic to $\mathbb{Z}/n \times \mathbb{Z}/n$. As such, the action of Galois on this group determines a homomorphism

$$\rho_n : \text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \rightarrow \text{Aut}(\mathbb{Z}/n \times \mathbb{Z}/n)$$

which we'll call the mod- n Galois representation. Almost always we'll pick a basis for $E[n]$, as doing so allows us to identify $\text{Aut}(\mathbb{Z}/n \times \mathbb{Z}/n)$ with $\text{GL}_2(\mathbb{Z}/n)$. We'll often abuse notation even further, by conflating an element $\sigma \in \text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ with the matrix obtained in this process.

It's a standard consequence of properties of the Weil pairing $e_n(\cdot, \cdot)$ that, should a field K satisfy $E[n] \subset K$, necessarily as well $\mu_n \subset K$. Galois will also act on μ_n , and it's natural to ask whether we can interpret the action on the roots of unity in terms of the representation of the previous paragraph. We'll refer to the answer frequently, so I'll record it here:

Basic Fact 2.2.1. *Suppose $\{P, Q\}$ is a spanning set for $E[n]$. If $\sigma \in \text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$, identify $\rho_n(\sigma)$ with a matrix in $\text{GL}_2(\mathbb{Z}/n)$. Then σ acts on μ_n by raising to the $\det \rho_n(\sigma)$ power.*

Proof. Suppose $\rho_n(\sigma)$ corresponds to the matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$, i.e., $\sigma(P) = aP + bQ$ and $\sigma(Q) = cP + dQ$. Since P and Q generate $E[n]$, applying the Weil pairing to them will produce a primitive n -th root of unity. Denote this by $\zeta = e_n(P, Q)$. Then act

on both sides by σ . All the conjugates of ζ are powers of ζ , so on the left we get a mystery power: $\sigma(\zeta) = \zeta^w$. On the right, we apply the fact that the Weil pairing is Galois-equivariant, alternating, and bilinear:

$$\begin{aligned}
\sigma(\zeta) &= \zeta^w = \sigma(e_n(P, Q)) \\
&= e_n(\sigma(P), \sigma(Q)) \\
&= e_n(aP + bQ, cPQ) \\
&= e_n(aP, cP) e_n(aP, dQ) e_n(bQ, cP) e_n(bQ, dQ) \\
&= e_n(P, Q)^{ad} + de_n(Q, P)^{bc} \\
&= e_n(P, Q)^{ad-bc} = \zeta^{ad-bc}.
\end{aligned}$$

□

Although we phrased this result in terms of a specific generating set for $E[n]$, this wasn't essential: any other generating set will pair to a different primitive n -th root of unity, and σ must act the same way on each primitive n -th root of unity. If χ_n denotes the n -th cyclotomic character, then our Basic Fact says that $\det \rho_n = \chi_n$.

We are searching for subgroups of $E[n]$ that are isomorphic to μ_n as Galois modules. Recall that the maximal μ -type subgroup represents (the dual of) the group $\text{Ker}(E/\mathbb{Q})$ of interest. We'll say that a point $P \in E[n]$ "transforms according to the (n -th) cyclotomic character" if $\sigma(P) = \chi_n(\sigma)P$ for all $\sigma \in \text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$. This is saying in symbols that the cyclic subgroup generated by P is isomorphic to μ_N , as the map $P \mapsto \zeta$ for any primitive n -th root of unity ζ will be a Galois-equivariant bijection. Such a point P has $\phi(n)$ Galois conjugates and the Galois group acts on it

by multiplication by scalars. Thus P will be defined over an abelian extension of \mathbb{Q} of degree $\phi(n)$. This is not a major surprise, as we saw $\text{Tors}(E(\overline{\mathbb{Q}}))^{\chi_n} \subseteq \text{Tors}(E(\mathbb{Q}^{\text{ab}}))$.

If $P \in E[n]$ is a point that transforms according to the cyclotomic character and we choose P as one of our elements for a generating set for $E[n]$, then the matrices appearing in the image of the Galois representation ρ_n will all have the form

$$\begin{bmatrix} \chi_n(\sigma) & b \\ 0 & d \end{bmatrix}.$$

Our basic fact says that the determinant of this matrix should be $\chi_n(\sigma)$, so evidently $d = 1$. The following construction uses this logic to give us a source of points that transform according to the cyclotomic character.

Proposition 2.2.2. *Let E/\mathbb{Q} be an elliptic curve that possesses a \mathbb{Q} -rational point P of order N . If $\phi : E \rightarrow E/\langle P \rangle$ is the isogeny emanating from E whose kernel is the subgroup generated by P , then the quotient curve $E/\langle P \rangle$ possesses a point that transforms under the cyclotomic character.*

Proof. Let Q be an independent point of order n on E , so that $\{P, Q\}$ forms a generating set for $E[n]$. Since P is rational, Galois acts on $E[n]$ via matrices of the form

$$\begin{bmatrix} 1 & b \\ 0 & d \end{bmatrix}.$$

As above, the basic fact tells us that $d = \chi_n(\sigma)$. Thus an element $\sigma \in \text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ acts on Q by the formula $\sigma(Q) = bP + \chi_n(\sigma)Q$. If we apply the homomorphism ϕ , P is killed. Since the kernel of ϕ consists of rational points, Velú's formulas imply

that ϕ has rational coefficients. Therefore it commutes with the action of Galois, and we see that

$$\sigma(\phi(Q)) = \phi(\sigma(Q)) = \phi(bP + \chi_n(\sigma)Q) = \phi(\chi_n(\sigma)(Q)) = \chi_n(\sigma)\phi(Q),$$

so $\phi(Q)$ transforms according to the cyclotomic character. \square

Notice the similarities of the two centered matrices above. This construction shows that the isogeny ϕ is reversing the roles of rational subgroups (\mathbb{Z}/n) and μ -type subgroups. Despite being phrased in terms of elliptic curves defined over \mathbb{Q} , the result holds for curves over arbitrary fields, as Velú's formulas aren't limited to isogenies over \mathbb{Q} . This construction admits a converse; these two propositions together fully characterize when an elliptic curve possesses a subgroup isomorphic to μ_n .

Proposition 2.2.3. *Let E/\mathbb{Q} be an elliptic curve possessing a point Q which transforms according to the n -th cyclotomic character. There exists an elliptic curve E' with a rational point $P \in E'[n](\mathbb{Q})$ which is isogenous to E via the quotient-by- P map.*

Proof. Let ψ be the isogeny emanating from E whose kernel is the subgroup generated by Q . By our assumption on Q , we have that $R \in \ker \psi$ implies $\sigma(R) \in \ker \psi$ for all points R and all automorphisms $\sigma \in \text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$. This means that the subgroup $\ker \psi$ is defined over \mathbb{Q} , and consequently the coefficients of ψ are defined over \mathbb{Q} . The codomain of ψ will then also be defined over \mathbb{Q} . This will be our E' .

Now let R represent an independent point of order n on E . Then as we've seen the Galois representation ρ_n has an image consisting of matrices of the form

$\begin{bmatrix} x & b \\ 0 & 1 \end{bmatrix}$, and so by the logic of the previous proof $\psi(R)$ will be a point of order n with rational coordinates. Now consider the isogeny dual to ψ , which we'll call $\widehat{\psi}$ for the moment. We know $\widehat{\psi} \circ \psi = [n]$, so $\widehat{\psi}(\psi(R)) = nR = 0$ as R was a point of order n . However R is not in the kernel of ψ , since we constructed that isogeny to have kernel precisely $\langle Q \rangle$. Therefore $\psi(R)$ is in the kernel of $\widehat{\psi}$. By considering degrees, we see that $\deg \psi = \deg \widehat{\psi} = n$, so the kernel of $\widehat{\psi}$ can only have n points in it and must be exactly $\langle \psi(R) \rangle$. Then $\widehat{\psi}$ is the quotient-by- $\psi(R)$ isogeny, so by taking P to be $\psi(R)$ we achieve our goal. \square

Our construction shows that the μ -type subgroups of an elliptic curve are related to rational torsion on isogenous ones. In light of Mazur's classification of the possibilities for $\text{Tors}(E(\mathbb{Q}))$, it follows that $\# \text{Ker}(E/\mathbb{Q})$ can only be divisible by the primes 2, 3, 5, and 7. We also learn a basic fact that, while $\# \text{Ker}(E/\mathbb{Q})$ is finite for any specific E , it can grow arbitrarily large:

Proposition 2.2.4. *Let p be a prime number. There exists a number field K and an elliptic curve E/K with the property that $p \mid \# \text{Ker}(E/K)$.*

Proof. We need only to find a K -rational point on some elliptic curve, since the construction above will convert that into a μ_p on a quotient. To attain such a point, we could take an elliptic curve over \mathbb{Q} and base change to a field K where the p -rank of $\text{Tors}(E(K))$ is equal to one. The “other” subgroup of order p inside $E(K)$ will then be μ -type.² Now the Theorems 3.6 and 5.7 in González-Jiménez

²Note that if we base change to $\mathbb{Q}(E[p])$, μ_p will live in our base field, and there will be no difference between μ -type and rational. This is perhaps more convenient—why worry about a

and Najman in [22] says that there exists an elliptic curve E/\mathbb{Q} whose mod- p Galois representation has a small enough image that a point $P \in E[p]$ is defined over a number field smaller than $Q(E[p])$. This is the example we needed. \square

In fact, combing through the results of [22] slightly more carefully, a $\text{Ker}(E/K)$ divisible by p is attainable by a base-change from an elliptic curve defined over \mathbb{Q} to a number field of degree at most $\frac{p^2-1}{3}$. If $p \equiv 1 \pmod{3}$, we can do dramatically better and achieve this over a field of degree $2(p-1)$.

Non-example 2.2.5. Let E/\mathbb{Q} be an elliptic curve that has no rational 5-torsion, but attains a point P of order 5 over $\mathbb{Q}(\sqrt{5})$. Such a point must look like $(a, b\sqrt{5})$ with a and b rational. Since $\sqrt{5} = \zeta_5 - \zeta_5^2 - \zeta_5^3 + \zeta_5^4$, elements of the Galois group with cyclotomic character 1 and 4 will fix $\sqrt{5}$ while elements with character 2 and 3 will negate it. Thus the action of Galois on $\langle P \rangle$ is in fact determined by the *square* of the cyclotomic character: $\sigma(P) = (\chi(\sigma))^2 P$. Owing to our basic fact and the knowledge that $\chi^5 = \chi$, it follows that the Galois representation of E takes the form

$$\begin{bmatrix} \chi^2 & * \\ 0 & \chi^3 \end{bmatrix}.$$

In this case when we pass to the quotient $E/\langle P \rangle$, we get a point that transforms according to the *inverse* of the cyclotomic character. Thus this group is not μ -type: it is off by a twist. Nonetheless, P is defined over an abelian number field. This shows that the containment $\text{Tors}(E(\overline{\mathbb{Q}}))^\times \subseteq \text{Tors}(E(\mathbb{Q}^{\text{ab}}))$ can be strict. The Galois action when you can slaughter it?—but we can do better in terms of minimizing the degree of the number field needed.

subgroup $\langle P \rangle$ is, however, defined over \mathbb{Q} . So the requirement that E' possesses a rational point cannot be taken out of our construction.

According to Najman [44, Theorem 2c], there is exactly one such elliptic curve: Cremona's 50A3, with short Weierstrass model

$$y^2 = x^3 - 97875x + 14208750.$$

The 5-torsion point in question has coordinates $(255, 1080\sqrt{5})$.

The intuition you should take away from our theory and the above example is that there is an inverse relationship between the degree of definition of an abelian point and its membership in $\text{Ker}(E/K)$. The most extreme case is when P is rational, that is, defined over an abelian extension of K of smallest possible degree ($[K(P) : K] = 1!$). Then it is an étale cover of the quotient by P . On the other extreme, if P is defined over an abelian extension of K of largest possible degree, it possesses an étale cover. In the middle are curves like 2.2.5, which possess points defined over an abelian extension of K of “intermediate” degree. Then the quotient by such a point is also defined over an intermediate extension.

Each μ -type subgroup of E can be detected as the kernel of a rational isogeny emanating from E . This means if we seek to calculate $\text{Ker}(E/\mathbb{Q})$ we should start by looking for rational isogenies. Suppose E possesses a rational isogeny whose codomain E' possesses a rational point of order p . If we reduce the setup modulo a common prime ℓ of good reduction, then $\#E'(\mathbb{F}_\ell)$ will be a multiple of p . Since the number of points modulo ℓ is an isogeny invariant, $\#E(\mathbb{F}_\ell)$ will be a multiple of p as well. Thus, for all common primes of good reduction, $p \mid \#E(\mathbb{F}_\ell)$. According to

Katz, the converse holds as well:

Theorem 2.2.6 (Katz [31]). *Let E be an elliptic curve over a number field K , and let $m \geq 2$ be an integer. Suppose that there exists a set of primes of density one, all of whose members are primes of good reduction for E , and for which the congruence $\#E(\mathbb{F}_p) \equiv 0 \pmod{m}$ holds. Then there exists an elliptic curve E' , defined over K and K -isogenous to E , with the property that $\#\text{Tors}(E'(K)) \equiv 0 \pmod{m}$.*

2.3 Algorithm

Katz's result implies that we can try to detect rational isogenies by checking for common factors of $\#E(\mathbb{F}_\ell)$ as ℓ varies. This observation motivates the first step of the following algorithm.

Algorithm 2.3.1. *Given an elliptic curve E/\mathbb{Q} , this algorithm determines all elliptic curves E'/\mathbb{Q} possessing an isogeny $E' \rightarrow E$ whose kernel consists entirely of points defined over \mathbb{Q} .*

1. Calculate $\#E(\mathbb{F}_\ell)$ for the smallest fifty primes ℓ for which E has good reduction. Let d be the greatest common divisor of these cardinalities.
2. In what follows, assume p is a prime dividing d . For each such prime, calculate $\psi_p(x)$, the p -th division polynomial, and factor it.
3. If $\psi_p(x)$ has no factor of degree $\frac{p-1}{2}$, then p does not divide $\#\text{Ker}(E/\mathbb{Q})$. In what follows, assume $f(x)$ is a factor of ψ_p of degree $\frac{p-1}{2}$. For each such f , do steps 4 and 5.

4. Calculate the roots of f and let Q be a point of E whose x -coordinate is a root of f . Using Velu's formulas, calculate the isogeny emanating from E whose kernel is $\langle Q \rangle$.
5. Find the codomain of ϕ and check to see if it possesses a rational point of order p . If it does, remember ϕ .
6. Output all isogenies remembered in step 5.

Proof of correctness. As mentioned above, if a prime p merits our consideration it must divide $\#E(\mathbb{F}_\ell)$ for all primes ℓ . If a prime doesn't survive the first step, E won't have any rational isogeny of degree p , let alone one to a curve that has a rational point.

If ϕ is one of our desired isogenies, its kernel consists of points of order p , and therefore its kernel polynomial divides the p -th division polynomial. Since the kernel is supposed to be isomorphic to μ_p , a group with $p - 1$ nonidentity elements, its kernel polynomial will have degree $\frac{p-1}{2}$. This explains the statement in step 3. Once we've found candidate polynomials f , it remains to check whether they correspond to isogenies that possess all the correct properties. We do that in steps 4 and 5. Note step 5 can be accomplished quickly with the Lutz-Nagell theorem. \square

Some remarks are in order:

1. To calculate all of $\text{Ker}(E/\mathbb{Q})$ we should then check whether our p -covers have further p -covers. This is done by recursively calling the algorithm until nothing is found. Note that in these second rounds of running the algorithm we have

already decided on the prime p , so we can skip step 1. We are also guaranteed to find the isogeny dual to the one we discovered at the previous step. Since $\psi \circ \widehat{\psi} = [\deg \psi] = [p]$, we should only include that isogeny if the entirety of $E[p]$ is rational. For curves over \mathbb{Q} this is only possible when $p = 2$.

2. It may happen that E/\mathbb{Q} possesses a point of order p ; in this case p will certainly come up as a factor of every $\#E(\mathbb{F}_\ell)!$ This will also always happen in a “second round” detailed in the previous remark. In this case, step 1 does nothing to increase our confidence that E possess an étale cover of degree p . But the division polynomial will still check for a subgroup with the desired property—the only quirk here is that the division polynomial will possess many linear factors, corresponding to the rationally defined p -torsion.
3. If the ground field is not \mathbb{Q} then the action of the absolute Galois group on μ_n may be smaller than the action of $\text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$. This necessitates altering the degrees of the factors we examine in step 3. For example, if we are calculating $\text{Ker}(E/\mathbb{Q}(\sqrt{5}))$, then the Galois group acts on μ_5 via ± 1 only, so the question of whether $5 \mid \#\text{Ker}(E/\mathbb{Q}(\sqrt{5}))$ should be answered by searching for *linear* factors of $\psi_5(t)$ rather than quadratic ones.

Example 2.3.2. Take $K = \mathbb{Q}$ and E the curve $y^2 = f(x) = x^3 - 16875x - 9956250$. This curve has a cardinality which is a multiple of five over every prime p satisfying $10 < p < 1000$. Factoring the 5-division polynomial for E reveals a quadratic factor:

$$\psi_5(x) = (x^2 + 150x - 34875) (\text{degree } 5) (\text{degree } 5).$$

This quadratic polynomial is intriguing. Its discriminant is $2^4 \cdot 3^4 \cdot 5^3$, so it splits over $\mathbb{Q}(\sqrt{5})$. In that field its roots are $-75 \pm 90\sqrt{5}$. The quantity $f(-75 + 90\sqrt{5})$ is not a square, but extracting a further square root of this quantity gives us the field

$$M = \mathbb{Q} \left(\sqrt{-1822500 + 3645000\sqrt{5}} \right)$$

over which E attains a 5-torsion point. Asking Sage to apply Velu's formulas to quotient by this subgroup produces an isogeny to the curve $y^2 = x^3 - 2446875x + 1776093750$ which, lo and behold, is defined over \mathbb{Q} . As a sanity check, Sage confirms that this curve has 5 torsion points over \mathbb{Q} . Note also that $\text{Gal}(M/\mathbb{Q}) = \mathbb{Z}/4$, so the point we constructed in this process is indeed defined over an abelian number field.

We can reduce the entire construction modulo a prime of good reduction and achieve an admissible cover modulo that prime. As such, although the algorithm as stated calculates $\text{Ker}(E/\mathbb{Q})$, in fact it is calculating $\text{Ker}(E/\mathbb{Z}[\frac{1}{N}])$, where N is the product of all the primes of bad reduction. More work will need to be done to determine the situation at a prime of bad reduction. But in some cases, as in the next example, we can circumvent this issue.

Example 2.3.3. Let $K = \mathbb{Q}(\sqrt{29})$, $\epsilon = \frac{5+\sqrt{29}}{2}$ its fundamental unit, and E/K the elliptic curve with equation

$$y^2 + xy + \epsilon^2 y = x^3 - 5\epsilon^2 x - \epsilon^2 - 7\epsilon^4.$$

This curve comes from Kagawa [30]; it has discriminant ϵ^{14} , and so has good reduction everywhere over \mathcal{O}_K . The curve possesses no rational torsion over K . Checking

the cardinality of the reduction modulo the fifty smallest primes of \mathcal{O}_K , we discover all of them are divisible by 3. This prompts us to examine the 3-division polynomial of E :

$$\psi_3(x) = 3 \cdot \left(x + \frac{1}{3}\right) \cdot \left(x^3 + \frac{-243 - 45\sqrt{29}}{2}x - 9828 - 1825\sqrt{29}\right).$$

The point with x -coordinate $-\frac{1}{3}$ is defined over a quadratic extension of K . Sage calculates that the quotient isogeny that kills this point maps from E to the curve isomorphic³ to the one with equation

$$y^2 + xy + \epsilon^2 y = x^3.$$

This curve has K -rational torsion isomorphic to $\mathbb{Z}/3$, and so we deduce $\text{Ker}(E/K) \cong \mathbb{Z}/3$. In fact, as this curve has good reduction everywhere, we can promote this statement all the way down to the ring of integers:

$$\text{Ker}(E/\mathcal{O}_K) \cong \mathbb{Z}/3.$$

According to Zhao [57], there is actually one example over a smaller field: the curve

$$y^2 = x^3 + \left(-5134860 + 2096280\sqrt{6}\right)x + \left(-6324738336 + 2582063568\sqrt{6}\right)$$

has $\text{Ker}(E/\mathbb{Z}[\sqrt{6}]) = \mathbb{Z}/3$.

For $p = 5$, let $K = \mathbb{Q}(\sqrt{37})$ and $\epsilon = 6 + \sqrt{37}$ the fundamental unit. The curve

$$y^2 - \epsilon y = x^3 + \frac{3\epsilon}{2}x^2 - \frac{1669\epsilon + 139}{2}x - 7(5449\epsilon + 451)$$

³This curve was known to Tate; it was the first example written down of a curve with good reduction everywhere.

has good reduction everywhere and possesses a cover of order five, namely⁴

$$y^2 - \epsilon y = x^3 + \frac{3\epsilon + 1}{2}x^2 + \frac{11\epsilon + 1}{2}.$$

⁴This example is in Kagawa [29], wherein he notes that the covering curve in this case was known to Shimura. I believe this is the smallest example, but I cannot claim to have confirmed this. Comalada's tables [13] limit the possibilities for a smaller example to being over $\mathbb{Q}(\sqrt{d})$ for $d = 26, 29$, or 33 .

Chapter 3: $g = 1$, second way: modular polynomials

Herein we present a second plan based on the j -invariant. Suppose, as above, that we have high confidence (possibly by checking the divisibility of $\#E(\mathbb{F}_\ell)$ for many ℓ) that an elliptic curve E ought to possess an étale cover of degree p . Call the isogenous curve E' .

3.1 Modular polynomials

When two elliptic curves are related by a cyclic-isogeny, their j -invariants are related by a polynomial called the *modular polynomial*. This is a manifestation of the fact that the field of modular functions for $\Gamma_0(m)$ is generated over \mathbb{C} by $j(\tau)$ and $j(p\tau)$. See Cox [15, §11.B] for the theory behind the modular polynomial. We'll only need the polynomial for p prime, so our definition will be artificially limited.

For p a prime, let $C(p)$ denote the set of $p + 1$ matrices consisting of $\begin{bmatrix} 1 & a \\ 0 & p \end{bmatrix}$ for $0 \leq a \leq p - 1$, together with the one extra matrix $\begin{bmatrix} p & 0 \\ 0 & 1 \end{bmatrix}$. These are the matrices that move a point in the upper half-plane τ to the points corresponding to the lattices which are p -isogenous to $\langle 1, \tau \rangle$. (See §1.5 for a refresher.)

Definition 3.1.1. Let p be prime and $C(p)$ the set of matrices as above. The p -th

modular polynomial $\Phi_p(x, y)$ is the element of $\mathbb{Z}[x, y]$ satisfying the relationship

$$\Phi_p(x, j(\tau)) = \prod_{\sigma \in C(p)} (x - j(\sigma\tau)).$$

From the definition, the modular polynomial knows when two lattices are isogenous: when $\langle 1, \tau \rangle$ is p -isogenous to $\langle 1, \tau' \rangle$, then $\Phi_p(j(\tau), j(\tau')) = 0$. It is also clear from the definition that once we have plugged in $y = j(E)$, the resulting polynomial is monic in x .

The curve E' that we seek will have a j -invariant that is a rational root of the single-variable polynomial $\Phi_p(x, j(E))$. It's computationally feasible to produce this polynomial for, say, $p < 300$ (see Sutherland [53]). Then we can use the rational root theorem to search for candidate j -invariants for E' . It's possible to generate an elliptic curve from a j -invariant, so we might expect that this process will reconstitute a curve isogenous to E .

3.2 Twisting theory of elliptic curves

The elephant in the room is that the j -invariant only determines an elliptic curve up to isomorphism over $\overline{\mathbb{Q}}$. Over \mathbb{Q} there are many curves with the same j -invariant and only one sliver of them will be isomorphic over \mathbb{Q} to the curve we're searching for. The small glimmer of hope we can cling to is that two curves with the same j -invariant will be quadratic twists of each other, and as such we only need to decide which twist to take. The next two propositions make this precise and explain what we should look for.

Definition 3.2.1. Let E be an elliptic curve defined over a field k of characteristic

zero; suppose E has the Weierstrass model $y^2 = x^3 + a_4x + a_6$. The **quadratic twist of E by d** is the elliptic curve with model $dy^2 = x^3 + a_4x + a_6$.

Equivalently, by multiplying through by d^3 and scaling x and y , we can say the quadratic twist has equation $y^2 = x^3 + d^2a_4x + d^3a_6$. Recall that isomorphisms of short Weierstrass equations are enacted by a change of variables of the form $x \mapsto u^2x$, $y \mapsto u^3y$; this produces a new cubic polynomial of the form $x^3 + u^4a_4x + u^6a_6$. As such, a quadratic twist is “half” of an isomorphism. Twisting twice by the same d is the same as performing an isomorphism (i.e., doing nothing). This also means that we may limit our twisting constant to the classes of $\mathbf{k}^\times/(\mathbf{k}^\times)^2$.

Proposition 3.2.2. *Let E and E' be two elliptic curves both defined over \mathbb{Q} . Suppose $j(E) = j(E')$, with $j(E) \neq 0, 1728$. Then E and E' are related by a quadratic twist.*

Proof. Suppose E has the short Weierstrass model $y^2 = x^3 + Ax + B$ and E' has the model $y^2 = x^3 + Cx + D$. Our assumption that the common j -invariant isn't 0 or 1728 implies A , B , C , and D are all nonzero. Since E and E' are isomorphic over $\overline{\mathbb{Q}}$, there exists a $u \in \overline{\mathbb{Q}}^\times$ satisfying $A = u^4C$ and $B = u^6D$. The quotient $\frac{B}{A} = u^2\frac{D}{C}$ is in \mathbb{Q} , implying $u^2 \in \mathbb{Q}$. Then B is the quadratic twist of A by u^2 . \square

Note that $j \neq 0, 1728$ corresponds to $\text{Aut}(E) = \pm 1$, that is, E has “no extra endomorphisms.”¹ We'll come back to this idea later in genus two.

¹A curve with $j = 0$ defined over \mathbb{Q} has the model $y^2 = x^3 + Ax$ for some $A \in \mathbb{Z}$. These curves can be quartically twisted to the curve $y^2 = x^3 + dAx$. Similarly, a curve with $j = 1728$ defined over \mathbb{Q} has the model $y^2 = x^3 + B$ for some $B \in \mathbb{Z}$. This curve can be sextically twisted to the

Proposition 3.2.3. *Let $E : y^2 = x^3 + Ax + B$ be an elliptic curve with good reduction at a prime $p \geq 5$, and suppose d is a squarefree integer such that $p \mid d$. The quadratic twist of E by d has additive reduction at p . Conversely, if $p^6 \nmid \Delta$, then the twist of E by d has good reduction at p if and only if p divides both A and B .*

Recall from Proposition 1.2.3.1 that when E is in short Weierstrass form, if a prime divides Δ and the coefficients A and B the reduction is additive at that prime.

Proof. Write $d = pr$ with $(p, r) = 1$. Assume that the Weierstrass equation for E has been chosen to be globally minimal. By this I mean that the discriminant of E , call it $\Delta = 4A^3 + 27B^2$, cannot be made smaller by any change of variables. The minimal discriminant is only divisible by primes of bad reduction for E . Our assumption implies $p \nmid \Delta$, meaning p doesn't divide at least one of A or B . Following the above, the quadratic twist has Weierstrass equation $y^2 = x^3 + p^2r^2Ax + p^3r^3B$. The discriminant of the twist is $p^6r^6\Delta$ which is divisible exactly by p^6 . This factor can't be removed by an isomorphism (isomorphisms change the discriminant by a twelfth power), so p is a prime of bad reduction for the twist. Since the twist becomes isomorphic to E over $\mathbb{Q}(\sqrt{d})$, the bad reduction must be potentially good, implying it's additive.

For the converse, start with the fruitful case: $p \mid A$ and $p \mid B$. Write $A = p^k A'$ and $B = p^m B'$. The discriminant of the curve is

$$\Delta = 4p^{3k}(A')^3 + 27p^{2m}(B')^2.$$

curve $y^2 = x^3 + dB$.

The first term has a valuation which is a multiple of 3 while the second term has a valuation which is a multiple of 2. As such to get p^6 out of the sum we must have $p^2 \mid A$ and $p^3 \mid B$, with one of these divisibilities being exact. Suppose that $p^2 \parallel A$ (the other option proceeds similarly). The discriminant then factors as

$$p^6 (4(A')^3 + 27p^{2(m-3)}(B')^2) = p^6 \Delta',$$

with $p \nmid \Delta'$. The twist has equation $y^2 = x^3 + p^4 A'x + p^{m+3} B'$. Performing the change of variables $x' = p^2 x$, $y' = p^3 y$ leads to a model with discriminant Δ' , so the twist has good reduction at p .

In the upsetting case, $p^6 \parallel \Delta$ but p divides neither A nor B . The quadratic twist of E will have a model with $p^{12} \mid \Delta$, but only p^2 and p^3 dividing its coefficients a_4 and a_6 . An isomorphism of E with itself will alter these coefficients by a fourth and sixth power respectively, so this model for the twist is in fact minimal. (You can't clear p from A or B , hence you cannot decrease the p -adic valuation of Δ with an isomorphism.) The factor of p^{12} in the discriminant is essential, and the twist has bad reduction at p . □

We should make a few remarks here as this result was unexpectedly delicate.

1. Our proof made use of the fact that \mathbb{Q} has class number one, so that E had a global minimal model. This isn't strictly necessary: one could invert primes away from p in order to kill all the ideal classes, or start from an equation that's only minimal at p , since the p -adic valuation at Δ is the point of contention.
2. Twisting is not a panacea for bad reduction. Typically twisting takes you

from a curve with bad reduction to another curve with bad reduction. As an example, both $y^2 = x^3 + 5x + 5$ (discriminant $-2^4 \cdot 5^2 \cdot 47$) and its quadratic twist by $d = 5$ with equation $y^2 = x^3 + 5^3x + 5^4$ (discriminant $-2^4 \cdot 5^8 \cdot 47$) have bad reduction at $p = 5$. Both of these equations are globally minimal and the reduction at 5 for each curve is additive. This demonstrates that the conditions outlined in the converse are needed.

3. The bad case of the converse does happen. As an example, $y^2 = x^3 + 7x + 792$ (discriminant $-2^6 \cdot 5^6 \cdot 271$) has a factor of 5^6 in its discriminant which cannot be subdued with a twist by 5. (Of course, by sight 5 doesn't divide either of the coefficients 7 or 792.)
4. In writing the converse we really should consider Δ up to twelfth powers: p^6 is not the only power of p that could be remedied by a twist. The case of higher p -adic valuations is typically a mashup of the two cases. In the notation of the proof, this would manifest itself as a factor of p^{12} dividing Δ' , and these factors of p^{12} may either be essential or be removable by an isomorphism as in the previous remark. So really the condition should be $p^{6+12k} \parallel \Delta$, and the first step should be to pass to an equation minimal at p . If the valuation of the minimal discriminant at p is 6, then proceed as in the proof.

3.3 Algorithm

Owing to the fact that twisting twice by the same constant will return you to the curve you started with, this proposition gives us some clues as to what primes

will divide the appropriate twisting constant. This insight leads to the following algorithm.

Algorithm 3.3.1. *Let E/\mathbb{Q} be an elliptic curve without complex multiplication. The following steps will calculate the model of an elliptic curve defined over \mathbb{Q} and rationally isogenous to E via an isogeny of degree p .*

- 0 *Become confident that E possesses an étale cover of degree p .*
1. *Calculate $j(E)$ and use it to create the polynomial $f(x) = \Phi_p(x, j(E))$.*
2. *Use the rational root test to search for rational roots of $f(x)$. If none exist, no such cover exists; output nothing.*
3. *Let j_0 be a root discovered in the previous step. Generate an elliptic curve E_j defined over \mathbb{Q} which has j_0 as its j -invariant.*
4. *Factor the discriminants $\Delta(E)$ and $\Delta(E_j)$. Let d be the product of every prime that divides one of these two discriminants but not the other.*
5. *Twist E_j by d . If $E_j^{(d)}$ is an étale cover of E of degree p , output $E_j^{(d)}$. (To do this, check if it possesses a rational point of degree p and if the quotient by that point is isomorphic to E .)*
6. *Systematically traverse every product d of primes of bad reduction for E . For each, test if $E_j^{(d)}$ is an étale cover of E of degree p . If so, output $E_j^{(d)}$.*

Proof of correctness. By the propositions above, if j_0 is a rational root of $f(x)$ then an isogenous curve E' exists, and that curve will have j -invariant j_0 . Since E and

E' are isogenous, they share the same reduction type at every prime. A curve with the same j -invariant as E' must be a quadratic twist of E' . Say the twist is by d .

If p is a prime of good reduction for E that divides d , then E' will have good reduction at p while E_j will not. Such a prime will then be a divisor of the discriminant of E_j but will not divide the discriminant of E . These primes are tallied in step 4. We need to do slightly more work than this, though—it's possible that the twist from E' to E_j included twisting at a prime of bad reduction for E' which becomes good after a twist. In that case p will divide $\Delta(E)$ but not $\Delta(E_j)$. We include those primes in our twisting constant d as well. The work above shows that the twist we generate in step 5 will have the correct j -invariant and the correct primes of bad reduction, and hence the correct conductor. It therefore could be the curve we are searching for.

If it is not the curve we sought, it is because the “mystery twist” included factors that were primes of bad reduction for E . We traverse every possibility in step 6. The isogenous curve must be among the set of curves considered. \square

We avoided the situation of having to twist curves with $j = 0$ or $j = 1728$ in a heavyhanded way by decreeing that our starting curve cannot have complex multiplication. Those cases cannot come up if $\text{End}(E) \otimes \mathbb{Q} \not\cong \mathbb{Q}(i)$ or $\mathbb{Q}(\mu_3)$.²

This algorithm has the advantage of never leaving the rational numbers. On the other hand, it involves calculating some scary objects: Φ_p is notoriously monstrous, and we're being asked to factor potentially massive numbers in step 4. Nev-

²I believe that the algorithm works in these cases if you replace “quadratic twist” in step 5 with “quartic twist” or “sextic twist” respectively, but am not confident in this.

ertheless, for typical inputs this algorithm is practical. In the situation described above where the ground field is \mathbb{Q} , the list of possible isogeny degrees is small, and Φ_p has already been calculated for every such degree. So we can look the polynomial up instead of computing it.

The curve with a given j -invariant for step 3 also can be computed from a formula:

$$E_j : y^2 + xy = x^3 - \frac{36}{j - 1728} x - \frac{1}{j - 1728}.$$

This will typically be the “start” curve of Step 3.

Example 3.3.2. Let E be the elliptic curve $y^2 = x^3 - 1323x - 42714$, with j -invariant $-\frac{2401}{6}$ and discriminant $-2^{13} \cdot 3^{13} \cdot 7^2$. At each of the 21 two-digit primes $E(\mathbb{F}_p)$ is a multiple of 7, so we have reason to believe that a curve 7-isogenous to E may be out there somewhere. Sage calculates the polynomial $\Phi_7(x, j(E))$ without difficulty:

$$\begin{aligned} \Phi_7(x, j(E)) = & x^8 + \frac{213040394078304357485519137}{279936} x^7 + \frac{17661906410035531644497191135081212918101}{46656} x^6 \\ & + \frac{1410555319940980331987073409241420228562791475007867211}{23328} x^5 \\ & - \frac{1152523601114246426547158277599677214082609530047623552106759}{8748} x^4 \\ & + \frac{3606479300246324042807156700608374611249151304897342977760220202403}{46656} x^3 \\ & + \frac{20857846576370879649306663336267794477998841073858975391514312896705021}{11664} x^2 \\ & - \frac{21509466497902086531812835080145216777088176888375619725536996617023937151}{34992} x \\ & + \frac{392517463848357692086411098083947242717793357417123038753321102313521034771201}{1679616}. \end{aligned}$$

This polynomial has one rational root, $j_0 = -\frac{6329617441}{279936}$. Let E_j be the elliptic curve from the formula above which has this j -invariant. For simplicity, I’ll complete

the square and give E_j a short Weierstrass form; we get the model

$$y^2 = x^3 - \frac{170899670907}{6813346849}x + \frac{341799341814}{6813346849}.$$

The discriminant of E_j is

$$-1 \times 2^{19} \times 3^{19} \times 7^2 \times 197^{-6} \times 419^{-6} \times 967^6,$$

which suggests twisting E_j by $197 \times 419 \times 967$ since those primes of bad reduction are unique to E_j . That curve is

$$y^2 = x^3 - 159806402368755723x + 25511200349980609846701126.$$

Sage confirms that this curve is isogenous to the original E . It has a superfluous factor of 967^{12} in its discriminant;³ after an isomorphism we get a minimal model for E' , $y^2 = x^3 - 182763x + 31201254$.⁴

3.4 Alternative method to find $j(E')$: via \mathbb{C}

There is an alternative to using the polynomial $\Phi_m(X, Y)$ for large m . As noted, these polynomials become unruly as m grows, and the operations we would like to do with them may prove implausible if we can't even get the polynomial into working memory! As such in this section we present an alternative to acquiring the j -invariant of the isogenous curve, based on passage through the complex numbers.

³This could have been avoided if we twisted by $\frac{197 \cdot 419}{967}$ instead. Since twisting by p alters the discriminant by a factor of p^6 , we could see this factor of 967^{12} coming and we could have taken evasive maneuvers.

⁴ E has Cremona label 294b1 and E' has Cremona label 294b2.

Suppose the j -invariant of E is written as $\frac{m}{n}$ in lowest terms. Then the polynomial $\Phi_p(x, \frac{m}{n})$ we considered in the previous section will be monic in x of degree $p + 1$ and have a constant term whose denominator is at worst n^{p+1} . The rational roots of this polynomial will have denominator bounded by n^{p+1} , so if we have some method for calculating them to precision finer than $\frac{1}{n^{p+1}}$, we can round off to determine them as exact rational numbers.

Such a method is possible via passage into the complex numbers. Recall the following facts from §1.5. E can be represented by the quotient of the complex numbers by a lattice Λ with a basis $\{1, \tau\}$ for a $\tau \in \mathbb{C}$ which lives in the typical fundamental domain for the action of $\mathrm{SL}(2, \mathbb{Z})$ on the upper half-plane. The p -torsion points of $E(\mathbb{C})$ are the cosets of Λ with representatives $\frac{1}{p}(a\tau + b)$ for $0 \leq a, b \leq p-1$. There are $p + 1$ different subgroups of $E[p]$ of order p . One is generated by $\frac{1}{p}$ and the other p of them are generated by $\frac{\tau+i}{p}$ for $0 \leq i < p$. The $p + 1$ curves that are p -isogenous to E are the quotients of \mathbb{C} by the $p + 1$ lattices with bases $\left\{\frac{1}{p}, \tau\right\}$ and $\left\{1, \frac{\tau+i}{p}\right\}$ for $0 \leq i < p$. Note that these lattices are not guaranteed to be in standard form as written. Also note the first of these is homothetic to the lattice with basis $\{1, p\tau\}$. It's true that the numbers in the upper half-plane may not live in the fundamental domain, though if desired we could use the action of $\mathrm{SL}(2, \mathbb{Z})$ to move them there. This won't be strictly necessary in what follows.

As suggested in [6] we can calculate the j -invariant of a lattice efficiently by using the eta-quotient formula

$$j(\tau) = \left(\frac{(\eta(\frac{\tau}{2})/\eta(\tau))^{24} + 16}{(\eta(\frac{\tau}{2})/\eta(\tau))^8} \right)^3.$$

So if we have a method to calculate τ from the equation of E to arbitrary precision, we can then apply the centered equation to the values from the previous paragraph to estimate the j -invariants of the curves that are p -isogenous to E .

Where does τ come from? The lattice corresponding to E come from evaluating “periods,” that is, integrals of invariant differentials, around paths forming a basis for the homology group $H_1(E, \mathbb{Z})$. Algorithm 7.4.7 in Cohen [12], reproduced below for the situation of interest here, describes how to calculate these complex periods ω_1 and ω_2 of E in terms of the arithmetic-geometric mean of the roots of the cubic polynomial defining E . These periods form a basis for a lattice homothetic to Λ . Calculating the AGM can be done to arbitrary precision (Sage has a built-in function for this, which I’ll freely use in the example to come). From these periods we can get our hands on τ .

Algorithm 3.4.1. Given a Weierstrass equation of the form $y^2 = x^3 + Ax + B$ with $A, B \in \mathbb{Q}$, this algorithm computes a basis $\{\omega_1, \omega_2\}$ of the period lattice for E . The basis created will have $\omega_2 \in \mathbb{R}$, $\text{Im}(\omega_1/\omega_2) > 0$, and $\text{Re}(\omega_1/\omega_2) = 0$ or $-\frac{1}{2}$.

- (a) Calculate Δ , the discriminant of E . Depending on whether Δ is positive or negative, go to step (b) or (c).
- (b) If $\Delta > 0$, let e_1, e_2 , and e_3 be the three roots of $f(x)$, ordered so that $e_1 > e_2 > e_3$. Return the quantities

$$\omega_2 = \frac{\pi}{\text{AGM}(\sqrt{e_1 - e_3}, \sqrt{e_1 - e_2})} \quad \text{and} \quad \omega_1 = \frac{\pi\sqrt{-1}}{\text{AGM}(\sqrt{e_1 - e_3}, \sqrt{e_2 - e_3})}$$

- (c) If $\Delta < 0$, let e_1 be the unique real root of f . Set $a = 3e_1$ and $b = \sqrt{3e_1^2 + A}$.

Return the quantities

$$\omega_2 = \frac{2\pi}{\text{AGM}(2\sqrt{b}, \sqrt{2b+a})} \quad \text{and} \quad \omega_1 = -\frac{\omega_2}{2} + \frac{\pi\sqrt{-1}}{\text{AGM}(2\sqrt{b}, \sqrt{2b-a})}$$

Some comments on this algorithm:

1. The value of τ we want is ω_1/ω_2 , since the lattice with basis $\{\omega_1, \omega_2\}$ is homothetic to the lattice with basis $\{1, \omega_1/\omega_2\}$.
2. Owing to the fact that the j -invariant has a power series expansion in terms of $q = \exp(2\pi i\tau)$, j is real-valued on the lines $\text{Im}(\tau) = 0$ and $\text{Im}(\tau) = -\frac{1}{2}$. On those two lines q is real and positive or real and negative respectively, and as the power series for j has all positive integer coefficients it follows that j is real and positive or real and negative on those two lines respectively. Since the input to the algorithm was an elliptic curve over \mathbb{Q} , it makes sense that the value of τ that comes from this algorithm will live on one of those two lines.
3. This algorithm isn't guaranteed to give us a τ in the fundamental domain. The function $j(\tau)$ takes the arc between $\exp(\frac{2\pi i}{3})$ and $\exp(\frac{\pi i}{4})$ to the real interval $[0, 1728]$. So if we input an elliptic curve whose j -invariant lies between those two real numbers we must get something outside of the fundamental domain. (It will be "below" the fundamental domain, inside the unit circle.)
4. If one cares about isogenies of high degree, they'll necessarily be led to considering elliptic curves over number fields larger than \mathbb{Q} . Nothing changes in the above algorithm if we were to ask for E to be defined over a subfield of

\mathbb{R} instead of \mathbb{Q} . In fact step (b) of the algorithm can be made to work if E is defined over a non-real subfield of \mathbb{C} , though caution is needed to deal with the issue of the arithmetic-geometric mean being multi-valued. The precision needed should be determined in light of the fact that the j -invariants sought will be elements of this number field.

Example 3.4.2. Let's revisit the curve $y^2 = x^3 - 1323x - 42714$ from the previous section. This curve had j -invariant $\frac{-2401}{6}$, so when searching for the j -invariant of the curve which is 7-isogenous to E we will need at precision of at least 6^8 . As $6^8 \approx 10^{6.23} \approx 2^{20.68}$, seven digits or 21-bits of precision past the decimal point will suffice. In this example I worked to 50 digits.

Following Cohen's algorithm for the periods, the unique real root of f is $e_1 = 47.2010357\dots$. The quantities a and b are $a = 141.6031071\dots$ and $b = 73.2175752\dots$, leading to the periods

$$\omega_2 = 0.3686775\dots \quad \text{and} \quad \omega_1 = -0.1843387\dots + 0.4030362\dots i.$$

Dividing gives us $\omega_1/\omega_2 = \tau = -\frac{1}{2} + 1.0931945\dots i$. As a sanity check, plugging τ into our eta quotient gives us $j(\tau) = -400.1666667$ plus a vanishingly small imaginary part which can be attributed to rounding error. Now we calculate the j -invariants for the curves 7-isogenous to E over \mathbb{C} :

$$\begin{aligned} j\left(\frac{\tau}{7}\right) &= -248712190019162.57933688 - 132646662910561.41711712i \\ j\left(\frac{\tau+1}{7}\right) &= -248712190019162.57933688 + 132646662910561.41711712i \\ j\left(\frac{\tau+2}{7}\right) &= 1100561.13579798 + 340843.890361726i \end{aligned}$$

$$\begin{aligned}
j\left(\frac{\tau+3}{7}\right) &= 172.09442973 + 314.88559198i \\
j\left(\frac{\tau+4}{7}\right) &= -22610.94479084 + 0.00000000i \\
j\left(\frac{\tau+5}{7}\right) &= 172.09442973 - 314.88559198i \\
j\left(\frac{\tau+6}{7}\right) &= 1100561.13579798 - 340843.890361726i
\end{aligned}$$

$$j(7\tau) = -761032003141124103427.30178156 + 0.00000000i$$

The polynomial $\Phi_7(x, j(E))$ had two real roots.⁵ Multiplying $j\left(\frac{\tau+4}{7}\right)$ by 6^8 makes it look like an integer:

$$6^8 \cdot j\left(\frac{\tau+4}{7}\right) \approx -37977704646.000000000000 \dots,$$

while the second one does not look like an integer

$$6^8 \cdot j(7\tau) \approx 1.278241528987882302102150909.15150771669 \dots.$$

The former j -invariant is $-\frac{6329617441}{279936}$ which is the one calculated earlier, so it will lead to the isogenous curve via the procedure of the previous section. As the other one is not rational to within our tolerance, we conclude it isn't an element of \mathbb{Q} at all.

⁵We could have foreseen this: it has degree eight and the j -invariant we're searching for is rational, so there must be a second real root.

Chapter 4: Middlebrow arithmetic of curves of genus two

In §5 we'll demonstrate that the method of §3.4 can be made to work in genus two. This chapter is devoted to the arithmetic of Jacobians that we'll need. We'll also indicate what the theory of Jacobians over \mathbb{C} looks like.

4.1 Genus two curves and their Jacobians

A curve of genus two is a curve whose complex points are topologically a surface with genus two. It is a fact stemming from the Riemann-Roch theorem that a curve of genus two will always be the zero locus of an equation of the form $C : y^2 = f(x)$, where f is either degree five or six without repeated roots. (See for example [9, §1.3].¹) We say C is defined over a field k if $f(x) \in k[x]$.

Two warnings are in order.

1. The curve is **not** the projective closure of the graph of $y^2 = f(x)$ in \mathbb{P}_k^2 . That

¹This is the equivalent of “short Weierstrass form.” The “long Weierstrass form” looks like

$$y^2 + h(x)y = f(x)$$

for a polynomial f of degree five or six and a polynomial h of degree not exceeding three. Transferring from long Weierstrass form to short Weierstrass form cannot be done over a field of characteristic 2, 3, or 5; we'll avoid those.

object has a single point $[0 : 1 : 0]$ at infinity, which is necessarily singular. A nonsingular model for C can be created by blowing up this singular point. This amounts to embedding C into \mathbb{P}^3 via $(x, y) \mapsto [1 : x : y : x^2]$. Note that in the patch of \mathbb{P}^3 where the first homogeneous coordinate is nonzero, the information in the final coordinate is redundant, and we can pretend that we're only looking at the affine coordinates (x, y) . As such we often work as though we are in \mathbb{A}^2 , although under the hood we are secretly in a 3-dimensional open set inside \mathbb{P}^3 .

Suppose $f(x)$ has leading coefficient a_5 if it is degree five and a_6 if it is degree six. In the former case, the singularity at infinity is a cusp, and so there is still one point at infinity $[0, 0, 0, a_5]$ on the model in \mathbb{P}^3 . If f is degree six the singular point in \mathbb{P}^2 is a node and so there are two points $[0 : 0 : 0 : \pm\sqrt{a_6}]$ at infinity in \mathbb{P}^3 . These points are sometimes called ∞^+ and ∞^- .

This warning has a sub-warning: the single point at infinity is always \mathbf{k} -rational (i.e., a member of $C(\mathbf{k})$) if f is quintic. The two points at infinity are not necessarily \mathbf{k} -rational if f is sextic: they are if and only if the leading coefficient a_6 is a square.

2. For the purposes of algebraic geometry, i.e., when $\mathbf{k} = \bar{\mathbf{k}}$, we can always assume f has degree five. But if we work in the setting where \mathbf{k} is not algebraically closed, we can transform the degree-six case into the degree-five case only when the degree six polynomial possesses a root $e_1 \in \mathbf{k}$. In that case,

$$y^2 = (x - e_1)g(x)$$

with g of degree five can be birationally mapped to the model

$$v^2 = u^5 g \left(e_1 + \frac{1}{u} \right)$$

by using the transformation

$$u = \frac{1}{x - e_1}, \quad v = \frac{y}{(x - 1)^3}.$$

On the other hand, if f is of degree five, we can always promote it to degree six by using any transformation of the shape $u = \frac{1}{x-a}$, $v = \frac{y}{(x-a)^3}$, so long as $f(a) \neq 0$.

As you may have inferred from the above warning, two equations determine the same curve when the variables are related by fractional linear transformations of the form

$$x = \frac{aX + b}{cX + d} \quad y = \frac{eY}{(cX + d)^3},$$

where $ad - bc \neq 0$ and $e \neq 0$ (this guarantees that the transformations are invertible). The curve $y^2 = f(x)$ has good reduction at a prime if the equation makes sense mod p ; this occurs when $p \nmid \text{disc}(f)$, i.e., f has simple roots mod p —this is the same as the genus 1 situation.

The points of C cannot be made into a group the way we did for the points on an elliptic curve. The issue is that a line will intersect the curve in six points, which is too many to impose a binary operation based on collinearity.² Instead, the

²From a philosophical perspective, it's a good thing that this doesn't work. The groups that we would get would be very small, due to Faltings's theorem [19] that a curve of genus higher than one has only finitely many points with coefficients in any fixed number field.

numerology suggests that *pairs* of points on C might assemble themselves into an abelian group. A cubic polynomial $y = c(x)$ will intersect the graph of $y^2 = f(x)$ in six points. If four points in the plane are specified, a unique cubic runs through them; as such the remaining two points of intersection of the cubic with the curve will determine a new pair of points. We can think of this cubic polynomial as being determined by two pairs of points of C rather than four individual ones. Doing this shows that we have a binary operation that takes in two pairs of points and returns a third pair of points. (An example follows—see below.) Since the operation is agnostic to the order of the points within each pair, the group operation is defined on the symmetric square $(C \times C)/\{(P, Q) \sim (Q, P)\}$. As usual we treat the pair $\{P, P\}$ as imposing the condition that the cubic polynomial meets the curve C to multiplicity two at P , and intersections at infinity³ are actually determined in the model in \mathbb{P}^3 . Much as for elliptic curves, three pairs of points sum to zero when they all lie on the graph of a cubic.

There's one issue that needs to be overcome: just as with elliptic curves, if $(x, y) \in \mathbf{k}(C)$ then $(x, -y) \in \mathbf{k}(C)$ as well, and it's typically not possible for a the graph of a function to go through two points with the same x -coordinate. As such we have to exclude pairs of the form $\mathfrak{D} = \{(x, y), (x, -y)\}$ from our domain. This amounts to *blowing down* the set of pairs in \mathfrak{D} ; they determine a copy of C sitting diagonally inside the symmetric square. In the case $\deg f(x) = 5$, $\{\infty, \infty\}$ counts

³The short version is that if f is degree five, then $y = c(x)$ of degree three intersects $y^2 = f(x)$ at infinity precisely when the leading coefficient of c vanishes, i.e., c is actually quadratic. It intersects $y^2 = f(x)$ at infinity to multiplicity two when c is actually linear.

as a point together with its “negative”; when $\deg f(x) = 6$, $\{\infty^+, \infty^-\}$ is the pair that is on the offending copy of C . Blowing this curve down amounts to imposing an equivalence relation whereby all of the pairs of points in \mathfrak{D} are collapsed to one point, leaving everything else untouched. This common point is the identity element of the group. We can take $\{\infty, \infty\}$ as a representative for it when f is quintic and $\{\infty^+, \infty^{-1}\}$ when f is sextic. With this definition, the negative of a pair $\{(a, b), (c, d)\}$ is $\{(a, -b), (c, -d)\}$ since those four points join together to create two elements of \mathfrak{D} . The group created in this way is called the **Jacobian** of C .

Toy example 4.1.1. Let C be the hyperelliptic curve

$$y^2 = f(x) = (x - 3)(x - 1)(x + 2)(x + 4)(x + 6).$$


We’ll add the pairs $\{(-6, 0), (-4, 0)\}$ and $\{(0, 12), (3, 0)\}$ on $\text{Jac}(C)$. We first need the cubic polynomial passing through all four of these points. Typically this requires Lagrange interpolation, but since three of the four points are on the x -axis we know the polynomial must look like $c(x) = \lambda(x + 6)(x + 4)(x - 3)$ for some choice of λ . Owing to the fact that $c(0) = 12$ we get $\lambda = \frac{1}{6}$. Next we search for the other points of intersection of $y = c(x)$ with $y^2 = f(x)$. This is accomplished by plugging $c(x)$ in for y , i.e., solving $c(x)^2 - f(x) = 0$. This will tell us the six x -coordinates of the points of intersection. We know ahead of time this polynomial will be divisible by $(x + 6)(x + 4)x(x - 3)$. Factoring, we get

$$c(x)^2 - f(x) = \frac{1}{36} (x + 6)(x + 4)x(x - 3)(x^2 - 29x + 42),$$

with roots $\alpha_1 = \frac{29 - \sqrt{1009}}{2} \approx -1.3824$ and $\alpha_2 = \frac{29 + \sqrt{1009}}{2} \approx 30.3823$. Plugging into $c(x)$ recovers the y -coordinates. Remembering that the third pair of intersections

are the negative of the sum, we get that

$$\{(-6, 0), (4, 0)\} + \{(0, 12), (3, 0)\} = \left\{ (\alpha_1, -2850 + 90\sqrt{1009}), (\alpha_2, -2850 - 90\sqrt{1009}) \right\}.$$



(The other point of intersection is far to the right of the pictured window.)

This example illustrates two truths: first, the obvious fact that the coordinates of the sum of two points can be much more complicated than the points we started with. Second, the coordinates of the points comprising the sum will either be rational or conjugate over a quadratic extension of the field of definition for C .⁴ This means that $\text{Jac}(C)(\mathbb{Q})$ consists of “rational pairs”, that is, Galois orbits of size two. This is different from “pairs of rational points.” We’ll see this phenomenon in a different guise in §3 when we work over a finite field.

Note that the Jacobian makes sense as long as the intersection theory of cubics with the curve C makes sense, which only requires that C be non-singular. In particular, we infer that the Jacobian will have good reduction mod p whenever C

⁴The x -coordinates satisfy a quadratic equation with coefficients in $\mathbb{k}[x]$, making them conjugates. The y -coordinates are values of $c(x) \in \mathbb{k}[x]$ at conjugate points, hence are conjugate as well.

does. (This hints at the deeper fact that C and $\text{Jac}(C)$ comprise the same amount of information: one determines the other. This is known as Torelli's theorem [40, § 12].)

4.2 The Riemann-Roch formula; Jacobians as groups of divisors

In slightly higher language, the Jacobian can be understood as a group of divisor classes of C . In this section we'll state the version of the Riemann-Roch theorem causing this phenomenon.

The notation is as follows. A **divisor** D is a formal sum of points on C ; the sum of its coefficients is its **degree** $\deg D$. It is called **effective** if all its nonzero coefficients are positive. If $f \in k(C)$ is a rational function defined on C , the **principal divisor** $\text{div } f$ associated to it is $\sum_{P \in C} \text{ord}_f(P)[P]$, where ord_f counts the number of zeros or poles of f at P with multiplicity. Principal divisors have degree zero. Two divisors are **linearly equivalent** if their difference is principal. The group of divisors modulo linear equivalence is the **Picard group** $\text{Pic}(C)$. Inside $\text{Pic}(C)$ is the subgroup $\text{Pic}^0(C)$ consisting of divisor classes of degree zero.

Definition 4.2.1. The **Jacobian** of a curve C is the group $\text{Pic}^0(C)$.⁵

For a divisor D , $\mathcal{L}(D)$ is the vector space

$$\mathcal{L}(D) = \{f \in k(C) : \text{div } f + D \text{ is effective}\},$$

and the dimension of this vector space is $\ell(D)$. The number $\ell(D)$ is invariant under linear equivalence. Since scaling functions doesn't change their divisors, the projec-

⁵Cognoscenti will realize I'm abusing the self-duality of the Jacobian here. Since we won't leave the world of curves I'm taking liberties.

tivization of the vector space $\ell(D)$ is the set of effective divisors linearly equivalent to D . A **canonical divisor** K is the divisor of a nonzero differential on C . Canonical divisors are linearly equivalent and have degree $2g - 2$, where g is the genus of the curve.

Theorem 4.2.2 (Riemann-Roch). *Let C/k be a curve of genus two over a field k .*

Then

$$\ell(D) - \ell(K - D) = \deg D - 1.$$

In particular, suppose $\deg D = 2$. Then $\ell(D) = 1$ except when D is linearly equivalent to K , in which case $\ell(D) = 2$.

Proof. The Riemann-Roch theorem proper is in Hartshorne [23, Theorem IV.1.3]. For the second statement, $K - D$ will have degree zero. The only effective divisor of degree zero is $\text{div}(1)$, the empty sum. Since $\text{div}(f) + (K - D)$ will always be degree zero, there are elements in $\mathcal{L}(K - D)$ precisely when $K - D$ is linearly equivalent to zero, i.e., D is linearly equivalent to K . In that case $\ell(K - D) = 1$ and so the Riemann-Roch formula produces $\ell(D) = \ell(K) = 2$. Otherwise $\mathcal{L}(K - D)$ is empty, so the Riemann-Roch formula produces $\ell(D) = 1$. \square

Since the projectivization of a one-dimensional vector space is a point, it follows that if D is a divisor of degree two not in the canonical class, there is a unique effective divisor linearly equivalent to D . From now on we may conflate a noncanonical divisor class D of degree two with the unique effective divisor in that class. The map $D \mapsto D + K$ exhibits a bijection between $\text{Pic}^0(C)$, our Jacobian, and the set $\text{Pic}^2(C)$ of divisor classes of degree two. We impose a group structure on $\text{Pic}^2(C)$

by transporting the structure inherent in $\text{Pic}^0(C)$. The advantage of working with $\text{Pic}^2(C)$ is that every class has a canonical representative, namely the unique effective divisor in that class. The lone exception is the canonical class, which is the identity element of the group (note $0 \in \text{Pic}^0(C)$ mapped to $K \in \text{Pic}^2(C)$). Some caution is needed because many effective divisors live in the canonical class; these are all the same point of the Jacobian since they are all linearly equivalent.

When C is presented as the model $y^2 = f(x)$ with f of degree five or six, the linearly equivalent effective divisors that make up the canonical class are the divisors of the form $[(a, b)] + [(a, -b)]$. The group law presented above on pairs of points matches the group law for divisors. Starting from two pairs of points $\{P_1, P_2\}$ and $\{P_3, P_4\}$, we found the graph of a cubic polynomial $y - c(x)$ that passed through those four points. If the third pair of intersections are $\{P_5, P_6\}$, then $\text{div}(y - c(x)) = \sum_{i=1}^6 [P_i]$, so this divisor is principal. Rearranging and adjusting by points in the canonical divisor class, It follows that $[\widetilde{P_5}] + [\widetilde{P_6}]$ is the effective divisor that is linearly equivalent to $[P_1] + [P_2] + [P_3] + [P_4]$, where $\widetilde{(x, y)} = (x, -y)$. Blowing down \mathfrak{D} identified all of the effective divisors in the canonical class; those are all of the form $[(x, y)] + [\widetilde{(x, y)}]$.

4.3 Étale covers of genus two curves

We have gone on at some length discussing the Jacobian of a curve C of genus two. The following proposition explains why.

Proposition 4.3.1. *Let C be a curve of genus two defined over a field k with $C(k)$*

nonempty. Suppose $D \rightarrow C$ is an étale cover with abelian structure group. There exists an isogeny $\phi : B \rightarrow \text{Jac}(C)$ with the property that D is the pullback of ϕ .

Proof. This is Proposition 9.1 in Milne [40]. See also Coombes and Grant [14] for a slightly different point of view of the same underlying truth. \square

The proposition tells us that if we can find étale covers of the Jacobian, those will in turn determine all the étale covers of the curve. This puts us back in a familiar setting: we are now searching for isogenies that target $\text{Jac}(C)$.

Warning 4.3.2. Under most circumstances the variety B in the theorem is going to be the Jacobian of another curve C' of genus two. However C' and D are unrelated. The Riemann-Hurwitz formula tells us that a curve D with a degree d étale map to a genus two curve must itself have genus $d + 1$.

We again will build from the paper of Katz:

Theorem 4.3.3 (Katz, [31]). *Let A/\mathbb{Q} be an abelian surface. The $n = \ell$ case of Theorem 2.2.6 is true.*

Explicitly, suppose $\#A(\mathbb{F}_p)$ is divisible by ℓ for a set of primes p of density one. Then there exists an abelian surface A'/\mathbb{Q} which is isogenous to A via an ℓ -power isogeny with the additional property that $\ell \mid \# \text{Tors}(A'(\mathbb{Q}))$.

This tells us that we can gain confidence in the fact that there exists an isogeny targeting $\text{Jac}(C)$ by calculating the number of points of $\text{Jac}(C)$ modulo p for many primes. For small primes p there is a convenient formula for the cardinality of the Jacobian over \mathbb{F}_p , which is proved in Appendix A.

Proposition 4.3.4. *Let C/\mathbb{Q} be a curve of genus two and let p be a prime of good reduction for C . Then*

$$\# \operatorname{Jac}(C)(\mathbb{F}_p) = \frac{1}{2} \left(\#C(\mathbb{F}_{p^2}) + \#C(\mathbb{F}_p)^2 \right) - p.$$

We now can determine the size of the Jacobian by counting points on the hyperelliptic curve itself, which is practical for small p even by a naïve exhaustive search: for each $x \in \mathbb{F}_p$, calculate $f(x)$ and decide whether it is a square mod p . If so, there exists a y whereby $(x, \pm y) \in C(\mathbb{F}_p)$. Otherwise that x -coordinate contributes nothing. Cassels and Flynn [9, § 8.2] derive the formula for $\# \operatorname{Jac}(C)(\mathbb{F}_p)$ with their bare hands by counting how many pairs of points of C there are over \mathbb{F}_p and how many conjugate pairs of points C has over \mathbb{F}_{p^2} .

4.4 Jacobians over \mathbb{C}

Similar to the situation for elliptic curves, we can represent every Jacobian of a genus two curve over \mathbb{C} as the quotient of \mathbb{C}^2 by a full-rank lattice Λ ; this time Λ will be the span of four vectors $\{\omega_1, \omega_2, \omega_3, \omega_4\}$. As with elliptic curves, these four vectors are the integrals of invariant differential forms along a basis for the homology group $H_1(C, \mathbb{Z})$.

The notions of homothety and isogeny are the same: two lattices Λ_1 and Λ_2 are homothetic if there is a linear map $M : \mathbb{C}^2 \rightarrow \mathbb{C}^2$ satisfying $M\Lambda_1 = \Lambda_2$. They are isogenous if such an M exists whereby $M\Lambda_1 \subset \Lambda_2$. Homothety descends to an isomorphism of abelian varieties on the quotient, while isogeny descends to a homomorphism with a finite kernel.

There is a new wrinkle in higher dimensions: there are relatively few complex tori which are abelian varieties. The restriction comes from the fact that not every discrete subgroup of \mathbb{C}^n is the set of poles of a meromorphic function.

Definition 4.4.1. Let $\Lambda \subset \mathbb{C}^n$ be a lattice. A **polarization** for Λ is a Hermitian form H on \mathbb{C}^n whose imaginary part E is integer-valued when restricted to Λ .

E recovers H by the formula $H(u, v) = E(iu, v) + iE(u, v)$, so one sometimes sees E referred to as the polarization. Note that E is a skew-symmetric bilinear (i.e., symplectic) real form, so if we identify \mathbb{C}^2 with \mathbb{R}^4 then this \mathbb{R}^4 is a symplectic space. Riemann gave conditions that will tell us whether or not a given symplectic form E is a polarization on a lattice Λ .

Theorem 4.4.2 (Riemann). *Let $\Lambda \subset \mathbb{C}^n$ be a lattice. Write Π for the $n \times 2n$ matrix whose columns are the $2n$ basis vectors of the lattice. A skew-symmetric integer matrix A determines a polarization for Λ if and only if the following two conditions hold:*

1. $\Pi A^{-1} \Pi^{tr} = 0$
2. $i \Pi A^{-1} \overline{\Pi}^{tr}$ is positive-definite.

Proof. See Birkenhake and Lange [1, § 4.2]. □

If Λ is a polarized lattice with symplectic pairing E , a process similar to the process used to row-reduce integer matrices shows that there exists a basis for Λ of the shape $\{a_1, \dots, a_n, b_1, \dots, b_n\}$ with the properties that $E(a_i, a_j) = E(b_i, b_j) = 0$,

while $E(a_i, b_j) = d_i \delta_{ij}$ for some integers d_1, \dots, d_n satisfying $d_1 \mid d_2 \mid \dots \mid d_n$. As such we can assume by changing the basis for Λ that E has the form

$$\left[\begin{array}{c|c} 0 & D \\ \hline -D & 0 \end{array} \right]$$

where D is a diagonal matrix with diagonal entries d_1, \dots, d_n as above. If we can take $D = I_n$, E is called a **principal polarization** and Λ is **principally polarized**. In any case, $\{a_1, \dots, a_n, b_1, \dots, b_n\}$ is called a **symplectic basis** for Λ .

A polarization determines a map between a complex torus and its “dual torus” of degree $\prod d_i$. We will not need the theory of the dual because of the following paragraph.

Jacobians of curves are automatically principally polarized; this comes from the interpretation of the lattice associated to $\text{Jac}(C)$ as arising from the homology group $H_1(C, \mathbb{Z})$, where a topological intersection pairing exists. This means Jacobians are isomorphic to their dual varieties.

Suppose $\Lambda = \langle \omega_1, \omega_2, \omega_3, \omega_4 \rangle$ is a lattice in \mathbb{C}^2 . The linear transformation with matrix $[\omega_3 \ \omega_4]^{-1}$ will produce a homothety which transforms this lattice into one with a basis $\{\tau_1, \tau_2, \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}\}$. Let $\Omega = [\tau_1 \ \tau_2]$, so that post-homothety our lattice is $\Omega \mathbb{Z}^2 \oplus \mathbb{Z}^2$. The Riemann conditions applied to the standard matrix for the principal polarization,

$$J = \left[\begin{array}{c|c} 0 & I_2 \\ \hline -I_2 & 0 \end{array} \right],$$

become the following.

Proposition-Definition 4.4.3. *An 2×2 matrix Ω determines a principally polarized lattice $\Omega\mathbb{Z}^2 \oplus \mathbb{Z}^2$ precisely when*

1. *Ω is symmetric and*
2. *The imaginary part of Ω is positive-definite.*

*The set of such matrices are called the **Siegel upper half-space**. A matrix in the Siegel upper half-space will occasionally be referred to as a **period matrix**.*

Much as in the one-dimensional case the upper half-space is much larger than the space of principally polarized abelian varieties. The issue is again that a lattice has multiple bases. However at this stage we've locked in that the symplectic form on Λ should be given by $E(u, v) = u^{\text{tr}} J v$, and because of this we can only choose amongst symplectic bases for Λ . This limits us to acting by a matrix in $\text{Sp}(4, \mathbb{Z})$, the group of 4×4 matrices with determinant 1 that preserve the pairing in the sense that $M J M^{\text{tr}} = J$. Changing basis and then applying a homothety to return to the Siegel upper half-space produces a “fractional linear transformation” much as in the one-dimensional case. Explicitly, if $M \in \text{Sp}(4, \mathbb{Z})$ is written in block form as

$$M = \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array} \right],$$

then $M \cdot \Omega = (A\Omega + B)(C\Omega + D)^{-1}$. Note the similarity to the action of $\text{SL}(2, \mathbb{Z})$ in the one-dimensional setting.

One possible choice for a fundamental domain for the action of $\text{Sp}(4, \mathbb{Z})$ on the Siegel upper half-space is the set of those period matrices $\Omega = \begin{bmatrix} x_1 + iy_1 & x_3 + iy_3 \\ x_3 + iy_3 & x_2 + iy_2 \end{bmatrix}$ which satisfy the following three conditions:

1. The real part of Ω is “size-reduced,” i.e., $-\frac{1}{2} \leq x_i < \frac{1}{2}$ for all i .
2. The imaginary part of Ω is “ $GL(2, \mathbb{Z})$ -reduced,” i.e., $0 \leq 2y_3 \leq y_1 \leq y_2$.⁶
3. For all $M = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \in \mathrm{Sp}(4, \mathbb{Z})$, we have $|\det(C\Omega + D)| \geq 1$.

Streng [52, § 5.3] gives an algorithm for moving a given point in the upper half-space into this fundamental domain. We’ll use this algorithm in the next chapter as we attempt to recover a curve from its period matrix.

⁶This condition is Gauss’s condition for the binary quadratic form $y_1X^2 + 2y_3XY + y_2Y^2$ to be reduced—it is a somehow natural condition to impose on a positive-definite matrix.

Chapter 5: Calculating a cover in genus two

5.1 Igusa invariants

Playing the role of the j -invariant in genus two are a trio of **Igusa invariants**. These form the three coordinates on the moduli space of curves of genus two—they determine the isomorphism type of a curve over the algebraic closure. Since Torelli’s theorem forms a bridge between curves of genus two and principally polarized abelian surfaces, we can also view the Igusa invariants as markers for orbits of the $\mathrm{Sp}(4, \mathbb{Z})$ action on the Siegel upper half-space.

Definition 5.1.1. Let $y^2 = f(x)$ be the equation for a curve of genus two, with $f(x) = \sum_{i=0}^6 a_i x^i$ a degree six polynomial. Write $\alpha_1, \dots, \alpha_6$ for the roots of f . Given $\sigma \in S_6$, write (ij) for the difference $\alpha_{\sigma(i)} - \alpha_{\sigma(j)}$. The quantities

$$\begin{aligned} I_2 &= a_6^2 \sum_{\sigma}^{15} (12)^2 (34)^2 (56)^2 \\ I_4 &= a_6^4 \sum_{\sigma}^{10} (12)^2 (23)^2 (31)^2 (45)^2 (56)^2 (64)^2 \\ I_6 &= a_6^6 \sum_{\sigma}^{60} (12)^2 (23)^2 (31)^2 (45)^2 (56)^2 (64)^2 (14)^2 (25)^2 (36)^2 \\ I_{10} &= a_6^{10} \prod_{i < j} (\alpha_i - \alpha_j)^2 \end{aligned}$$

are the four **homogeneous Igusa-Clebsch invariants** associated with C . The

notation means to let σ vary over S_6 , including in the sum each unique expression encountered. The upper index on each sum indicates how many terms are included. For I_2 , there are fifteen ways to partition the six roots into three groups of two, and each of those partitions determines a unique expression to the sum. I_4 tallies the ten ways there are to segregate the six roots into two groups of three. I_6 counts the six bijections between the two groups of three in each of those ten segregations. Finally I_{10} is the discriminant of f .

The Igusa-Clebsch constants are defined assuming $f(x)$ is sextic. This is not a restriction: as we saw in the last chapter, every genus two curve has a model of the form $y^2 = \text{degree six}$, while a model of degree five may or may not exist depending on whether $f(x)$ possesses a root in the ground field. It is clear from the symmetry involved that the Igusa-Clebsch invariants are defined over the same field as $f(x)$ itself.

Following Streng [52, § 2.1], set $I'_6 = \frac{1}{2}(I_2I_4 - 3I_6)$. This will be convenient for improving the rate of convergence of some series in what's to come.

Removing the homogeneity produces the invariants we seek.

Definition 5.1.2. The three **absolute Igusa invariants** are¹

$$i_1 = \frac{I_4I'_6}{I_{10}}, \quad i_2 = \frac{I_2I_4^2}{I_{10}}, \quad \text{and} \quad i_3 = \frac{I_4^5}{I_{10}^2}.$$

¹This is to some extent a matter of taste—for instance, van Wamelen [54] uses $\frac{I_2^5}{I_{10}}$, $\frac{I_2^3I_4}{I_{10}}$, and $\frac{I_2^2I_6}{I_{10}}$, while the “most standard” choice, that of Spallek, also involves some factors of 2. The crucial fact is that the invariants are independent elements of the algebra of homogeneous degree zero elements in $\mathbb{Q}[I_2, I_4, I_6, I_{10}^{-1}]$. We'll follow Streng's definition in what follows.

Note that passing from a tuple (I_2, I_4, I'_6, I_{10}) to $(\Delta I_2, \Delta^2 I_4, \Delta^3 I'_6, \Delta^5 I_{10})$ will not change the absolute invariants derived from them, hence why the Igusa-Clebsch invariants are called homogeneous. Igusa proved [26] the absolute invariants of f are isomorphism invariants of the curve $y^2 = f$ over the algebraic closure. If f has a model over k , then its homogeneous Igusa-Clebsch invariants will be elements of k , and therefore the absolute Igusa invariants will also be elements of k . The converse is not true, however: while every triple of elements of k determine a genus two curve, not all of those curves can be defined over k .

An algorithm of Mestre [38] accepts four numbers (the homogeneous invariants) in a field k and outputs a conic and a cubic. If the conic has a rational point, then there exists a curve defined over k with those homogeneous invariants; the intersection of the conic with the cubic determines an equation for the curve. So, provided the curve exists in the first place, Mestre's algorithm promises to construct a curve from its homogeneous invariants. This algorithm is implemented in Sage.

5.2 Theta constants

The Igusa-Clebsch invariants from the previous section were formulated in a way that makes them easy to calculate from the equation for the curve. We ultimately want to find genus two curves from points in the Siegel upper half-space, so we won't have an equation at hand. This will require us to calculate the invariants in another way.

Definition 5.2.1. An element $c = (c_1, c_2, c_3, c_4) \in \{0, \frac{1}{2}\}^4$ is called a **theta char-**

acteristic. Let E be the function $E(z) = \exp(\pi iz)$. The **theta constant with characteristic** c is the complex-valued function on the Siegel upper half-space $\theta[c]$ defined by

$$\theta[c](\Omega) = \sum_{n=(n_1, n_2) \in \mathbb{Z}^2} E\left(\begin{bmatrix} n_1 + c_1 & n_2 + c_2 \end{bmatrix} \Omega \begin{bmatrix} n_1 + c_1 \\ n_2 + c_2 \end{bmatrix} + 2 \begin{bmatrix} n_1 + c_1 & n_2 + c_2 \end{bmatrix} \begin{bmatrix} c_3 \\ c_4 \end{bmatrix}\right).$$

These are sometimes called “theta-null values” because a second parameter z has been set to zero in the above series. Not all of these constants define interesting values: six of them are zero. Using Dupont’s notation in [17], write $\theta[c] = \theta_{16c_2+8c_1+4c_4+2c_3}$. A theta characteristic is even or odd depending on the parity of the integer $4(c_1c_3 + c_2c_4)$. Six theta characteristics are odd and ten are even. The six theta constants corresponding to the six odd theta characteristics are zero. This is due to anti-symmetry in the series.² The ten nonzero (even) theta constants are $\theta_0, \theta_1, \theta_2, \theta_3, \theta_4, \theta_6, \theta_8, \theta_9, \theta_{12}$, and θ_{15} .

The theta constants are the building blocks of the homogeneous invariants. To state the formulas, first let T denote the set of the ten even theta constants. In order to use Streng’s I'_6 , we need to go one further step. Define

$$S = \left\{ C \subset T : \#C = 4, \sum_{c \in C} c \in \mathbb{Z}^4 \right\}.$$

The set S has cardinality fifteen and consists of so-called “Göpel quadruples”. Call a set $\{b, c, d\} \subset T$ *syzygous* if it is a subset of a Göpel quadruple; there are sixty of

²As an example, $\theta_5 = \theta[\frac{1}{2}, 0, \frac{1}{2}, 0]$ is zero because the term indexed by (n_1, n_2) is the negative of the term indexed by $(-n_1, n_2)$. If you feel compelled to check this, keep in mind that $E(z_1) = -E(z_2)$ when z_1 and z_2 differ by an odd integer.

these. Now set

$$h_4 = \sum_{c \in T} \theta[c]^8$$

$$h_6 = \sum_{\substack{b, c, d \in T \\ \text{syzygous}}} \pm \theta[b]^4 \theta[c]^4 \theta[d]^4$$

$$h_{10} = \prod_{c \in T} \theta[c]^2$$

$$h_{12} = \sum_{C \in S} \prod_{c \in T \setminus S} \theta[c]^4.$$

The signs in h_6 are determined by the fact that h_6 is a modular form for $\mathrm{Sp}(4, \mathbb{Z})$, normalized so that the sign on the term $\theta[0, 0, 0, 0]^4 \theta[0, 0, 0, \frac{1}{2}]^4 \theta[0, 0, \frac{1}{2}, 0]^4$ is $+$. Streng gives a four-line formula for the sign of each term on page 68 of [52].³

Proposition 5.2.2 (Igusa [27], pg. 848). *Let Ω be a point in the Siegel upper half-space. If $h_{10}(\Omega) \neq 0$, the principally polarized abelian variety corresponding to Ω is the Jacobian of a curve C defined over \mathbb{C} whose homogeneous Igusa-Clebsch invariants satisfy*

$$I_2(C) = \frac{h_{12}(\Omega)}{h_{10}(\Omega)}, \quad I_4(C) = h_4(\Omega), \quad I'_6(C) = h_6(\Omega), \quad I_{10}(C) = h_{10}(\Omega).$$

The theta constants are infinite series, but Streng estimated their tails in order to decide how many terms are needed to reach a desired precision.

Proposition 5.2.3 (Streng [52], Algorithm II.7.15). *Choose $s \in \mathbb{N}$. To estimate $\theta[c](\Omega)$ with error at worst 2^{-s} , first compute $R = \lceil \sqrt{0.51s + 2.55} \rceil$. Then calculate*

³This hints that perhaps this “direct” way of defining the h_i was an obfuscation of what’s going on. One can reach the h_i via formulas involving Eisenstein series for $\mathrm{Sp}(4, \mathbb{Z})$. Since we ultimately want to calculations, going deeper into the theory would lead us afield.

the sum determining the theta constant, only including the terms for those $(n_1, n_2) \in \mathbb{Z}^2$ satisfying $|n_i| \leq R$; for summand meeting this criterion, calculate it to within a precision of $s + 1 + \lfloor 2 \log_2(2R + 1) \rfloor$.

Armed with these propositions we now have some hope to come back to a hyperelliptic curve from a point in the Siegel upper half-space. Of course, much as in the one-dimensional case we could get a curve that is off by a twist. This will be addressed in the next section.

5.3 Twisting theory of curves of genus two

In general, twists of an object C defined over a field K correspond to the first Galois cohomology set $H^1(\text{Gal}(\overline{K}/K), \text{Aut}(C))$. The automorphism group of a curve of genus two is always finite, and a list of what groups are possible is well-known.⁴ These were worked out in characteristic zero by Bolza [3] and extended to arbitrary characteristic by Igusa [26].

Not surprisingly, then, the possibilities for what twists are possible depends on what automorphisms exist. In his thesis Cardona [8] gives models for twists depending on the automorphism group of C . The most generic situation is when

⁴Be aware that the Jacobian of C can have lots of automorphisms: $\text{End}(\text{Jac}(C))$ could be the ring of integers in a quartic CM-field. Dirichlet's theorem says that such a ring will have infinitely many units, giving infinitely many automorphisms of $\text{Jac}(C)$. We are asking for invertible morphisms $C \rightarrow C$ here, not $\text{Jac}(C) \rightarrow \text{Jac}(C)$. In fact the automorphism group of a curve of genus $g \geq 2$ over a field of characteristic zero is always finite and has cardinality at most $84(g - 1)$. This is a theorem of Hurwitz; see Hartshorne [23, Ex. IV.2.2].

$\text{Aut}(C) = \mathbb{Z}/2$, with the only automorphism being the hyperelliptic involution $(x, y) \mapsto (x, -y)$. One might expect, mirroring the situation in genus one, that when there are no extra automorphisms the only options will be quadratic twists, and indeed this is the case.

Proposition 5.3.1. *Let C/K be a curve of genus two defined over a field K , with $\text{Aut}(C) = \mathbb{Z}/2$. Suppose C has the model $y^2 = f(x)$. Then the set of twists of C is in bijection with $K^\times / (K^\times)^2$. Given $d \in K^\times / (K^\times)^2$, the twist of C by d has equation $dy^2 = f(x)$.*

Proof. The hyperelliptic involution is defined over K , so $\text{Aut}(C)$ is a trivial Galois module. Therefore $H^1(K, \text{Aut}(C)) = H^1(K, \mathbb{Z}/2) = K^\times / (K^\times)^2$. Classes in this H^1 correspond to quadratic extensions of K . Explicitly, the cocycle corresponding to d is $\xi : \sigma \mapsto \frac{\sigma(\sqrt{d})}{\sqrt{d}}$. The matrix $M = \begin{bmatrix} 1/\sqrt{d} & 0 \\ 0 & 1/\sqrt{d} \end{bmatrix}$ satisfies $\sigma(M) = \xi(\sigma)M$, so the fractional linear transformations determined by this matrix (as in §4.1) will determine a curve with the appropriately-twisted action.⁵ The transformation works out to $x \mapsto x$, $y \mapsto \sqrt{d}y$. So $y^2 = f(x)$ becomes $dy^2 = f(x)$. \square

By multiplying through by d and changing y to dy , we can also represent the quadratic twist of $y^2 = f(x)$ by d by $y^2 = df(x)$. If $\deg f = 5$, we have the other option of changing variables by $x \mapsto dx$, $y \mapsto d^{-3}y$ to turn the curve with equation $dy^2 = \sum_{i=0}^n a_i x^i$ with the equation $y^2 = \sum_{i=0}^n d^{n-i} a_i x^i$.

There's one other possible problem that needs to be addressed. We've discussed twisting curves C , but we really want to twist the Jacobian of C , since the

⁵The correspondence between cocycles and equations for twists is summarized well in [51, X.2].

cover that we're creating is a cover of the Jacobian. Thankfully these are the same.

Proposition 5.3.2. *Let $C^{(d)}$ be the quadratic twist of the genus two curve C by the nonsquare d . Then $\text{Jac}(C^{(d)}) = \text{Jac}(C)^{(d)}$.*

Proof. See Petersen [45, p. 5]. □

It's easy to compute the effect that a quadratic twist has on our invariants, since the roots of the polynomial $f(x)$ don't change. The only change relevant for the invariants is the leading coefficient diminishing by a factor of $\frac{1}{d}$. This means I_2 will be changed by a factor of d^{-2} , I_4 will be changed by d^{-4} , I_6 by d^{-6} , and I_{10} by d^{-10} . We can use this fact to twist curves without having a model for them, i.e., before applying Mestre's algorithm.

A quadratic twist is again “half an isomorphism,” since twisting twice produces the same thing as the change of variables determined by the matrix $\begin{bmatrix} d & 0 \\ 0 & d \end{bmatrix}$. So if C/\mathbb{Q} , when it comes time to twist C we need to answer the following question for each prime number p : should p be included as a factor of d or not? The alternate model for the quadratic twist makes it clear that a weak version of Proposition 3.2.3 holds in genus two: if C has good reduction at p , its twist by p will not. This means that if we know ahead of time that C should possess a twist with good reduction at p , while C itself has bad reduction at p , we should try twisting by p to “clear” the bad reduction.

5.4 (p, p) -isogenies

The final ingredient needed to create an algorithm is a determination of those points in the Siegel upper half-space that represent abelian varieties that are isogenous to the one we're starting with. Unlike in the one-dimensional situation, there is a restriction on the type of finite subgroup that one can quotient by.

Proposition 1 (Poonen, [46]). *Suppose A is a principally-polarized abelian surface with no extra endomorphisms and $H \subset A$ is a finite subgroup whose order is not a square. Then the abelian variety $B = A/H$ cannot be principally polarized.*

Proof. The composition of an isogeny ϕ with its dual should be the multiplication-by- $\deg \phi$ isogeny. Using a putative principal polarization for B in the sequence

$$A \xrightarrow{\phi} B \cong B^\vee \xrightarrow{\widehat{\phi}} A^\vee \cong A$$

produces such an isogeny on A . But because A has no extra endomorphisms, the only isogenies $A \rightarrow A$ are multiplication-by- n isogenies which have degree n^4 . Meanwhile the centered isogeny $A \rightarrow A$ has degree $(\deg \phi)^2$, but by assumption this quantity is not a fourth power—contradiction. \square

This proposition tells us that, in general, the kernels of our isogenies need to have square order. Abelian surfaces are either split, that is, isogenous to the product of two elliptic curves, or they are simple. If an abelian surface is isogenous to a split surface it is itself split. Jacobians of curves all carry a principal polarization, so if we start with a simple Jacobian with no extra endomorphisms it cannot carry an

isogeny of non-square degree to another Jacobian by the above proposition. If we look for isogenies of smallest degree possible, this restricts us to looking for kernels of isogenies that are abstractly isomorphic to $\mathbb{Z}/\ell \oplus \mathbb{Z}/\ell$.⁶ In what follows we'll focus on the latter.

Definition 5.4.1. *Two abelian surfaces A and B are called (ℓ, ℓ) -isogenous if there is an isogeny $A \rightarrow B$ between them whose kernel is a maximal isotropic $\mathbb{Z}/\ell \oplus \mathbb{Z}/\ell$ inside $A[\ell]$. (Isotropic is with respect to the Weil pairing.)*

The Weil pairing in higher dimensions is a pairing $A[n] \times \widehat{A}[n] \rightarrow \mu_n$. As the abelian varieties we care about are Jacobians which are principally polarized, there is an isomorphism $\widehat{\text{Jac}(C)} \xrightarrow{\sim} \text{Jac}(C)$, and so the Weil pairing can be defined on $\text{Jac}(C)[n] \times \text{Jac}(C)[n]$. This is explained in Milne [41, §13], as well as why kernels of (ℓ, ℓ) -isogenies are necessarily isotropic (Proposition 13.8). The polarization induces a symplectic pairing on $(\mathbb{Z}/n)^4$ which can be identified with the Weil pairing. A subspace of $\text{Jac}(C)[n]$ is called isotropic if the pairing restricted to that subspace is the zero pairing. The largest an isotropic subspace can be is half of the dimension of the ambient vector space. For us, that means that we're looking for two-dimensional subspaces as our kernels. These will be abstractly isomorphic to $\mathbb{Z}/\ell \oplus \mathbb{Z}/\ell$.

We now need representatives for homothety classes of lattices which are (ℓ, ℓ) -isogenous to a given lattice Λ . Much as in the one-dimensional situation, with a

⁶Why not \mathbb{Z}/ℓ^2 ? It is true that its order is a square, but a map with cyclic kernel followed by its dual would produce a map $\text{Jac}(C) \rightarrow \text{Jac}(C)$ of order ℓ^4 whose kernel is abstractly $\mathbb{Z}/\ell^2 \times \mathbb{Z}/\ell^2$. Again, we know every endomorphism of degree ℓ^4 : there is $[\ell]$ and $[-\ell]$ and nothing else. The kernels of those two endomorphisms have exponent ℓ , not ℓ^2 .

homothety we can rearrange it so that the isogeny is given by the 2×2 matrix $\begin{bmatrix} \ell & 0 \\ 0 & \ell \end{bmatrix}$. With this convention, each isogenous lattice will correspond to a superlattice $\Lambda' \supset \Lambda$ for which $\Lambda/p\Lambda'$ is a maximal isotropic subspace of $\Lambda/\ell\Lambda = (\mathbb{Z}/\ell)^4$. The next proposition enumerates those subspaces.

Proposition 5.4.2. *There are $\ell^3 + \ell^2 + \ell + 1$ maximal isotropic subspaces of the symplectic space \mathbb{F}_ℓ^4 .*

Proof. The astute reader may have noticed that the stated formula is the cardinality of \mathbb{P}^3 . We'll exhibit an (indirect) bijection between the set of isotropic planes and the set of lines. Given a line L , its orthogonal complement is three-dimensional. Since every vector pairs with itself to zero, there are two dimensions' worth of vectors in L^\perp that are linearly independent from L . Dividing by scalars leaves us with $\ell + 1$ isotropic planes that involve L .

On the other hand, each isotropic plane is a plane, so it possesses $\ell + 1$ lines. These last two facts taken together imply that there are an equal number of lines and isotropic planes. \square

We now need to exhibit these lattices. One way to approach this, as is done in [5], is to define the group

$$\Gamma_0^{(2)}(\ell) = \left\{ \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array} \right] \in \mathrm{Sp}(4, \mathbb{Z}) : C \equiv 0 \pmod{\ell} \right\}.$$

This is the group of transformations which, taken mod ℓ , stabilize the isotropic plane generated by the reduction of the first two basis vectors. Witt's extension theorem [11, Theorem 7.5] implies that $\mathrm{Sp}(4, \mathbb{F}_\ell)$ will permute its maximal isotropic

subspaces transitively. As such, the group $\Gamma_0^{(2)}(\ell)$ has index $\ell^3 + \ell^2 + \ell + 1$ in $\mathrm{Sp}(4, \mathbb{Z})$, and the planes we're searching for are going to be spanned modulo ℓ by the first two columns of the various coset representatives for $\mathrm{Sp}(4, \mathbb{Z}) / \Gamma_0^{(2)}(\ell)$. There are coset representatives in [5], but only the first ℓ^3 of them determine distinct cosets. Here is a correct list.

Proposition 5.4.3. *The $\ell^3 + \ell^2 + \ell + 1$ cosets of $\Gamma_0^{(2)}(\ell)$ inside $\mathrm{Sp}(4, \mathbb{Z})$ can be organized into four types:*

$$\text{Type (I, } a, b, c) \quad \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ a & b & 1 & 0 \\ b & c & 0 & 1 \end{bmatrix} \quad a, b, c \in \mathbb{F}_p \quad (\text{cardinality } \ell^3)$$

$$\text{Type (II, } a, b) \quad \begin{bmatrix} 1 & 0 & 0 & 0 \\ a & 0 & 0 & -1 \\ b & -a & 1 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \quad a, b \in \mathbb{F}_p \quad (\text{cardinality } \ell^2)$$

$$\begin{array}{ccc}
\text{Type (III, } a) & \begin{bmatrix} 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ a & 0 & 1 & 0 \end{bmatrix} & a \in \mathbb{F}_p \quad (\text{cardinality } \ell)
\end{array}$$

$$\begin{array}{ccc}
\text{Type (IV)} & \begin{bmatrix} 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} & (\text{cardinality } 1)
\end{array}$$

Proof. That these matrices are in $\text{Sp}(4, \mathbb{Z})$ can be checked by computing symplectic inner products of the columns. The first column must pair with the third column to 1 and the second column must pair with the fourth column to 1, while all other pairings must come out zero. It's easy to check that this is the case for all of our purported coset representatives.

That no two matrices on our list are in the same $\Gamma_0^{(2)}(\ell)$ orbit involves deciding whether the first two columns ever determine the same plane modulo ℓ . Based on the position of the pivots, we only have to investigate this question within each type. Note that since our planes are spanned by the columns, we are only allowed to use column operations to change bases for the subspaces. With that said, no column operation can transform a plane of a given type into another plane of a that type because the pivot rows are always $[1 \ 0]$ or $[0 \ 1]$.

One could also verify this classification mechanically by checking that $h_1^{-1}h_2 \notin \Gamma_0^{(2)}(\ell)$ for any h_1, h_2 on the list. (This is the observation that led to the discovery

that something was wrong with the classification in [5].) □

5.5 Tentative algorithm

Algorithm 5.5.1. *The following algorithm takes as input a curve C/\mathbb{Q} of genus two for which $\text{Aut}(C) \cong \mathbb{Z}/2$. The program either runs forever or outputs each genus two curve C'/\mathbb{Q} for which there exists a (p, p) -isogeny $\text{Jac}(C') \rightarrow \text{Jac}(C)$.*

1. *Become convinced that $\text{Jac}(C)$ could possess such a cover. (This involves checking whether $p \mid \# \text{Jac}(C)(\mathbb{F}_\ell)$ for many primes ℓ .)*
2. *Calculate the big period matrix Π_C for $\text{Jac}(C)$. In what follows we'll assume that Π_C is a 2×4 matrix with complex entries. Let Λ be the lattice spanned by its columns.*
3. *Fix a precision s .*
4. *Using the coset representatives from 5.4.3, do the following:*
 - (a) *Construct the superlattice of Λ corresponding to that coset representative;*
 - (b) *Using Streng's formulas, calculate the theta characteristics of the superlattice;*
 - (c) *From the theta characteristics, calculate the Igusa invariants i_1 , i_2 , and i_3 . Check whether they are real to within our tolerance s .*
 - (d) *If they are real, test whether they are apparently rational, e.g., by computing their continued fraction and looking for a convergent of larger than 100 digits. If they are, remember the triple (i_1, i_2, i_3) .*

(e) *If nothing is remembered, repeat step (d) with precision doubled. Continue this until something is found.*

5. *For those Igusa invariants remembered in step 4:*

(a) *Use Mestre’s algorithm to construct a curve C' of genus two having those Igusa invariants.*

(b) *Calculate the conductors of C and C' . Quadratically twist C' by each prime that divides only one of these values.*

(c) *Find a reduced model of C' , so that C and C' have the same primes of good reduction.*

(d) *Test whether C and C' now have the same zeta function at the first many primes of good reduction. If so, remember C' and stop. Otherwise go on to step (e).*

(e) *Systematically further twist C' by every product of primes of bad reduction. After each twist, check whether C' and C have zeta functions that agree at every prime. When they do, remember C' and stop.*

6. *Output every curve C' discovered in step 5.*

While it is true that the invariants of the isogenous curves satisfy some sort of “modular equations,” those equations have rational coefficients and have never been calculated in a form that would allow us to state an analogue of the estimate we used in §3.3 in the genus two setting. See Milio [39] for a discussion on whether a different set of invariants might help to overcome this issue.

Note that because a modular polynomial exists, if a curve C' exists its Igusa invariants will eventually be found in step 4. The logic is much as in the dimension 1 setting: a bound for the denominators of the isogenous invariants is determined by the degree of the polynomial and the denominators of the coefficients thereof. The lack of an explicit bound for the denominators obviously hampers our ability to give a runtime analysis, or to say ahead of time what precision is needed. On the other hand, if no such cover exists, the program will run forever. Note that how long one needs to test each apparently-real Jacobian could vary from lattice to lattice; this isn't explicitly built into the algorithm but deserves mention here.

For a discussion of how to complete Step 5(e), see §6.1.

5.6 Example

In this section we'll examine the hyperelliptic curve C/\mathbb{Q} defined by

$$y^2 = -99x^6 - 1782x^5 - 23463x^4 - 30888x^3 - 14652x^2 - 12672x - 8976.$$

The discriminant of the sextic is $-2^{21} \cdot 3^{19} \cdot 5^3 \cdot 11^{22} \cdot 17$, so C has good reduction away from those five primes. In fact, Magma reports the conductor of C is $2^3 \cdot 3^8 \cdot 5 \cdot 17$, so the bad reduction at 11 is an illusion. The homogeneous Igusa-Clebsch invariants of C are

$$\begin{aligned} (I_2, I_4, I_6, I_{10}) = & (2^{10} \cdot 3^4 \cdot 11^2 \cdot 31 \cdot 47, 2^{14} \cdot 3^9 \cdot 5 \cdot 11^8 \cdot 1667, \\ & -2^{21} \cdot 3^{12} \cdot 7 \cdot 11^{10} \cdot 258173, -2^{41} \cdot 3^{19} \cdot 5^3 \cdot 11^{22} \cdot 17). \end{aligned}$$

This curve is the curve $\widetilde{C}_{2,1,3}$ from the paper of Bruin, Flynn, and Testa [7]. In it, they prove that the Jacobian of C is targeted by a $(3, 3)$ -isogeny from the Jacobian of the curve

$$C_{2,1,3} : y^2 = 13x^6 + 42x^5 + 129x^4 + 200x^3 + 276x^2 + 192x + 112,$$

whose invariants are

$$(I_2, I_4, I_6, I_{10}) = (-1 \cdot 2^{10} \cdot 3^2 \cdot 619, 2^{14} \cdot 3^5 \cdot 5 \cdot 491, \\ -1 \cdot 2^{21} \cdot 3^6 \cdot 7 \cdot 11 \cdot 149011, -1 \cdot 2^{43} \cdot 3^9 \cdot 5 \cdot 17^3).$$

We will try to reconstruct this cover using the algorithm from §3, ported to the genus two setting. We are only trying to verify an isogeny that is independently known to exist because there are some hangups in proving the output of the algorithm is actually isogenous to the starting curve; see §6.1 for a deeper discussion.

For the moment we'll feign ignorance about $C_{2,1,3}$ and begin the algorithm. For every prime p satisfying $20 < p < 500$, the Jacobian of C possesses a point of order three. Nevertheless, Magma tells us that the torsion subgroup of $\text{Jac}(C)$ over \mathbb{Q} is trivial. This motivates us to search for a curve C' whose Jacobian possesses a $(3, 3)$ -isogeny targeting $\text{Jac}(C)$.

Calculations in this section are done to 500 digits of precision (approximately 1692 bits), though of course printed approximations will be significantly shorter.

Magma computes the big period matrix as follows, by estimating integrals of holomorphic differentials on C :

$$\Pi = \begin{bmatrix} 0.0088 + 0.0964i & 0.1233 + 0.0542i & 0.0316 + 0.0000i & -0.0379 + 0.0000i \\ 0.2066 + 0.0120i & -0.1770 + 0.0325i & -0.0760 + 0.0000i & 0.1204 + 0.0000i \end{bmatrix}.$$

It's an easy check that for the standard symplectic form

$$J = \left[\begin{array}{cc|cc} & & 1 & 0 \\ & 0 & 0 & 1 \\ \hline -1 & 0 & & \\ 0 & -1 & & 0 \end{array} \right],$$

Riemann's two conditions $\Pi J^{-1} \Pi^{\text{tr}} = 0$ and $i \Pi J^{-1} \overline{\Pi}^{\text{tr}} > 0$ hold to within our tolerance. So for the lattice Λ spanned by the four columns $\{w_1, w_2, w_3, w_4\}$ of Π , the basis of w_i is already symplectic. (As noted, this was to be expected based on how Magma calculated the matrix in the first place.)

It's now time to bring in the 40 lattices that are $(3, 3)$ -isogenous to Λ . As mentioned in the previous section, these are indexed by the coset representatives for $\Gamma_0^{(2)}(3)$. The superlattice of interest can be derived from a coset representative by using the columns of the coset representative as coefficients in a linear combination of w_1, w_2, w_3 , and w_4 , then dividing the first two of those four vectors by 3. This sets it up so that the isogeny is realized by multiplication by 3. As an explicit example, consider the coset representative of type $(I, 2, 1, 2)$. The matrix is reproduced below along with the basis for the superlattice it corresponds to.

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 2 & 1 & 1 & 0 \\ 1 & 2 & 0 & 1 \end{bmatrix}$$

$$\Lambda' = \left\langle \frac{w_1 + 2w_3 + w_4}{3}, \frac{w_2 + w_3 + 2w_4}{3}, w_3, w_4 \right\rangle.$$

Continuing with this example, by multiplying the first two basis vectors by the inverse of the 2×2 matrix consisting of the latter two basis vectors, we create a homothetic lattice of the form $\Omega\mathbb{Z}^2 \oplus \mathbb{Z}^2$. Ω is the following period matrix Ω in the Siegel upper half-space:

$$\Omega = \begin{bmatrix} 2.667 + 2.715i & 2.167 + 1.746i \\ 2.167 + 1.746i & 1.333 + 1.191i \end{bmatrix}.$$

This matrix is not in the fundamental domain; using Streng's algorithm II.5.3 [52] we discover that the fractional linear transformation associated with $B \in \text{Sp}(4, \mathbb{Z})$ moves Ω to a period matrix τ in the fundamental domain.

$$B = \begin{bmatrix} 2 & -3 & 2 & 0 \\ 0 & 1 & -3 & -2 \\ -1 & 1 & 1 & 1 \\ -1 & 2 & -1 & 0 \end{bmatrix}, \quad \tau = \begin{bmatrix} 0.3063 + 1.230i & -0.4187 + 0.2994i \\ -0.4187 + 0.2994i & 0.3010 + 1.435i \end{bmatrix}.$$

Next we calculate the ten theta constants associated to τ , using Streng's estimate for the number of terms needed and the precision required for each term. From this we achieve estimates for the quantities h_4 , h_6 , h_{10} , and h_{12} , and then we can calculate approximations to the absolute invariants i_1 , i_2 , and i_3 .

$$h_4 \approx 3.8270394 + 0.5064270i$$

$$h_6 \approx 4.4333813 - 1.0686329i$$

$$h_{10} \approx 0.0017842635 + 0.0005356077i$$

$$h_{12} \approx -0.0020507352 - 0.0096667904i$$

$$i_1 \approx 8716.620697448 - 3650.360656750i$$

$$i_2 \approx -21421.44921783 - 36630.55012339i$$

$$i_3 \approx 246357160.0004 + 18404949.8119i .$$

As the h_i were calculated to within 10^{-500} , we can confidently say that the Jacobian corresponding to this point in the Siegel upper half-space is not defined over the real numbers.

This process is then repeated for the other 40 coset representatives. Six of them give points that are apparently the Jacobians of curves defined over \mathbb{R} . For those, the Igusa-Clebsch invariants were tested to see if they were close to a rational number of comparatively small height. This was accomplished by means of continued fractions. Despite a lack of a rigorous bound for the denominators of the Igusa invariants, we can achieve practical results as long as we have calculated the h_i to precision far exceeding the size of the denominators of the rational numbers we are searching for. In such a situation the continued fraction expansion of our estimates will lay bare the fact that our estimate is a small rational number together with a very small perturbation coming from rounding error. This small perturbation will manifest itself as an astoundingly large convergent to the continued fraction. Truncating the continued fraction at the enormous convergent will give the rational number we intended.

In our example, six of the 40 lattices had Igusa invariants whose imaginary parts were on the order of magnitude of 10^{-500} , our working precision. Of those six real Jacobians, one of them was defined over the rational numbers. For type

$(I, 0, 2, 1)$, the Igusa invariant i_1 had continued fraction expansion beginning

$$[3601; 4, 1, 4, 6, 3, 14, 1, 4, 2, 1, 6, 1, 2, 477\text{-digit convergent}]$$

strongly suggesting that $i_1 = [3601; 4, 1, 4, 6, 3, 14, 1, 4, 2, 1, 6, 1, 2] = \frac{9058680441}{2515456}$. Similarly gargantuan continued fraction convergents for i_2 and i_3 give us hope that the Jacobian associated with this period matrix is the Jacobian of a hyperelliptic curve defined over \mathbb{Q} . The Igusa invariants of this curve are presented below. By exploiting homogeneity we also can write down the homogeneous invariants. To do this we use that $I_2 = \frac{h_{12}}{h_{10}}$. Then we scale the quadruple (I_2, I_4, I_6, I_{10}) in a preliminary fashion so that $I_2 = 1$. The rest of the invariants should be rational, and we check this with continued fractions. Then after we have rational invariants we can clear denominators and write down an integer tuple of homogeneous invariants. This gives

$$(I_2, I_4, I_6, I_{10}) = (14826, 66285, 309793869, 509379840),$$

$$(i_1, i_2, i_3) = \left(\frac{9058680441}{2515456}, \frac{20145933765}{157216}, \frac{7801287023203417125}{1581879721984} \right).$$

Breaking the fourth wall for a moment, note that this set of homogeneous Igusa invariants matches that for $C_{2,1,3}$ up to homogeneity.

Magma's implementation of Mestre's algorithm took in the homogeneous Igusa-Clebsch invariants and returned a degree six polynomial $f_{\text{start}}(x)$ which is too large to print here. We will have to hack this down to size. There are two principles that we can use to get underway. First, we can multiply or divide f freely by squares, since those can be absorbed into y^2 on the other side of the equation. Second, if

a prime divides the leading coefficient, we can clear that prime with a quadratic twist. Suppose the curve had a model $y^2 = px^6 + \dots$. Then twisting by p creates a curve with the model $py^2 = px^6 + \dots$ and dividing by p clears the leading coefficient. Factoring the coefficients of the polynomial from Mestre's algorithm revealed primes dividing every coefficient, which can also be cleared with the same twist. There were also primes which sat in the coefficients in the pattern

$$a_6x^6 + p^2a_5x^5 + p^4a_4x^4 + \dots + p^{12}a_0,$$

for which we can perform an isomorphism $y \mapsto p^6y$, $x \mapsto p^2x$ in order to clear. Lastly there was one prime (619) which sat in the coefficients in the pattern

$$p^6a_6x^6 + p^5a_5x^5 + \dots + a_0,$$

which can be absorbed into x . Having made these cosmetic changes, we're down to the model

$$\begin{aligned} y^2 = & 694626095965745028230294955391585x^6 \\ & - 10481319775165116968116978940843831518950x^5 \\ & + 65897603877401152296211405922072845784012150625x^4 \\ & - 220964253553870508421121288893336483711391122393600000x^3 \\ & + 416770493527637700115454733118980958248490124593143822750000x^2 \\ & - 419247665361735107956943612244433762183404391074557037760182000000x \\ & + 175724824257935421232943831835203148453580610817862878046361615940000000. \end{aligned}$$

The discriminant of the sextic is $-2^{33} \cdot 3^{69} \cdot 5^{121} \cdot 17^{13} \cdot 23^{30} \cdot 79^{30} \cdot 972777^{30}$, indicating that we still have work to do at these latter three primes. These last

three primes do not divide the conductor of this curve, suggesting that we should not twist it any further. This polynomial reduces to the sixth power of a monomial modulo 23, 79, and 972777. For instance, mod 23, we get $f \equiv (x + 18)^6$. This suggests that we can improve the model at 23 without a quadratic twist. Following van Wamelen [54], in this situation the bad reduction can be cleared by replacing $f(x)$ with $23^{-6}f(23x - 18)$. The resulting polynomial is separable at 23. (Note that this is an admissible change of variables.)

Repeating this process for $p = 79$ and $p = 972777$, we now have a model with discriminant $-2^{33} \cdot 3^{69} \cdot 5^{121} \cdot 17^{13}$. Magma tells us the conductor of this curve is $2^3 \cdot 3^8 \cdot 5 \cdot 17$, matching the curve we started with, so there's some hope that we are in fact done at this point. Indeed, the Jacobian of the present curve,

$$\begin{aligned} y^2 = & 694626095965745028230294955391585x^6 - 1876945031945052561176765681403210x^5 \\ & + 2113200824518592887128993300221025x^4 - 1268903854932889653881486911346200x^3 \\ & + 428587192513964428603625675261400x^2 - 77205526485657988588419249701760x \\ & + 5794905354192578677033254702160, \end{aligned}$$

has a zeta function which matches that of C for every prime p which satisfies $20 < p < 1000$. This gives us high confidence that the two curves have isogenous Jacobians. Magma tells us that the Jacobian of this curve has torsion $\mathbb{Z}/3 \oplus \mathbb{Z}/3$ over \mathbb{Q} , as we expected. To actually prove that they are isogenous, let's check to make sure $C_{2,1,3}$ is in the isomorphism class of the previous centered equation. Asking Magma if the two curves are $\text{GL}_2(\mathbb{Q})$ -equivalent returns a satisfying answer of yes; the change of variables needed is as follows. If $g(x)$ is the sextic above and

$f_{2,1,3} = 13x^6 + 42x^5 + \cdots + 112$ is the sextic defining the target curve, then with

$$(a, b, c, d) = \left(1, \frac{11043}{101776}, \frac{451987}{203552}, \frac{49041}{203552}\right),$$

we get the equivalence

$$f_{2,1,3} = 13 \cdot \frac{4445619213153221742415895855104}{34393327766239643096923828125} \cdot g\left(\frac{ax+b}{cx+d}\right) \cdot (cx+d)^6.$$

Chapter 6: Future directions

6.1 Proof of correctness for the method of §5.5

The algorithm we presented in §5.5 was admittedly a proof-of-concept. There are two hurdles that need to be cleared in order to prove correctness:

1. **A bound for the denominators of isogenous Igusa invariants.** This is the least troubling of the issues. In §3.4 we used the fact that j -invariants of ℓ -isogenous elliptic curves are related by the modular polynomial $\Phi_\ell(x, y)$, which has integer coefficients and a certain shape. From this we concluded that we could detect if a complex number is the j -invariant of a curve ℓ -isogenous to E , provided we could calculate it to precision greater than $n^{-\ell-1}$, where n is the denominator of $j(E)$.

The Igusa invariants of Jacobians that are (ℓ, ℓ) -isogenous to a given Jacobian are known to satisfy a “modular polynomial” [5], though these polynomials are very cumbersome. Is it possible to analyze them, or variants of modular polynomials, to decide upon a precision needed in order to determine whether given complex numbers are Igusa invariants of isogenous Jacobians?

As mentioned, this is more a practical issue than a theoretical one: the algo-

rithm can be proved correct without knowing the specific tolerance needed. But it would be nice to remove the language about the program running forever—it would certainly be preferable to have the algorithm output a negative answer if no cover actually exists.

2. **Proving apparently isogenous Jacobians are isogenous.** If one runs the algorithm on a general genus-two curve C , a curve C' is produced which has a Jacobian which is apparently isogenous to $\text{Jac}(C)$. They have the same conductor and their reductions at primes of good reduction seem always to have the same zeta function.

The Serre-Faltings method says that two Galois representations are isomorphic once their reductions agree at some finite set of primes. This has been made more-or-less explicit by Livné in the case that ρ is a two-dimensional representation (i.e., the Tate module of an elliptic curve) in [35]. The finite needed is a set of primes whose relative Frobenius elements generate a large subset of the Galois group of the maximal quadratic extension of \mathbb{Q} unramified outside the primes of bad reduction for E . As far as I could see, there is no comparable description of the set needed to show two four-dimensional representations are isomorphic. Such a description is needed in order to know how many local zeta functions to calculate before the algorithm can terminate with a certificate that two Jacobians are isogenous. It would be interesting to try to mimic Livné in dimension 4.

6.2 Equations for covering curves

As mentioned in Warning 4.3.2, the Riemann-Hurwitz formula tells us that the pullback of an isogeny of degree n targeting the Jacobian of a degree two curve should have genus $n+1$. At no point in this thesis do we attempt to give equations for these covers. It would be interesting to come up with a procedure that explicitly calculates equations for these higher-genus curves. This could be useful for determining the exact number of points on a curve of genus two using the Chabauty method, as is done in Coombes and Grant [14].

6.3 What (ℓ, ℓ) -isogenies are possible over \mathbb{Q} ?

We saw that Mazur’s theorem implies that the only primes that could divide $\#\text{Ker}(E/\mathbb{Q})$ are 2, 3, 5, and 7. The parametrizations of Flynn [20] seems to suggest that a Jacobian can only possess a $(5, 5)$ -isogeny in the presence of $\sqrt{5}$. Is this the case? If so, Algorithm 5.5.1 would only need to be run for $p = 2$ and $p = 3$ when the input curve is defined over \mathbb{Q} .

6.4 Isogenies besides (ℓ, ℓ) -isogenies

We saw that an isogeny whose domain is a suitably generic¹ Jacobian of a genus two curve and whose degree is not a square will have as its codomain an abelian surface which is not principally polarizable. Taking duals, the same statement is

¹absolutely simple, no extra endomorphisms

true when the Jacobian is the target of the isogeny. How can we describe these non-principally polarized varieties? I believe that the “challenge curve” $C_{1109} = C$ of Bloom, with LMFDB label `1109.a.1101.1` [36], has a Jacobian which possesses such a cover. The reduction mod p of its Jacobian is always a multiple of five, but the algorithm which calculates $(5, 5)$ -isogenies turned up nothing. Thus there should be a non-principally polarized abelian surface, defined over \mathbb{Q} and possessing a \mathbb{Q} -rational point of order five, for which the quotient is $\text{Jac}(C)$. As Igusa invariants only describe the moduli space of principally polarized abelian varieties, something else is needed to get our hands on this surface.

If we’re successful in calculating equations for covers, we could potentially determine the genus 6 curve which is an étale cover of C . This non-principally polarized variety will then accept a map from the Jacobian of the covering curve. Could that be leveraged somehow? What if we only care about the Galois representation of the cover at $p = 1109$ (a prime of bad reduction)?

6.5 Reducible Jacobians

Some curves of genus two have Jacobians which are *reducible*: they are isogenous to a product of elliptic curves. Two elliptic curves can be “glued together” to form the Jacobian of a genus two curve precisely when their n -torsion subgroups are isomorphic as $\text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ -modules, subject to a certain irreducibility criterion. This is a result of Kani (see, for instance, [21]). Explicit formulas are known for gluing two curves together along their 2-torsion [25] or their 3-torsion [4], and in [37] it

is suggested that the groundwork has been laid for generating formulae for gluing elliptic curves along isomorphic 5-torsion.

Can we use these constructions to create genus two curves with specified torsion over \mathbb{Q}^{ab} , or even better, specified primes dividing $\# \text{Ker}(C/\mathbb{Q})$? For instance, without using much thought we get the following result.

Proposition 6.1. *Let G be one of the groups which occurs in Chou’s classification [10] of possibilities for $\text{Tors}(E(\mathbb{Q}^{\text{ab}}))$. Then G occurs as a subgroup of $\text{Tors}(\text{Jac}(C)(\mathbb{Q}^{\text{ab}}))$ for some curve C of genus two.*

Proof. By results of Rubin and Silverberg [50], there is an infinite, explicit family of elliptic curves whose 2, 3, or 5-torsion is as a given elliptic curve. Let E be a curve with $\text{Tors}(E(\mathbb{Q}^{\text{ab}})) = G$. Crucially note that no member of Chou’s list is divisible by 30, so we can choose a prime $p \in \{2, 3, 5\}$ which does not divide $\#G$. Let E' be a curve from the Rubin-Silverberg family which satisfies $E[p] \cong E'[p]$ as Galois-modules with the property that E and E' are not isogenous over \mathbb{Q} . (Recall that \mathbb{Q} -isogeny classes are very small—at most eight curves!—so this is easy to arrange.) The lack of an isogeny between the elliptic curves guarantees that the irreducibility condition is satisfied, so Kani’s results say that $E \times E'$ possesses a (p, p) -isogeny to a Jacobian of a genus two curve. Call this curve C . Now the relevant torsion on E is prime to p and so descends to the quotient, hence is torsion of $\text{Jac}(C)$ defined over \mathbb{Q}^{ab} . \square

Pushing these ideas further would require determining something about what torsion is possible among curves occurring in the Rubin-Silverberg family. Is this feasible?

6.6 Torsion of Jacobians over \mathbb{Q}^{ab}

Mazur’s result on the possible torsion structures of an elliptic curve over \mathbb{Q} is a powerful tool that underlies a lot of the work that has been done for torsion of elliptic curves since. In particular, the isogeny theorem that states the possible orders of cyclic isogenies between elliptic curves defined over \mathbb{Q} has Mazur’s result as one of its cornerstones. This theorem in turn is a key ingredient in Chou’s classification of possible torsion structures for $E(\mathbb{Q}^{ab})$.

In genus two there is no corresponding theory; in fact, the situation is almost embarrassing—not only is there no finite list for the possibilities of $\text{Tors}(\text{Jac}(C))(\mathbb{Q})$, no one has yet been able to rule out any single number as never occurring. Howe in [24] says that a point of order n has been found on a genus 2 Jacobian for all $1 \leq n \leq 36$ except $n = 31$, as well as $n \in \{39, 40, 45, 48, 60, 63, 70\}$. Of course torsion over \mathbb{Q} certainly exists over \mathbb{Q}^{ab} . No attempt to collect data over \mathbb{Q}^{ab} has been done, as far as we are aware. It would be interesting to see to what extent we could use the promise of a result like the genus 1 isogeny theorem to attempt to find examples of torsion over abelian number fields. Even wandering in the dark base-changing curves to quadratic extensions might offer something new.

It’s a personal belief that $n \in \{42, 56, 72\}$ should be possible among the family of curves with reducible Jacobians; if the questions from the previous section had satisfactory answers then there would be hope of gluing two elliptic curves together along their 5-torsion, one of which possessed a point of order 8 and the other a point of order 9, for instance.

6.7 Determining $\text{Ker}(E/\mathbb{Z})$ or $\text{Ker}(E/\mathcal{O}_K)$

In §2.3 we presented an algorithm that calculates $\text{Ker}(E/\mathbb{Q})$, and mentioned in passing that it “really” calculates $\text{Ker}(E/\mathbb{Z}[\frac{1}{N}])$, where N is the product of the primes of bad reduction. Is it possible to say what is happening at the primes of bad reduction in order to provide a statement regarding $\text{Ker}(E/\mathbb{Z})$? We saw that examples with $\#\text{Ker}(E/\mathbb{Z}[\sqrt{29}]) > 1$ are possible, but is it possible over \mathbb{Z} ? Katz and Lang in [32] go through the example of the modular curve $C = X_0(N)$, which naturally has $X_1(N)$ as a cover, albeit not an étale one. The maximal étale subcover between $X_1(N)$ and $X_0(N)$ is called the *Shimura cover*, and is a cyclic torsor of $X_0(N)$ over \mathbb{Q} . However, they show $\text{Ker}(X_0(N)/\mathbb{Z})$ is trivial. So I don’t know an example of a nontrivial element of Ker over \mathbb{Z} .

The key to an answer may come from Bloch’s theory; see the following section.

6.8 Relationship to bigger theories

Two other connections should be pointed out before we close this section, both with similar flavor of passage to a higher codimension.

1. Although Katz and Lang discuss “geometric class field theory” in their paper defining Ker , the usual class group of algebraic number theory is not a Ker . However in Bloch’s paper [2] a *quotient* of π_1^{ab} does generalize the class group for arithmetic surfaces, which Bloch calls $F_0K_0(X)$. This group is that part of $K_0(X)$, the Grothendieck group of coherent \mathcal{O}_X -modules, whose support is

only on a set of closed points of X . Bloch proves that this group is finite; can it be calculated?

2. An intriguing paper of Rössler and Szamuely [49] point out that Ribet's result [48] that $\text{Tors}(A(\mathbb{Q}^{\text{ab}}))$ is finite can be interpreted in two interesting ways: first, as a statement about Galois-fixed points of certain étale cohomology groups, and second as a statement about the torsion of a Chow group $\text{CH}^1(C_{K(\mu)})$. They conjecture CH^i should also have finite torsion for $i > 1$. Is this accessible for $i = 2$?

Appendix A: Cardinality of $\text{Jac}(C)(\mathbb{F}_p)$

In this appendix we'll calculate the cardinality of the Jacobian of a curve over a finite field, in the process proving Proposition 4.3.4. The basic facts we need are the Weil conjectures and the relationship between the étale cohomology of a curve and its Jacobian. Throughout we assume that C is a curve of genus g over the finite field \mathbb{F}_q ; we do not demand that q is prime.

Definition A.1. Let X/\mathbb{F}_q be a variety. Its **zeta function** is the power series

$$Z(X, t) = \exp \left(\sum_{n=1}^{\infty} \#X(\mathbb{F}_{q^n}) \frac{t^n}{n} \right).$$

The following conjectures, proven by Dwork, Grothendieck, and Deligne, cover the properties of this function. (See Milne [42] for a longer introduction and the proof.)

Theorem A.2 (Weil Conjectures). *Suppose X/\mathbb{F}_q is a variety of dimension n .*

1. *The power series $Z(X, t)$ is a rational function in t . This rational function has a factorization of the form*

$$Z(X, t) = \frac{P_1(t) P_3(t) \cdots P_{2n-1}(t)}{P_0(t) P_2(t) \cdots P_{2n}(t)}$$

with each $P_i(t)$ a polynomial with integer coefficients. Write $P_i(t) = \prod (1 - \alpha_{ij}t)$ for complex numbers α_{ij} . $P_0(t)$ is $1 - t$ and $P_{2n}(t)$ is $1 - q^n t$.

2. There exists a number $E = E(X)$ whereby

$$Z(X, q^{-n}t) = \pm q^{\frac{nE}{2}} T^E Z(X, t).$$

This functional equation implies that α_{ij} is equal to $q^n \alpha_{(2n-i),k}^{-1}$ for some k .

3. The complex absolute value of a_{ij} is $q^{\frac{i}{2}}$ for every i, j .

When $X = C$ is a curve, the zeta function has only one mystery: the polynomial $P_1(t)$ in the numerator. This polynomial has an interpretation as the characteristic polynomial of the Frobenius automorphism of C acting on the first étale cohomology $H_{\text{ét}}^1(C, \mathbb{Q}_\ell)$. We won't pursue this here, except to note that this implies that $\deg P_1(t) = 2g$. As $i = 1$ is in the “middle” as far as Conjecture 2 is concerned, the $2g$ inverse roots of $P_1(t)$ must take the shape $\alpha_1, \dots, \alpha_g, \frac{q}{\alpha_1}, \dots, \frac{q}{\alpha_g}$.

Proposition A.3. *For a variety C of dimension one, the polynomial $P_1(t)$ takes the shape*

$$1 + a_1 t + a_2 t^2 + \dots + a_g t^g + q a_{g-1} t^{g+1} + q^2 a_{g-2} t^{g+2} + \dots + q^g t^g.$$

Proof. The coefficient on t^n is the sum of all products of (reciprocal) roots of cardinality n . For $k < g$, there is an obvious bijection between subsets of the set of roots of cardinality k and subsets of the set of roots of cardinality $g + k$ —map a set to its complement. Since the product of all the roots is q^g , this map will take a summand in the coefficient on t^k and send it to q^g divided by that summand, which is a constituent of the coefficient of t^{g+k} .

Meanwhile there is an involution on the collection of subsets of the roots of P_1 of cardinality n induced by $\beta \mapsto \frac{q}{\beta}$. This map will send a product of k roots

to q^k divided by that product, for a reason similar to the argument in the previous paragraph: the product of all the elements of a subset and its involute is q^k . It follows that a_k , the coefficient on t^k , can be expressed as

$$\sum_{\substack{S \subset \{\text{roots}\} \\ \#S=K}} \prod_{\beta \in S} \frac{q^k}{\beta}.$$

while the coefficient on t^{2g-k} can be expressed as

$$\sum_{\substack{S \subset \{\text{roots}\} \\ \#S=K}} \prod_{\beta \in S} \frac{q^{g-k}}{\beta}.$$

These coefficients differ by a factor of q^{g-k} , as needed. \square

This formula implies that the numbers a_1, \dots, a_g determine the entire zeta function, from which we can recover the cardinality of C over every extension of \mathbb{F}_q . It goes the other way as well:

Corollary A.4. *Let C/\mathbb{F}_q be a curve. The numbers $\#C(\mathbb{F}_q), \dots, \#C(\mathbb{F}_{q^g})$ determine $\#C(\mathbb{F}_{q^n})$ for all $n \geq 1$.*

Proof. Take logarithms on both sides of the equality

$$\exp \left(\sum_{n=1}^{\infty} \#C(\mathbb{F}_{q^n}) \frac{t^n}{n} \right) = \frac{\prod_{i=1}^{2g} (1 - \beta_i t)}{(1-t)(1-qt)}$$

and expand in power series centered at $t = 0$. (We don't need the extra information from Conjecture 2 here.) The left side comes expanded for us: the coefficient of t^n is $\frac{\#C(\mathbb{F}_{q^n})}{n}$. On the right, the usual formula for $\log(1-x)$ implies the coefficient on t^n is

$$\frac{q^n}{n} + \frac{1}{n} - \sum_{i=1}^{2g} \frac{\beta_i^n}{n}.$$

Thus $\#C(\mathbb{F}_{q^n}) = q^n + 1 - \sum_{i=1}^{2g} \beta_i^n$. Supposing the numbers $\#C(\mathbb{F}_{q^n})$ are known for $n = 1, \dots, g$, we now have access to the first g power sums $\sum_{i=1}^{2g} \beta_i^n$. Since the elementary symmetric polynomials can be recovered from the power sums, this gives us access to the first g coefficients of $P_1(t)$. By Proposition A.3 above, we now know all of $P_1(t)$, and hence $Z(t)$ has been recovered. \square

We now have a reasonable handle on the cardinality of the set of \mathbb{F}_q -rational points of a curve. What of its Jacobian? There are two key facts that are needed to link the theory we've developed to the theory we're seeking.

Theorem A.5. *Let A be an abelian variety over an algebraically closed field. The étale cohomology ring $H_{\text{ét}}^*(A)$ is the exterior algebra on $H_{\text{ét}}^1(A)$.*

Theorem A.6. *Let C be a curve over an algebraically closed field and J its Jacobian. The groups $H_{\text{ét}}^1(C)$ and $H_{\text{ét}}^1(J)$ are isomorphic.*

Proof. These facts are contained in one place in Poonen's notes on rational points on curves [47]. \square

The second fact means that $P_{1,C}(t) = P_{1,J}(t)$. The first fact then says that, if $\alpha_1, \dots, \alpha_{2g}$ are the roots of $P_{1,C}(t)$, the roots of $P_{n,J}(t)$ are the various n -fold products of the α_i .

Corollary A.7. $\# \text{Jac}(C)(F_q) = P_{1,C}(1)$.

Proof. As in the proof of Proposition A.4, the number of points of the Jacobian is $q + 1$ minus the sum of the roots of the polynomials $P_{i,J}(t)$, with signs taken to

account for whether the term comes from the numerator or the denominator. We such we get

$$\# \text{Jac}(C)(\mathbb{F}_q) = q + 1 - \sum_{i=1}^{2g} \alpha_i + \sum_{i=1}^{2g} \alpha_i \alpha_j - \dots$$

Since the i -th elementary symmetric polynomial in the roots of $P_{1,C}(t)$ is the i -th coefficient of $P_{1,C}(t)$, this sum is actually a gory way to write $P_{1,C}(1)$. \square

In theory we can use Newton's identities relating power sums and elementary symmetric polynomials to derive an expression for the cardinality of the Jacobian in terms of cardinalities of the curve. These get complicated quite quickly, but for $g = 2$ it is manageable.

Proposition A.8. *Let C/\mathbb{F}_q be a curve of genus two. Then*

$$\# \text{Jac}(C)(\mathbb{F}_q) = \frac{\#C(\mathbb{F}_q)^2 + \#C(\mathbb{F}_{q^2})}{2} - q.$$

Proof. Write $P_1(t) = 1 + at + bt^2 + apt^3 + p^2t^4$ for the common P_1 shared by C and $\text{Jac}(C)$. Let $\alpha_1, \alpha_2, \alpha_3, \alpha_4$ be the four (reciprocal) roots of P_1 . Write e_i for the i -th elementary symmetric function in the α_i and p_i for the i -th power sum of the α_i . We have

$$-a = e_1 = \alpha_1 + \alpha_2 + \alpha_3 + \alpha_4$$

$$b = e_2 = \alpha_1\alpha_2 + \alpha_1\alpha_3 + \alpha_1\alpha_4 + \alpha_2\alpha_3 + \alpha_2\alpha_4 + \alpha_3\alpha_4$$

$$\#C(\mathbb{F}_q) = q + 1 - p_1 = q + 1 + a$$

$$\#C(\mathbb{F}_{q^2}) = q^2 + 1 - p_2 = q^2 + 1 - (a^2 - 2b),$$

the final equality coming from Newton's relation $p_2 = e_1p_1 - 2e_2$. The term a^2 doesn't appear in our "target" $P(1) = 1 + a + b + ap + p^2$, so we'll eliminate it by

calculating

$$\#C(\mathbb{F}_q)^2 + \#C(\mathbb{F}_{q^2}) = (q^2 + 2q + 1 + 2aq + 2a + a^2) + (q^2 + 1 + 2b - a^2),$$

which simplifies to $2q^2 + 2aq + 2b + 2a + 2 + 2q$. Subtracting $2q$ and dividing by 2 hits our target. \square

I've never seen formulas for higher genus, but they can in theory be worked out in a similar fashion. The complexity of the set of divisor classes of degree zero that are linearly equivalent to multiple effective divisors grows as g does; this implies (thinking motivically) that the formulas ought to get more and more elaborate. For instance, if C is a curve of genus three, its Jacobian satisfies

$$\# \text{Jac}(C)(\mathbb{F}_q) = \frac{2\#C(\mathbb{F}_{p^3}) + 3\#C(\mathbb{F}_p)\#C(\mathbb{F}_{p^2}) + \#C(\mathbb{F}_p)^3 - 6p\#C(\mathbb{F}_p)}{6} - 1. \quad (\text{A.1})$$

The reader can check this by taking the identities

$$-a = e_1 = \alpha_1 + \cdots + \alpha_6$$

$$b = e_2 = \sum_{i < j} \alpha_i \alpha_j$$

$$-c = e_3 = \sum_{i < j < k} \alpha_i \alpha_j \alpha_k$$

$$\#C(\mathbb{F}_q) = q + 1 - p_1 = q + 1 + a$$

$$\#C(\mathbb{F}_{q^2}) = q^2 + 1 - p_2 = q^2 + 1 - (a^2 - 2b)$$

$$\#C(\mathbb{F}_{q^3}) = q^3 + 1 - p_3 = q^3 + 1 - (-a^3 + 3ab - 3c)$$

and plugging them into the right side of the expression above. One should get

$$\# \text{Jac}(C)(\mathbb{F}_q) = 1 + a + b + c + pb + p^2a + p^3$$

once the dust has settled.

Bibliography

- [1] C. Birkenhake and H. Lange. *Complex abelian varieties*, volume 302. Springer Science & Business Media, 2013.
- [2] S. Bloch. Algebraic k -theory and classfield theory for arithmetic surfaces. *Annals of Mathematics*, 114(2):229–265, 1981.
- [3] O. Bolza. On binary sextics with linear transformations into themselves. *American Journal of Mathematics*, pages 47–70, 1887.
- [4] R. Bröker, E. Howe, K. Lauter, and P. Stevenhagen. Genus-2 curves and jacobians with a given number of points. *LMS Journal of Computation and Mathematics*, 18(1):170–197, 2015.
- [5] R. Bröker and K. Lauter. Modular polynomials for genus 2. *LMS Journal of Computation and Mathematics*, 12:326–339, 2009.
- [6] R. Bröker and K. Lauter. Evaluating igusa functions. *Mathematics of Computation*, 83(290):2977–2999, 2014.
- [7] N. Bruin, E. Flynn, and D. Testa. Descent via $(3, 3)$ -isogeny on jacobians of genus 2 curves. *Acta Arithmetica*, 3(165):201–223, 2014.
- [8] G. Cardona. Models racionales de corbes de genere 2. *Doctoral Theses, Universitat Politècnica de Catalunya*, 2001.
- [9] JWS Cassels and EV Flynn. *Prolegomena to a middlebrow arithmetic of curves of genus 2*, volume 230. Cambridge University Press, 1996.
- [10] M. Chou. Torsion of rational elliptic curves over the maximal abelian extension of \mathbb{Q} . *ArXiv e-prints*, November 2017.
- [11] P. L. Clark. Quadratic forms, chapter i: Witt’s theory. *URL <http://math.uga.edu/~pete/quadraticforms.pdf>*, 2013.

- [12] H. Cohen. *A course in computational algebraic number theory*, volume 138. Springer Science & Business Media, 2013.
- [13] S. Comalada. Elliptic curves with trivial conductor over quadratic fields. *Pacific Journal of Mathematics*, 144(2):237–258, 1990.
- [14] K. Coombes and D. Grant. On heterogeneous spaces. *Journal of the London Mathematical Society*, 2(3):385–397, 1989.
- [15] D. Cox. *Primes of the form $x^2 + ny^2$* . Wiley-Interscience, 1997.
- [16] J. Cremona. *Algorithms for Modular Elliptic Curves*. CUP Archive, 1997.
- [17] R. Dupont. Moyenne arithmético-géométrique, suites de borchardt et applications (ph.d. thesis). *These de doctorat, École polytechnique, Palaiseau*, 2006.
- [18] N. Elkies. The existence of infinitely many supersingular primes for every elliptic curve over \mathbb{Q} . *Inventiones mathematicae*, 89(3):561–567, 1987.
- [19] G. Faltings. Endlichkeitssätze für abelsche varietäten über zahlkörpern. *Inventiones mathematicae*, 73(3):349–366, 1983.
- [20] E. Flynn. Descent via $(5, 5)$ -isogeny on jacobians of genus 2 curves. *Journal of Number Theory*, 153:270–282, 2015.
- [21] G. Frey and E. Kani. Curves of genus 2 covering elliptic curves and an arithmetical application. In *Arithmetic algebraic geometry*, pages 153–176. Springer, 1991.
- [22] E. González-Jiménez and F. Najman. Growth of torsion groups of elliptic curves upon base change. *arXiv preprint arXiv:1609.02515*, 2016.
- [23] R. Hartshorne. *Algebraic geometry*, volume 52. Springer Science & Business Media, 2013.
- [24] E. Howe. Genus-2 jacobians with torsion points of large order. *Bulletin of the London Mathematical Society*, 47(1):127–135, 2015.
- [25] E. Howe, F. Leprévost, and B. Poonen. Large torsion subgroups of split jacobians of curves of genus two or three. In *Forum Mathematicum*, volume 12, pages 315–364. Walter de Gruyter GmbH & Co. KG., 2000.
- [26] J. Igusa. Arithmetic variety of moduli for genus two. *Annals of Mathematics*, pages 612–649, 1960.
- [27] J. Igusa. Modular forms and projective invariants. *American Journal of Mathematics*, 89(3):817–855, 1967.
- [28] D. Jeon, C. Kim, and A. Schweizer. On the torsion of elliptic curves over cubic number fields. *Acta Arithmetica*, 113:291–301, 2004.

- [29] T. Kagawa. Determination of elliptic curves with everywhere good reduction over (37) . *Acta Arithmetica*, 83(3):253–269, 1998.
- [30] T. Kagawa. *Elliptic curves with everywhere good reduction over real quadratic fields*. PhD thesis, Ph. D Thesis, Waseda University, 1998.
- [31] N. Katz. Galois properties of torsion points on abelian varieties. *Inventiones mathematicae*, 62(3):481–502, 1980.
- [32] N. Katz and S. Lang. Finiteness theorems in geometric class field theory. *Enseign. Math.*, 2:285–314, 1981.
- [33] M.A. Kenku. On the number of q -isomorphism classes of elliptic curves in each q -isogeny class. *Journal of Number Theory*, 15(2):199–202, 1982.
- [34] M.A. Kenku and F. Momose. Torsion points on elliptic curves defined over quadratic fields. *Nagoya Mathematical Journal*, 109:125–149, 1988.
- [35] R. Livné. Cubic exponential sums and galois representations. *Contemp. Math*, 67:247–261, 1987.
- [36] The LMFDB Collaboration. The l -functions and modular forms database. <http://www.lmfdb.org>, 2013. [Online; accessed 16 September 2013].
- [37] K. Magaard, T. Shaska, , and H. Völklein. Genus 2 curves that admit a degree 5 map to an elliptic curve. In *Forum Mathematicum*, volume 21, pages 547–566. Walter de Gruyter GmbH & Co. KG, 2009.
- [38] J-F. Mestre. Construction de courbes de genre 2 à partir de leurs modules. In *Effective methods in algebraic geometry*, pages 313–334. Springer, 1991.
- [39] E. Milio. A quasi-linear time algorithm for computing modular polynomials in dimension 2. *LMS Journal of Computation and Mathematics*, 18(1):603–632, 2015.
- [40] J. Milne. Jacobian varieties, arithmetic geometry (cornell and silverman, ed.), 1986.
- [41] J. Milne. Abelian varieties (v2.00), 2008. Available at www.jmilne.org/math/.
- [42] J. Milne. Lectures on étale cohomology (v2.21), 2013. Available at www.jmilne.org/math/.
- [43] D. Mumford. Abelian varieties, 1970.
- [44] F. Najman. Torsion of rational elliptic curves over cubic fields and sporadic points on $x_1(n)$. *Math. Res. Letters*, 23:245–272, 2016.
- [45] S. Petersen. The rank of hyperelliptic jacobians in families of quadratic twists. *Journal de théorie des nombres de Bordeaux*, 18(3):653–676, 2006.

- [46] B. Poonen. non principally polarized complex abelian varieties. MathOverflow. URL:<https://mathoverflow.net/q/17014> (version: 2010-03-03).
- [47] B. Poonen. Lectures on rational points on curves (2006 ver). URL <http://math.mit.edu/~poonen/papers/curves.pdf>, 2006.
- [48] K. Ribet. Appendix to finiteness theorems in geometric class field theory. *Enseign. Math.*, 2:315–318, 1981.
- [49] D. Rössler and T. Szamuely. Cohomology and torsion cycles over the maximal cyclotomic extension. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 2015.
- [50] K. Rubin and S. Silverberg. Families of elliptic curves with constant mod p representations. In *Conference on Elliptic Curves and Modular Forms*, pages 148–161. International Press, 1995.
- [51] J. Silverman. *The arithmetic of elliptic curves*, volume 106. Springer Science & Business Media, 2009.
- [52] M. Streng. *Complex multiplication of abelian surfaces (Ph.D. thesis)*. Mathematical Institute, Faculty of Science, Leiden University, 2010.
- [53] A. Sutherland. On the evaluation of modular polynomials. *The Open Book Series*, 1(1):531–555, 2013.
- [54] P. Van Wamelen. Examples of genus two cm curves defined over the rationals. *Mathematics of Computation of the American Mathematical Society*, 68(225):307–320, 1999.
- [55] J. Vélú. Isogénies entre courbes elliptiques. *C.R. Acad. Sc. Paris, Série A*, 273:238–241, 1971.
- [56] L. Washington. *Elliptic curves: number theory and cryptography, 2nd edition*. Chapman and Hall/CRC, 2008.
- [57] Y. Zhao. Elliptic curves over real quadratic fields with everywhere good reduction and a non-trivial 3-division point. *Journal of Number Theory*, 133(9):2901–2913, 2013.