Università di Pisa

Facoltà di Matematica

Counting Points on Finite Fields

A Classical Proof of Tate's Isogeny Theorem for Elliptic Curves

Tesi di Laurea Triennale
in Matematica

Candidato
Davide Ranieri

Relatore
Prof. Davide Lombardo
Università di Pisa

Anno Accademico 2023 - 2024

Contents

Introduction

The goal of this work is to give a motivated proof of an important theorem in the theory of elliptic curves over finite fields. This theorem, due to John T. Tate ([Tat66]) in its original form, was proven more generally for abelian varieties using difficult algebro-geometric machinery. We show, following a suggestion of Tate himself[1], that in the special case of dimension one it is possible to get a very close result using only classical tools. Specifically, we approach the theory of elliptic curves from the Zariski-Weil point of view of varieties, circumventing the need for schemes. The deepest fact from algebraic geometry proper that we shall require is the Riemann-Roch theorem. Most of our tools are standard and can be found in any one of the usual sources, such as J. Silverman's landmark text [Sil09]. We supplement this typical arsenal with a now superseded – but, crucially, accessible by our methods – result by Deuring ([Deu41]) concerning the lifting of endomorphisms from positive to zero characteristic. This theorem requires in turn some analytic ingredients in order to establish a certain integrality result. Namely, we will exploit well-known computations with modular forms to show that the j-invariants of two elliptic curves connected by a nonconstant isogeny satisfy a polynomial relation with integer coefficents. Although this last result also follows from more sophisticated methods, we stress that all our arguments use little more than what is considered 'basic' elliptic curve theory. Our work is laid down in three chapters. In the first, we give the primary definitions and recall the most important facts about the objects we shall be dealing with. We introduce elliptic curves along with their natural group structure, isogenies as the corresponding maps between them, and discuss how duality ties into the structure of endomorphism rings. The second chapter specialises to the realm of finite fields, where several unique phenomena such as the Frobenius and supersingularity emerge. We define the Tate module and state the homonimous theorem, along with a memorable consequence that helps reinforce the importance of said result: two elliptic curves over a finite field admit a nontrivial isogeny between them if and only if they have the same number of points over that field. Finally, in the last chapter, we briefly switch to the analytic point of view and outline the tools we need to borrow from the complex theory before diving into the proof of the lifting theorem. After this detour we deliver on our promise by giving a simple proof of Tate's theorem along with its corollary.

[1] "In case A' and A'' are elliptic curves this theorem is an easy consequence of results of Deuring, as Mumford pointed out to me four years ago."

Euclid alone has looked on Beauty bare.
Let all who prate of Beauty hold their peace,
And lay them prone upon the earth and cease
To ponder on themselves, the while they stare
At nothing, intricately drawn nowhere
In shapes of shifting lineage; let geese
Gabble and hiss, but heroes seek release
From dusty bondage into luminous air.
O blinding hour, O holy, terrible day,
When first the shaft into his vision shone
Of light anatomized! Euclid alone
Has looked on Beauty bare. Fortunate they
Who, though once only and then but far away,
Have heard her massive sandal set on stone.

Edna St. Vincent Millay

CHAPTER 1

Background Results on Elliptic Curves

1.1 Basics of curves

To begin our journey, we first recall some basic algebro-geometric notions. All results in this chapter are standard and most will be cited without proof. Where a reference is not provided, the interested reader may consult any one of the usual texts, such as [Mum99]. The vast majority of our material on elliptic curves, in this as well as the subsequent chapters, is taken from [Sil09].
This work aims to use only (or mostly) classical tools. We shall therefore refrain from employing scheme theory except at times in passing remark where we feel it can offer an alternative point of view for the trained reader.
The term "curve" will be always tacitly taken to mean "smooth geometrically connected projective curve" unless otherwise stated. We now recollect some standard definitions and facts we shall freely employ in the sequel.

Definition 1.1. For a divisor D on the curve C, denote by $H^0(D)$ the $\overline{K}$-vector space

$$\{f \in \overline{K}(C) : \mathrm{div} f + D \geq 0\}$$

and by $h^0(D)$ its dimension.

Proposition 1.2. *Let ω be a regular differential form on C. There is a natural way to associate to ω a divisor K_C such that any other choice of differential form yields a linearly equivalent divisor. The equivalence of class of K_C is known as the* canonical class *of C.*

Definition 1.3. Let C be a curve and K_C a divisor in its canonical class. The *(geometric) genus*[1] of C is defined to be $h^0(K_C)$.

Theorem 1.4 (Riemann-Roch)**.** *Let C be a curve of genus g, D any divisor on C and K_C a representative of the canonical class. Then*

$$h^0(D) - h^0(K_C - D) = \deg D - g + 1.$$

[1] This is a rather 'cheap' definition - one should instead work with the arithmetic genus $h^1(\mathcal{O}_C)$, but discussing this would lead us too far astray.

Proof. The standard proof uses sheaf cohomology and may be found in any suitably advanced text in algebraic geometry. For an alternative proof devised by Tate involving a clever analysis of the adèle ring of the function field of C see [Tat68]. □

The importance of this theorem cannot be overstated. As one of its many consequences, setting $D = K_C$ above yields

Corollary 1.5. $\deg K_C = 2g - 2$.

1.2 Elliptic curves and the group law

We now come to defining the central object of our study:

Definition 1.6. An *elliptic curve* over K is a curve E/K of genus 1 together with a distinguished point $O \in E(K)$.

Remark 1.7. It is important to include some K-point O in the data of the elliptic curve, even though we shall see that it does not actually matter which specific point is chosen. The reason for doing so is that, as an abstract curve, E might be defined over K without possessing a single K-rational point, and we do not wish to consider such a curve as being an elliptic curve over K.

The reader might be wondering what is so special about genus 1. There are many ways to spin this tale, but a simple way to shed some light is as follows. Mathematicians of the past became interested in the geometry of curves while trying to solve equations like $x^3 + y^3 = z^3$ or $x^4 + y^4 = z^2$ over various number systems, most notably the rationals. The simplest case, that of quadratic equations, is easy to handle: either there are no solutions, or finding one solution allows one to produce all the others via a straightforward geometric procedure. Indeed, choosing one smooth point and drawing lines through it, it is not difficult to show that each such line intersects the curve at exactly two points (possibly counted with multiplicity). Thus, varying the angle, we get a projective line's worth of points on the curve - we have just shown that all nonsingular conics have genus 0. The next simplest case is that of (smooth) plane cubics. An analogous projection argument shows

Theorem 1.8 (Genus-Degree Formula)**.** *Let C be a smooth plane curve of degree n and genus g. Then*

$$g = \frac{(n-1)(n-2)}{2}$$

Therefore our cubic has genus 1. While we cannot simply use the knowledge of one point to find all the others, we can use two points to find a third, since the generic line intersects the curve in three points (once more, care must be taken to make this vague statement formally correct). The keen reader may have noted that such a 'two in, one out' procedure is reminiscent of a group law – and indeed, we will soon see how a suitable modification equips the cubic with the structure of an abelian group defined purely using geometry. This structure is, among curves, entirely peculiar to the realm of genus 1, and is the chief reason why it is so worthwile studying.

Before turning to the construction of the group law, we need to find suitable models for the so-far abstract elliptic curves. Strictly speaking, they are not necessary to show the existence of the group structure on the curve; however, they allow the addition of two points to be described concretely as well as unlocking other theoretical results that shall be discussed later. We have stated above that any smooth plane cubic is a genus 1 curve, so that selecting one of its K-points makes it into an elliptic curve. Somewhat remarkably, the converse also holds: every elliptic curve may be written as a plane cubic in a special form.

Theorem 1.9 (Existence of the Weierstraß form)**.** *Any elliptic curve $(E/K, O)$ is isomorphic over K to the projective completion of the affine curve given by an equation of the form*

$$y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6 \tag{1.1}$$

with the isomorphism sending O to the unique point at infinity of the latter curve. Moreover if $\operatorname{char} K \neq 2$ *the equation may be taken to be of the form*

$$y^2 = x^3 + b_2x^2 + b_4x + b_6 \tag{1.2}$$

and if furthermore $\operatorname{char} K \neq 3$ *the equation may be taken to be of the form*

$$y^2 = x^3 + c_4x + c_6. \tag{1.3}$$

Proof sketch. Using Theorem 1.4 we compute

$$h^0(mO) = m$$

In particular no function on E has a single simple pole at O. This implies that there exist functions x and y whose pole divisors are $2(O)$ and $3(O)$ respectively. Now the divisor $3(O)$ is very ample, so it gives an embedding $E \hookrightarrow \mathbb{P}^2K$. On the other hand, we can easily produce seven elements of the six-dimensional space $H^0(6O)$, namely y^2, xy, y, x^3, x^2, x and 1. By dimension count there must be a linear dependence among them, which cuts out the desired equation for the embedding in the plane. The reduction to the two simpler forms is elementary algebraic manipulation. □

We are now ready to illustrate the group structure on an elliptic curve.

Construction 1.10. Let E/K be an elliptic curve embedded as a smooth cubic in the plane (possibly, but not necessarily, in Weierstraß form) with distinguished point O. Define a map

$$m : E(\overline{K}) \times E(\overline{K}) \longrightarrow E(\overline{K})$$

as follows: let P, Q be not necessarily distinct points on E. There is a unique line in the plane whose intersection with the curve is given, as a divisor, by $(P) + (Q) + (R)$ where R is some third point on E, possibly coinciding with P or Q. We can identify this line as 'the line through P and Q' with the caveat that when $P = Q$ it should be taken to mean 'the tangent line to the curve at P'. The same procedure is now applied with R and O in place of P and Q, yielding a new point S. Finally, we set

$$P + Q \stackrel{\text{def}}{=} m(P, Q) = S.$$

For later purposes we also denote by

$$i : E(\overline{K}) \longrightarrow E(\overline{K})$$

the map given by the second step in the previous definition, that is, by sending P to the third point on the line through P and O.

It is clear that the above operation is commutative. With more work, one establishes the following

Theorem 1.11.

1. *$(E, +)$ is an abelian group whose identity element is O and with inverses given by i.*

2. *m and i are morphisms of varieties, making E into an* algebraic group*.*

3. *The structure of algebraic group is independent of the choice of identity point and embedding, in the sense that another choice for these data yields isomorphic groups (via an algebraic isomorphism).*

Remark 1.12. For any field extension L/K and any $P, Q \in E(L)$ the line through them admits a parametric equation with coefficients in L by high-school algebra. Substituting this equation into the equation for E gives a cubic polynomial with coefficients in L, two of whose roots – corresponding to P and Q – are known to lie in L. Therefore the third intersection point also belongs to $E(L)$. A similar argument shows the same to be true in the limit case $P = Q$. This implies that $E(L)$ is a subgroup of $E(\overline{L})$, a fact that will turn out to be important later.

Let us delve a little deeper into the last statement of the above theorem. Defining addition in terms of intersections with lines is somewhat unsatisfactory - it requires fixing a plane embedding of E even though a posteriori said choice makes no difference. One would hope for an intrinsic construction depending solely on the abstract geometry of E.

Construction 1.13. Let (E, O) be as usual. Given points P, Q on E consider the divisor $D = (P) + (Q) - (O)$ of degree 1. By Theorem 1.4 there is, up to a scalar factor, a unique rational f such that $\mathrm{div} f + D \geq 0$. Counting degrees, we see that the left hand side must be an effective divisor of degree 1, i.e. a point R of E. We define the result of $P + Q$ to be R.

Checking that the two provided constructions result in the same group structure amounts to proving

Proposition 1.14. *The map*

$$\begin{aligned} E &\longrightarrow \mathrm{Pic}^0(E) \\ P &\longmapsto [(P) - (O)] \end{aligned}$$

is a group isomorphism.

The inverse morphism is given by a procedure analogous to our second construction above.

Remark 1.15. With more technology it is possible to describe the group law even more concisely. The choice of a point $O \in E(K)$ gives rise to the Jacobian embedding

$$E \hookrightarrow \mathrm{Jac}(E)$$

Since the genus of E equals $\dim \mathrm{Jac}(E) = 1$ the embedding is in fact an isomorphism of varieties, enabling us to transport the addition back to the original curve. This also explains Remark 1.7: while the Jacobian itself is defined over the same field as E, the embedding may not be. We shall not pursue this approach further because defining Jacobians over arbitrary fields is rather cumbersome.

Henceforth we shall suppress explicit mention of O unless required. We will also take the liberty to denote the identity point of any elliptic curve by O wherever doing so does not cause confusion.

1.3 Isogenies

One of the overarching themes in mathematics is that a good theory should be built up by identifying suitable objects and natural maps between them. We have seen that elliptic curves possess a dual nature of geometric and algebraic objects, so we are motivated to investigate the corresponding types of map:

Definition 1.16. An *isogeny* between two elliptic curves E_1, E_2 is a map

$$E_1 \longrightarrow E_2$$

which is simultaneously a morphism of varieties and groups.
We say that E_1 is *isogenous* to E_2 if there exists a nonconstant isogeny from E_1 to E_2.

It follows from the general theory of curves that an isogeny is either constant or surjective, and in the latter case all its fibres are finite of size bounded by its degree.
It is a remarkable fact that the definition of an isogeny is equivalent to a seemingly much weaker property:

Theorem 1.17. *Any morphism of varieties between elliptic curves mapping the identity to the identity is an isogeny*

Proof idea. The core of the proof is to establish the following commutative diagram:

$$\begin{array}{ccc} E_1 & \xrightarrow{\sim} & \mathrm{Pic}^0(E_1) \\ \downarrow{\scriptstyle \phi} & & \downarrow{\scriptstyle \phi_*} \\ E_2 & \xrightarrow{\sim} & \mathrm{Pic}^0(E_2) \end{array}$$

where the map on the right is given by functoriality of Pic^0. □

We have decided to highlight the commutative square above because it will be useful in the future and also for its conceptual significance. Another functorial correspondence which is crucial for our development is that given by function fields:

Theorem 1.18 (Curve-Field correspondence)**.** *Let K be a field. The functor $K(\cdot)$ sending a curve C/K to its field $K(C)$ of rational functions with coefficients in K is a (contravariant) equivalence between the categories of curves over K and field extensions[2] L of K of transcendence degree 1 such that K is algebraically closed in L. If a map $\phi : C_1 \to C_2$ is nonconstant (so as to have a well-defined degree) its degree equals the field-theoretic degree of the extension $K(C_2)/\phi^* K(C_1)$*

This motivates the following

Definition 1.19. If the field extension corresponding to an isogeny $\phi : E_1 \to E_2$ has a certain property (e.g. is separable, finite, Galois…) we say ϕ has the same property.

In particular great care will have to be placed in dealing with separability; because it is a phenomenon exclusive to positive characteristic it does not arise in the classical theory of curves over $\mathbb{C}$ (or, what amounts to the same thing, algebraically closed fields of characteristic zero). Therefore we will have to modify some results to account for this, and always be wary of the applicability of a line of thought guided by characteristic zero intuition. As an example, the well-known Riemann-Hurwitz formula only holds in the general setting for separable maps of curves.
A useful criterion for detecting separability is the following. Recall that, by Theorem 1.4, $\deg K_E = 0$ for an elliptic curve E. Using explicit charts (such as the ones given by a Weierstraß equation) one can show more: in fact, for a nonzero regular differential form ω (which is unique up to a scalar factor), $\operatorname{div}\omega = 0$. From this fact it is not hard to deduce

[2] Caveat: we must include the 0 morphism of rings as a valid map of K-fields to account for constant morphisms of curves

Proposition 1.20. *For a point P on E denote by τ_P the translation-by-P map, i.e. $Q \longmapsto Q + P$. Then*

$$\tau_P{}^*\omega = \omega$$

In the light of the above proposition we will sometimes refer to ω as an *invariant* differential form. The interplay between forms and isogenies is now summarised by

Theorem 1.21. *Let ω' and ω be invariant differential forms on E' and E respectively. Also, let $\phi : E' \longrightarrow E$ be an isogeny.*

1. *$\phi^*\omega = a\omega'$ for some $a \in \overline{K}$. This yields a map $\phi \longmapsto a_\phi$.*
2. *If $\psi : E' \longrightarrow E$ is another isogeny then*

$$a_{\phi+\psi} = a_\phi + a_\psi.$$

3. *ϕ is separable if and only if $a_\phi \neq 0$.*

Before moving on we mention a technical statement that will be needed at a later stage, namely, the existence of quotients by finite subgroups.

Theorem 1.22. *Let E/K be an elliptic curve and $H \subset E(K)$ a finite subgroup. There exist an elliptic curve E'/K, unique up to isomorphism, and a separable isogeny $\phi : E \longrightarrow E'$ such that* $\ker \phi = H$.

1.4 Dual isogenies and endomorphisms

So far we have considered the covariant mapping as in Theorem 1.17. However an isogeny also induces a map in the opposite direction

$$E_2 \xrightarrow{\sim} \mathrm{Pic}^0(E_2) \xrightarrow{\phi^*} \mathrm{Pic}^0(E_1) \xrightarrow{\sim} E_1$$

given by pullback of divisors. The composite map $\widehat{\phi}$ has, a priori, no reason to be more than a map of sets. Surprisingly, much more is true:

Theorem 1.23. *Let $\phi : E_1 \longrightarrow E_2$ be a nonconstant isogeny.*

1. *The map $\widehat{\phi}$ is an isogeny satisfying*

$$\phi \circ \widehat{\phi} = \widehat{\phi} \circ \phi = [\deg \phi]$$

   and it is uniquely determined by this property.

2. *If $\psi : E_2 \longrightarrow E_3$ is another isogeny then*

$$\widehat{\phi \circ \psi} = \widehat{\psi} \circ \widehat{\phi}.$$

3. *If $\psi : E_1 \longrightarrow E_2$ is another isogeny then*

$$\widehat{\phi + \psi} = \widehat{\phi} + \widehat{\psi}.$$

4. $\deg \widehat{\phi} = \deg \phi$ *and* $\widehat{\widehat{\phi}} = \phi$. *Furthermore, for an integer m, $\widehat{[m]} = [m]$, therefore* $\deg [m] = m^2$

Definition 1.24. We call $\widehat{\phi}$ the *dual isogeny* of ϕ. We also set $\widehat{0} = 0$ as a matter of convention.

Corollary 1.25. *If E_1 is isogenous to E_2 then E_2 is also isogenous to E_1. Thus isogeny is an equivalence relation (reflexivity and transitivity being obvious), and we shall from now on simply say E_1 and E_2 are isogenous.*

Remark 1.26. The emergence of the dual isogeny feels somewhat out of the blue in our elementary tractation. For an adequate explanation in terms of the Picard scheme and the theory of duality for abelian varieties see [Mum08]

Theorem 1.23 is a powerful computational tool. Indeed, it tells us that the degree function behaves in some sense like a positive definite quadratic form, since the pairing

$$\langle \phi, \psi \rangle = \deg(\phi + \psi) - \deg(\phi) - \deg(\psi) = \widehat{\psi} \circ \phi + \widehat{\phi} \circ \psi$$

is bilinear. It also imposes a strong restriction on the possible isomorphism type of the endomorphism ring of an elliptic curve, to which we now turn.
We consider the set $\mathrm{End}(E)$ of isogenies from E to itself. It possesses a ring structure given by pointwise addition on the curve and composition. The theorem may then be recast by saying that this ring is endowed with an antihomomorphism $\widehat{\cdot}$ to itself which is an involution and gives rise to a positive definite pairing as above. Such a map is analogous to conjugation in Cayley-Dickson algebras. Indeed, its existence implies that $\mathrm{End}(E)$ falls into three possible cases:

Theorem 1.27. *The endomorphism ring of E is, up to isomorphism, one of the following:*

- $\mathbb{Z}$
- *An* order *(i.e. subring of maximal rank) in an imaginary quadratic number field*
- *An order in a quaternion $\mathbb{Q}$-algebra*

In characteristic zero the last case never occurs, for theorem Theorem 1.21 provides an injection $\mathrm{End}(E) \hookrightarrow \overline{K}$. In the next chapter we will be able to say more about its occurrence in characteristic p.
To conclude this section we discuss an important computation concerning the multiplication-by-m endomorphisms. Leveraging Theorem 1.23 it is possible to refine the result $\deg[m] = m^2$ to give a description of its kernel, the m-torsion subgroup $E[m] \stackrel{\text{def}}{=} E(\overline{K})[m]$.

Proposition 1.28.

1. *If $(m, \mathrm{char}\, K) = 1$ then $E[m] \simeq \mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}$*
2. *If $\mathrm{char}\, K = p$ then one of the following occurs:*
   - $E[p^k] = 0 \ \forall k$
   - $E[p^k] = \mathbb{Z}/p^k\mathbb{Z} \ \forall k$

It will turn out that the dichotomy for p-primary torsion is intimately connected to the structure of $\mathrm{End}(E)$.

1.5 Isomorphism and the j-invariant

Whenever a new category $\mathcal{C}$ is introduced, one is naturally led to investigate the classification problem for that category: when are two objects in $\mathcal{C}$ isomorphic? Is it possible to find reasonably simple invariants for these objects? In algebraic and arithmetic geometry, these problems are usually far from tractable. However, for elliptic curves over a field K, there does exist a remarkably explicit complete invariant, at least for isomorphism over $\overline{K}$. We state the result for curves in 'short' Weierstraß form simply for convenience – there is a corresponding statement valid for all elliptic curves.

Theorem 1.29. *Let E be the projective completion of the affine curve*

$$y^2 = x^3 + Ax + B.$$

Then E is nonsingular if and only if the discriminant

$$\Delta = -16(4A^3 + 27B^2)$$

is nonzero. In this case E is an elliptic curve (by placing O at the point at infinity) and the quantity

$$j = -1728\frac{(4A)^3}{\Delta},$$

*known as the j-*invariant of *E, does not depend on the choice of Weierstraß equation for E. Moreover, two elliptic curves are isomorphic over $\overline{K}$ if and only if their j-invariants are equal.*

Remark 1.30. In general, over a non algebraically closed field K, two elliptic curves with the same j-invariant may not be isomorphic over K itself.

We may thus view j as a function from the set of $\overline{K}$-isomorphism classes of elliptic curves over K to K. It will prove useful later to know that this function is surjective:

Proposition 1.31. *Let $t \in K$. If $t \neq 0, 1728$ the equation*

$$y^2 + xy = x^3 - \frac{36}{t - 1728}x - \frac{1}{t - 1728}$$

defines an elliptic curve having j-invariant t.
Otherwise, one of the following two equations (depending on $\operatorname{char} K$*) has the desired property:*

$y^2 + y = x^3$	$\Delta = -27$	$j = 0$
$y = x^3 + x$	$\Delta = -64$	$j = 1728$

CHAPTER 2

Elliptic Curves over Finite Fields

2.1 The Frobenius and supersingularity

Consider a finite field $\mathbb{F}_q$, with $q = p^k$ a prime power. It is well-known that in characteristic p the map $x \longmapsto x^p$ is a field homomorphism; moreover, the construction of the finite fields implies that each finite extension $K/\mathbb{F}_q$ is Galois with cyclic Galois group generated by the q-th power map. This map is commonly referred to as the *Frobenius automorphism*, or simply the Frobenius. Let now X be a variety defined over $\mathbb{F}_q$. The Frobenius naturally acts on the set of points $X(\mathbb{F}_q)$ coordinate-wise. It is also straightforward to show that the resulting map is in fact a morphism of varieties. In the special case of elliptic curves we can be more precise:

Proposition 2.1. *Let E be an elliptic curve over $\mathbb{F}_q$. The map induced by the Frobenius*

$$\mathfrak{f} : E \longrightarrow E$$

is a purely inseparable isogeny of degree q.

We shall follow the literature standard and also call the above map Frobenius, taking care to make it as clear as possible which type we are referring to.
The Frobenius on an elliptic curve E encodes many of its arithmetic properties. We shall see that it is the main player in the proof of our main theorem. First, however, we return to the discussion initiated in the previous chapter.

Theorem 2.2. *Let E/K be an elliptic curve with* $\operatorname{char} K = p$. *The following are equivalent:*

1. $\operatorname{End}(E)$ *is an order in a quaternion algebra.*
2. $E[p] = 0$
3. $\widehat{\mathfrak{f}}$ *is (purely) inseparable*
4. $[p]$ *is purely inseparable and* $j(E) \in \mathbb{F}_{p^2}$

Definition 2.3. An elliptic curve with the above properties is said to be *supersingular*. Otherwise, it is *ordinary*.

Remark 2.4. The last statement in the equivalence implies that for fixed p there are only finitely many supersingular elliptic curves in characteristic p up to isomorphism (over the algebraic closure). It is then an interesting endeavour to write down the complete list of j-invariants of supersingular curves for a given prime p. They may be described as the roots

of the so-called *Atkin polynomials* with deep ties to the theory of modular forms which will not be pursued here; see [KZ98] for an overview of the topic.

Excluding the supersingular case, the remaining elliptic curves in positive characteristic have a commutative endomorphism ring. It turns out that in the case we shall be concerned with, that of curves over $\overline{\mathbb{F}}_p$, it is never trivial.

Proposition 2.5. *If the curve $E/\overline{\mathbb{F}}_p$ is ordinary* $\mathrm{End}(E)$ *is strictly larger than $\mathbb{Z}$.*

Idea of Proof. The main observation is that for an ordinary curve the Frobenius cannot be an integer. □

This fact will be exploited in the next chapter when we choose which endomorphisms to lift to positive characteristic.

2.2 The Tate Module and the Main Theorem

Recall that for any integer m which is nonzero in K the group $E[m]$ is a free $\mathbb{Z}/m\mathbb{Z}$-module of rank 2 equipped with a linear action of the absolute Galois group Γ_K. We wish to extract information about the curve from this representation; however, the presence of torsion introduces pathological behaviour that obscures the view. To 'rectify' the representation to characteristic zero we mimic the construction of the p-adic integers by taking limits:

Definition 2.6. Fix an elliptic curve E/K and a prime $\ell \neq \mathrm{char}\, K$ (if K has characteristic zero this condition is of course vacuous). The *ℓ-adic Tate module* $T_\ell(E)$ of E is defined to be the limit of the diagram

$$E[\ell] \xleftarrow{\cdot\ell} E[\ell^2] \xleftarrow{\cdot\ell} E[\ell^3] \xleftarrow{\cdot\ell} \cdots$$

It is clear from the definition, together with Proposition 1.28, that $T_\ell(E)$ is a free $\mathbb{Z}_\ell$-module of rank two, so that upon fixing a basis $T_\ell(E) \simeq \mathbb{Z}_\ell^2$. The Galois actions on the various $E[\ell^k]$ are compatible with multiplication by ℓ, therefore the Tate module inherits an action of Γ_K by $\mathbb{Z}_\ell$-linear automorphisms. Moreover, formation of the Tate module is functorial. Given an isogeny $\phi : E_1 \longrightarrow E_2$ we have $\phi\left(E_1[\ell^k]\right) \subset E_2[\ell^k]$, and passing to the limit we obtain an induced linear map

$$\phi_\ell : T_\ell(E_1) \longrightarrow T_\ell(E_2)$$

Although it would appear that we lose a lot of information when we pass from isogenies to linear maps, the functor T_ℓ is actually faithful. In fact we are even free to base change on the left hand side whilst retaining injectivity:

Proposition 2.7. *For any pair of elliptic curves E_1, E_2/K the map*

$$\mathrm{Hom}(E_1, E_2) \otimes \mathbb{Z}_\ell \longrightarrow \mathrm{Hom}(T_\ell(E_1), T_\ell(E_2))$$

is injective.

In full generality, this is all one can say. However, for particular fields of arithmetic interest stronger results are available. In the case of finite fields this is our main theorem, which we are now ready to state. For a finite extension $K/\mathbb{F}_q$, denote by the subscript K those isogenies/linear maps which are defined over K or, equivalently, equivariant with respect to the absolute Galois group of K. Then we have the following

Theorem 2.8 (Tate Isogeny Theorem, [Tat66]). *For E_1, E_2 and K as above the map*

$$\mathrm{Hom}_K(E_1, E_2) \otimes \mathbb{Z}_\ell \longrightarrow \mathrm{Hom}_K(T_\ell(E_1), T_\ell(E_2))$$

is an isomorphism.

The rest of this work will be devoted to the proof of this statement. As an application, we will it use to prove a fact which deserves a standalone status both for its surprising content and for its connection to another profound and beautiful branch of number theory.

Theorem 2.9 (Slogan). *Two elliptic curves over $\mathbb{F}_q$ are isogenous if and only if they have the same number of $\mathbb{F}_q$-rational points.*

We urge the reader to pause and ponder the non-trivial nature of both implications of the theorem we have just stated. On one hand, isogenies are in general not one-to-one, so preserving the number of points on the base field is all but guaranteed; on the other, the comparison of a simple numerical invariant should produce an isogeny out of thin air. Theorem 2.9 is a testament to the fruitful idea that point counting in finite field geometry is a very powerful tool, much more than we could possibly hope to convey here.

2.3 Computing the trace of the Frobenius

Construction 2.10. Fix an integer m such that $(m, \mathrm{char}\, K) = 1$. Let P, Q be m-torsion points on E. The degree 0 divisor

$$m\left((Q) - (O)\right)$$

sums to O, so by Proposition 1.14 it is the divisor of a function f. We compute

$$\mathrm{div}(f \circ [m]) = m\left(\sum_{mS=Q}(S) - \sum_{mR=O}(R)\right) = m\left(\sum_{R\in E[m]}(S+R) - (R)\right)$$

for an arbitrary S such that $mS = Q$. Again using Proposition 1.14, the divisor inside the parentheses is principal, for $S \in E[m^2]$ and there are m^2 terms being summed. Thus there exists a function g such that

$$\mathrm{div}(g^m) = \mathrm{div}(f \circ [m]).$$

The ratio

$$\frac{g^m}{f \circ [m]}$$

is then constant, so up to changing our choice of either f or g we may suppose it to be 1. Now consider the function

$$\frac{g(X+P)}{g(X)}.$$

As X ranges over E we always find

$$\frac{g(X+P)^m}{g(X)^m} = \frac{f(mX+mP)}{f(mX)} = 1.$$

The image of this function is then contained in the finite set μ_m of m-th roots of unity; in particular, it must be constant. We set $e_m(P, Q)$ to be equal to its constant value, obtaining a well-defined map

$$e_m : E[m] \times E[m] \longrightarrow \mu_m$$

Definition 2.11. The map constructed above is known as the *Weil pairing* on $E[m]$

The name is justified by

Theorem 2.12. *The Weil pairing enjoys the following properties:*

1. *It is bilinear, alternating and nondegenerate.*
2. *It is Galois-equivariant.*
3. *It is compatible with multiplication, in the sense that*

$$e_m(nP, Q) = e_{mn}(P, Q)$$

for all $P \in E[mn]$, $Q \in E[m]$.

4. *For any isogeny* $\phi : E_1 \longrightarrow E_2$ *the adjoint of* ϕ *with respect to* e_m *is* $\widehat{\phi}$*, i.e.*

$$e_m(\phi(P), Q) = e_m(P, \widehat{\phi}(Q))$$

for all $P \in E_1[m]$, $Q \in E_2[m]$.

Compatibility with multiplication ensures that we may take limits as before and extend the pairing to an *ℓ-adic Weil pairing* on $T_\ell(E)$, denoted simply $e(\cdot, \cdot)$.
A word of caution – even though we write the operation on the range of the Weil pairing multiplicatively (because it is inherited from the μ_{ℓ^k}'s), it corresponds to addition under the isomorphism with $\mathbb{Z}_\ell$. We are now in the position to establish the first half of theorem Theorem 2.9. First, we connect the degree of an endomorphism and the determinant of the associated linear mapping:

Proposition 2.13. *Let* $\phi \in \mathrm{End}(E)$. *Then* $\deg \phi = \det \phi_\ell$

Proof. After fixing a basis (v_1, v_2) of $T_\ell(E)$ it is enough to put together all we know about dual isogenies and the Weil pairing to compute

$$\begin{aligned} e(v_1, v_2)^{\deg \phi} = \\ = e([\deg \phi]v_1, v_2) = e(\widehat{\phi}_\ell(\phi_\ell(v_1)), v_2) = e(\phi_\ell(v_1), \phi_\ell(v_2)) = \\ = e(v_1, v_2)^{\det \phi_\ell} \end{aligned}$$

Because the pairing is nondegenerate and takes values in $\mathbb{Z}_\ell$ the exponents must be equal. □

The key idea is to now apply this formula to a well-chosen endomorphism.

Proof of the 'only-if' statement in Theorem 2.9. Suppose $\phi : E_1 \longrightarrow E_2$ is an isogeny between elliptic curves defined over $\mathbb{F}_q$ with Frobenius endomorphisms $\mathfrak{f}_1$, $\mathfrak{f}_2$ respectively.
The diagram

$$\begin{array}{ccc} E_1 & \xrightarrow{\phi} & E_2 \\ \downarrow{\scriptstyle \mathfrak{f}_1} & & {\scriptstyle \mathfrak{f}_2}\downarrow \\ E_1 & \xrightarrow{\phi} & E_2 \end{array}$$

commutes – ϕ, being an algebraic map, is locally defined by polynomial functions which are compatible with the Frobenius action. Furthermore, ϕ_ℓ admits an inverse up to a scalar factor, namely $\widehat{\phi}_\ell$. Thus, possibly after extending coefficients to $\mathbb{Q}_\ell$, ϕ_ℓ is a linear isomorphism between $T_\ell(E_1)$ and $T_\ell(E_2)$ conjugating the Frobenius on the first curve to that on the second one. In particular, $(\mathfrak{f}_1)_\ell$ and $(\mathfrak{f}_2)_\ell$ have the same trace and determinant. The latter is easily computed:

$$\det(\mathfrak{f}_i)_\ell = \deg \mathfrak{f}_i = q$$

The trace of the Frobenius may also be recovered from a determinant by evaluating the characteristic polynomial at 1:

$$\det(1-\mathfrak{f}_i)_\ell = 1 - \operatorname{tr}(\mathfrak{f}_i)_\ell + \det(\mathfrak{f}_i)_\ell = q + 1 - \operatorname{tr}(\mathfrak{f}_i)_\ell$$

On the other hand, $1-\mathfrak{f}_i$ is separable by Theorem 1.21: fixing an invariant form ω and recalling that the Frobenius is purely inseparable we obtain $(1-\mathfrak{f}_i)^*\omega = 1^*\omega = \omega$. Hence,

$$\det(1-\mathfrak{f}_i)_\ell = \deg(1-\mathfrak{f}_i) = \#\ker(1-\mathfrak{f}_i).$$

But the points killed by $1-\mathfrak{f}_i$ are precisely the points fixed by the Frobenius, i.e. the $\mathbb{F}_q$-rational points of E_i. Therefore we have just shown

$$\#E_1(\mathbb{F}_q) = q + 1 - \operatorname{tr}(\mathfrak{f}_1)_\ell = q + 1 - \operatorname{tr}(\mathfrak{f}_2)_\ell = \#E_2(\mathbb{F}_q)$$

which was the claim. □

2.4 Reduction

We now begin setting the stage for the proof of the isogeny theorem as well as the 'if' statement in the corollary. In classical number theory one commonly looks at diophantine equations modulo a certain prime to establish nontrivial properties like the existence or type of its solutions. This idea carries over to the realm of geometry: if a variety is cut out from projective space by polynomials with coefficients in $\mathbb{Z}$ we may consider the same equations in $\mathbb{F}_p$ to get a corresponding variety in characteristic p. More generally, the coefficients may be taken in the ring of integers of a local or global field. Since for any prime ideal $\mathfrak{p}$ it is always possible to multiply an equation with coefficients in the fraction field by a suitable element in order to make it $\mathfrak{p}$-integral, this condition is not restrictive. Moreover, since we will be working one prime at a time, we shall follow Silverman and assume our field to be local (hence complete). We now define what it means to reduce a point $P \in X(K)$ for a projective variety X over the local field K with residue field $F_\mathfrak{p}$. Suppose that $X \subset \mathbb{P}^n$ and P has homogenous coordinates $[x_0, \ldots, x_n]$. Letting $k = \min v(x_i)$ and choosing a uniformiser π for K we see that $[\pi^k x_0, \ldots, \pi^k x_n]$ is another representative of P with the property that all the entries are in the ring of integers $\mathcal{O}_K$ but not all of them lie in the maximal ideal $\mathfrak{p} = (\pi)$. Thus it is possible to consider the point $\widetilde{P} \in \mathbb{P}^n(F_\mathfrak{p})$ with coordinates obtained by simply considering the previous ones modulo $\mathfrak{p}$. Clearly, $\widetilde{P}$ depends only on P, since a different choice of uniformiser differs by a unit and projective changes of coordinates are compatible with reduction. Furthermore, $\widetilde{P}$ continues to satisfy the same equations P did, and therefore belongs to $\widetilde{X}(F_\mathfrak{p})$.
Although the map $\mathbb{P}^n(K) \longrightarrow \mathbb{P}^n(F_\mathfrak{p})$ is surjective, the same is not true for an arbitrary variety even if we allow field extensions. For example, consider the zero set of $x^p + p$ in $\mathbb{P}^1(\mathbb{Q}_p)$. The corresponding locus in $\mathbb{P}^1(\mathbb{F}_p)$ is the single point 0 (with multiplicity p); however, the equation $x^p + p \equiv 0 \mod p^2$ has no solutions with $x \equiv 0 \mod p$. The issue is that the characteristic-p point is nonsmooth. For smooth points it is possible to achieve a 'lift' using Hensel's celebrated lemma.

Lemma 2.14 (Hensel)**.** *Let K be a complete non-Archimedean field with ring of integers R and residue field F. For any polynomial $f \in R[x]$ and any* simple *root $\widetilde{\alpha} \in F$ of the corresponding reduced polynomial $\widetilde{f}$ there exists a unique root $\alpha \in K$ of f that reduces to $\widetilde{\alpha}$.*

Proof. The proof is a standard approximation argument; see, for example, [Neu99]. □

Remark 2.15. The assumption of simpleness on the root should be seen geometrically as the assertion that α is a smooth point of the 0-dimensional scheme given by the zero locus of $\widetilde{f}$.

Using 2.14 it is possible to show surjectivity of reduction for smooth points on plane projective curves (or, more generally, hypersurfaces) which is the case that interests us.

Proposition 2.16. *Let $C \subset \mathbb{P}^2$ be a plane curve over K and $\widetilde{P}$ a smooth point of $\widetilde{C}$. Then there exists a point P of C that reduces to $\widetilde{P}$.*

Proof. Let C be given by a homogenous polynomial $F \in R[X, Y, Z]$ whose coefficients are not all 0 modulo $\mathfrak{p}$. The condition that $\widetilde{P}$ be smooth translates to the nonvanishing of at least one of the partial derivatives of $\widetilde{F}$ at $\widetilde{P}$; without loss of generality, suppose that $\frac{\partial \widetilde{F}}{\partial X}(\widetilde{P}) \neq 0$. Choosing arbitrary lifts for the Y and Z coordinates of $\widetilde{P}$ and substituting them in F, we obtain a polynomial in one variable having a simple root modulo $\mathfrak{p}$ at the X coordinate of $\widetilde{P}$. Finally, we apply Hensel's lemma to it obtaining the X-coordinate of a point P satisfying the claim (this is indeed a well-defined projective point: not all its coordinates are zero since they are already valid in the residue field) □

We shall also refer to the above statement as 'Hensel's Lemma', being a generalisation of it.
We have described what it means to reduce a projective variety modulo a prime. There is an entirely analogous description of reduction for morphisms. To avoid descending too much into technicalities, the reader is encouraged to trust that morphisms, being algebraic maps, are compatible enough with reduction that the procedure cannot go 'too awry'. In particular, reducing an elliptic curve yields a curve that is not necessarily smooth, but its smooth points nevertheless carry an algebraic group structure defined by the very same construction as for the original curve. In any case, we shall always be considering curves which are also smooth over the residue field. When this holds we say that the curve has *good reduction*, and that it has *bad reduction* otherwise. We shall not explore the rich theory of reduction for elliptic curves further (see Silverman for further details); however, we state and prove one last result for later use.
Because the multiplication map on E is algebraic, the condition $nP = 0$ for a point $P \neq O$ may be phrased by imposing that the coordinates of P be roots of a certain polynomial. Concretely, for a curve in Weierstraß form, these are known as division polynomials and are denoted by ψ_n. It is possible to show that the degree of ψ_n is $\frac{n^2-1}{2}$ for odd n and $\frac{n^2}{2} - 1$ for even n.

Proposition 2.17. *Let E/K be an elliptic curve with good reduction at $\mathfrak{p}$. Then for any integer n relatively prime to* $\operatorname{char} F$ *the reduction map restricted to $E[n]$ is injective.*

Proof. First, note that since n is nonzero in F it is a fortiori nonzero in K. Therefore by Proposition 1.28 both $E[n]$ and $\widetilde{E}[n]$ have cardinality equal to n^2. Moreover, the above discussion implies that the image of $E[n]$ under the reduction map lands in $\widetilde{E}[n]$. Therefore showing injectivity is equivalent to showing surjectivity. This is accomplished by means of Hensel's lemma applied to ψ_n. We shall work with the 'short' Weierstraß form for the sake of ease, but the argument works in general with a little more care. For odd n the nonzero n-torsion points fall into pairs with matching x-coordinate and opposite y-coordinate. Since $\deg \psi_n = \frac{n^2-1}{2}$ all the roots must be simple, and they can be lifted using Lemma 2.14. We then invoke Hensel again on each of these to find the corresponding values for y, using the smoothness of E. The case of n even is identical, but one must first take care to exclude the 2-torsion points. These are easily accounted for, since they are precisely the points making y vanish. □

Remark 2.18. A more conceptual proof of Proposition 2.17 is given in Silverman. One first establishes the exact sequence

$$0 \longrightarrow E_0(K) \longrightarrow E(K) \longrightarrow \widetilde{E}_{ns}(F) \longrightarrow 0$$

and then shows the first term to be isomorphic to a certain formal group over the maximal ideal $\mathfrak{p}$. The theory of formal groups then informs us that such objects do not have torsion of order prime to the characteristic of the base field, proving the claim.

CHAPTER 3

Complex Elliptic Curves and Lifting

To finish the proof of the main results, we require an input from the analytic point of view of the theory of elliptic curves: using a theorem of Deuring ([Deu41]), we pass from positive to zero characteristic where establishing the existence of the required maps is much more tractable. In turn, the proof of said theorem hinges on a certain integrality result which we prove by means of the (basic) theory of modular forms.

3.1 Uniformisation of elliptic curves

Let us first describe the general situation for complex elliptic curves. Any smooth projective curve defined over $\mathbb{C}$ may be viewed as a complex analytic manifold of dimension one, that is, as a Riemann surface. This is in fact the historical origin of the algebraic theory we have been using so far. Another source of Riemann surfaces of (topological) genus one is supplied by one-dimensional complex tori, i.e. quotients $\mathbb{C}/\Lambda$ with Λ a lattice (full-rank discrete subgroup) inside $\mathbb{C}$. It is well-known that any such lattice is homothetic to one of the form $\mathbb{Z} + \tau\mathbb{Z}$ for some complex τ in the upper-half plane $\mathbb{H}$. Moreover, whenever two such lattices generated by τ and τ' are homothetic, the two values are related by a Möbius transformation of the form $\tau' = \frac{a\tau+b}{c\tau+d}$ with the corresponding matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ belonging to the modular group $\Gamma = \mathrm{SL}_2(\mathbb{Z})$. We shall return to this group in the next section.
We are thus led to investigate three seemingly different kinds of objects: elliptic curves, complex tori and complex lattices. It turns out that these three kinds of objects are essentially one and the same. More precisely, the following holds:

Theorem 3.1 (Uniformisation Theorem)**.** *The following categories are equivalent:*

- *Complex elliptic curves with isogenies as morphisms*
- *One dimensional complex tori $\mathbb{C}/\Lambda$ (for some lattice $\Lambda \subset \mathbb{C}$) with pointed holomorphic maps as morphisms*
- *Homothety classes of lattices $\Lambda \subset \mathbb{C}$, where a morphism $\Lambda_1 \longrightarrow \Lambda_2$ is given by a matrix $A \in M_2(\mathbb{Z})$ with nonzero determinant such that $A(\Lambda_2) \sim \Lambda_1$*

Under this equivalence, the degree of an isogeny corresponds to the absolute value of the determinant of the associated matrix.

Moreover, there exists a holomorphic function $j : \mathbb{H} \longrightarrow \mathbb{C}$ such that for an elliptic curve E and any lattice $\mathbb{Z} + \tau\mathbb{Z}$ in the corresponding homothety class one has

$$j(E) = j(\tau).$$

For a proof we refer the reader to [Sil94]. The function j appearing in the statement above is but the first link between elliptic curves and modular forms, a very small part of which we will recall now.

3.2 Modular functions and the modular polynomial

In the previous section we were led to consider the modular group Γ and its action on the upper-half plane given by Möbius transformations. It is then natural to define the following class of functions:

Definition 3.2. A *modular function* for Γ is a meromorphic Γ-invariant function which is meromorphic at infinity, namely, such that the behaviour of $f(z)$ as $\mathcal{I}(z) \longrightarrow \infty$ is that of a singularity of finite order.

The analytic function j, given by Theorem 3.1, is an example of a holomorphic modular function. As it turns out, it is essentially the only one.

Proposition 3.3. *The set of holomorphic modular functions is precisely $\mathbb{C}[j]$, the polynomial functions in j.*

This is but the first of many extraordinary properties of this function, which we cannot discuss here. We shall however need another fact about j for the upcoming proof. Being Γ-invariant, it is in particular 1-periodic; it therefore admits a Fourier expansion in terms of $q = e^{2\pi iz}$. Using the more general theory of modular forms (see, for example, [Sil94] or the classic [Shi94]) one carries out a standard computation, leading to

Proposition 3.4. *The Fourier expansion of j is of the form*

$$j(z) = \frac{1}{q} + \sum_{k=0}^{\infty} c_k q^k$$

where the c_k are integers.

Theorem 3.5. *For any $n > 0$ there esists a nonzero polynomial $F_n \in \mathbb{Z}[x, y]$ such that for any isogeny $E_1 \longrightarrow E_2$ of degree n between elliptic curves over $\mathbb{C}$ we have $F_n(j(E_1), j(E_2)) = 0$.*

Proof. We reproduce the argument presented in [Sil94]; non-analytic proofs of the theorem, which generalise to different fields, also exist but require more advanced tools.
Recall that, under the dictionary of Theorem 3.1, elliptic curves and isogenies correspond, respectively, to homothety classes of lattices and orbits of 2-by-2 integral nonsingular matrices under the action of the modular group. Furthermore, it is not hard to show that

$$S_n = \left\{ \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} : ad = n,\ 0 \leq b < d \right\}$$

is a set of representatives for the right action of Γ on the matrices having determinant n. Indeed, for any such matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ with $c \neq 0$ we write $-\frac{a}{c} = \frac{s}{r}$ in lowest terms and find

integers p, q such that $ps - qr = 1$ using Euclidean division. Then

$$\begin{pmatrix} p & q \\ r & s \end{pmatrix} \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} * & * \\ 0 & * \end{pmatrix},$$

so we may suppose that $c = 0$. By acting further with a matrix of the form

$$\begin{pmatrix} 1 & k \\ 0 & 1 \end{pmatrix} \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} = \begin{pmatrix} a & b + kd \\ 0 & d \end{pmatrix}$$

we can bring b in the range $[0, d-1]$. If two elements of S_n are equivalent we have

$$\begin{pmatrix} a & b \\ 0 & d \end{pmatrix} = \begin{pmatrix} p & q \\ r & s \end{pmatrix} \begin{pmatrix} a' & b' \\ 0 & d' \end{pmatrix}.$$

Expanding the product one finds that $r = 0$, $a = pa'$, $d = sd'$, and equating $psa'd' = ad = n = a'd'$ we obtain $p = s = \pm 1$. We may choose the positive sign up to possibly multiplying by $-I$. Finally, the top-right entry yields $b \equiv b' \mod d$ and since they are both between 0 and $d - 1$ we conclude that $b = b'$.

The conclusion of the theorem is thus established if we can show that the expression

$$F_n = \prod_{\alpha \in S_n} (X - j \circ \alpha)$$

is actually a polynomial in j and X with integer coefficients. This is accomplished in a number of steps.

First, expanding $F_n = \sum a_m X^m$, we wish to show that each of the a_m is a Γ-invariant holomorphic function on the upper-half plane. Let $\gamma \in \Gamma$. For each $\alpha \in S_n$ we have $\det(\alpha\gamma) = n$, and since S_n is a right coset there exists a unique $\delta_\alpha \in \Gamma$ such that $\delta_\alpha \alpha\gamma$ is again an element of S_n. Furthermore, as α ranges over the whole S_n, so does $\delta_\alpha \alpha\gamma$. Therefore

$$F_n \circ \gamma = \prod_{\alpha \in S_n} (X - j \circ \alpha\gamma) = \prod_{\alpha \in S_n} (X - j \circ \delta_\alpha^{-1} \delta_\alpha \alpha\gamma) =$$
$$= \prod_{\alpha \in S_n} (X - j \circ \delta_\alpha \alpha\gamma) = \prod_{\alpha \in S_n} (X - j \circ \alpha) = F_n.$$

In particular, the coefficients a_m are Γ-invariant. We now analyse their behaviour at infinity. Being 1-periodic holomorphic functions, the a_m also admit a Fourier expansion. Using Proposition 3.4 we find for $\alpha = \begin{pmatrix} a & b \\ 0 & d \end{pmatrix}$

$$j \circ \alpha = e^{-2\pi i \frac{az+b}{d}} + \sum_{k \geq 0} c_k e^{2k\pi i \frac{az+b}{d}},$$

so for a suitably large integer N we have $j \circ \alpha = o(q^{-N})$ as $q \longrightarrow 0$. Since each s_m is a symmetric function in the $j \circ \alpha$'s, the same is true for it. We have just shown that, for all m, s_m is a holomorphic Γ-invariant function which is meromorphic at infinity, i.e. a holomorphic modular function. These are completely characterised by Proposition 3.3 – they are precisely the polynomial functions $\mathbb{C}[j]$.

Next, let us look more closely at the Fourier coefficients of s_m. If we denote by ζ the distinguished n-th root of unity $e^{\frac{2\pi i}{n}}$ and observe that

$$e^{2\pi i \frac{az+b}{d}} = \zeta^{ab} q^{\frac{a^2}{n}}$$

the expansion above for $j\circ\alpha$ may be rewritten as an expansion in powers of $q^{\frac{1}{n}}$ with coefficients in $\mathbb{Z}[\zeta]$. Thus, the Fourier coefficients of each s_m also lie in $\mathbb{Z}[\zeta]$. We claim that in fact these coefficients lie in $\mathbb{Z}$. To see this, let G be the group $\mathrm{Gal}(\mathbb{Q}(\zeta)/\mathbb{Q}) \simeq (\mathbb{Z}/n\mathbb{Z})^\times$. Under this isomorphism the action of G is specified by $\sigma(\zeta) = \zeta^{r(\sigma)}$, with $r(\sigma)$ relatively prime to n. If we apply σ coefficient-wise to the expansion of $j \circ \alpha$ and compare with the original we obtain

$$\sigma\left(j \circ \begin{pmatrix} a & b \\ 0 & d \end{pmatrix}\right) = j \circ \begin{pmatrix} a & r(\sigma)b \\ 0 & d \end{pmatrix}.$$

Note that, as b runs through a set of representatives for the residue classes modulo d, so does $r(\sigma)b$. Moreover, we are free to change the top-right entry of α by a multiple of d without affecting $j \circ \alpha$ by multiplying on the left by a suitable element of Γ as we did above. It follows that the action of G permutes the various $j \circ \alpha$, and since each s_m is a symmetric function of these it is left fixed by G. Put differently, the Fourier coefficients of s_m lie in $\mathbb{Z}[\zeta]^G = \mathbb{Z}$. All that is left now is to show that $\mathbb{C}[j] \cap \mathbb{Z}((q)) = \mathbb{Z}[j]$. Let $f = \sum_{k=0}^{d} a_k j^k$ be a polynomial in j admitting an integral Fourier expansion. We substitute the Fourier expansion of j and look at the lowest order term, namely $\frac{a_d}{q^d}$. Using the uniqueness of Fourier coefficients, this yields $a_d \in \mathbb{Z}$. Finally, we use induction applied to $f - a_d j^d$ to conclude that all the a_d are integers. □

The above polynomial is sometimes referred to as the *modular polynomial*, though this term is reserved by some authors to denote the factor of F_n corresponding to the isogenies with cyclic kernel.

3.3 The lifting theorem – Completing the proof

We now turn to the main step towards the proof of Theorem 2.8.

Theorem 3.6 (Deuring Lifting Theorem)**.** *Let E be an elliptic curve over $\overline{\mathbb{F}}_p$ and ϕ an endomorphism of E. Then there exist a number field K, a prime $\mathfrak{p}$ of K lying above p, an elliptic curve $\mathcal{E}$ and an endomorphism φ of $\mathcal{E}$ both defined over K such that $\mathcal{E}$ and φ reduce respectively to E and ϕ modulo $\mathfrak{p}$.*

Proof. The following argument is taken from [Lan87], with minor corrections.[1]
We may suppose without loss of generality that ϕ be separable. If not, then $\phi^*\omega = 0$, where ω is the invariant differential on E as usual. But then $\phi + 1$ is separable by Theorem 1.21, and lifting the former is equivalent to lifting the latter. Let us choose a transcendental $t \in \mathbb{C}$ and take $\mathcal{E}_t$ to be an elliptic curve over $\mathbb{C}$ whose j-invariant is t. Composing reduction modulo p with the evaluation map $t \mapsto j(E)$ yields a well-defined map $\rho : \mathbb{Z}[t] \to \overline{\mathbb{F}}_p$, which we view as reduction modulo some prime of $\mathbb{Z}[t]$. Denote by n the degree of ϕ. We claim that $\ker\phi$ is the image under ρ of some subgroup $H \subseteq \mathcal{E}_t$ of order n. Write $n = p^k s$ with p not dividing s. By Proposition 2.16 and Proposition 2.17 we know reduction is injective on s-torsion and overall surjective, so there exists a subgroup of order s which injects into $\ker\phi$. As for the p-component of $\ker\phi$, there are two cases. If E is supersingular we are done, for there is no p-primary torsion at all. If E is ordinary the map $[p^k]$ on E has separable degree p^k. Equivalently, the p^k-th division polynomial is the p^k-th power of some separable polynomial ψ. Arguing as in the proof of Proposition 2.17, Hensel's lemma may then be applied twice, first to the roots of this polynomial, which are precisely the x-coordinates of the points in $E[p^k]$, and then to the curve itself to obtain the corresponding y-coordinates,

[1] The proof presented in the book contains a small error; namely, the procedure described there to make $\deg\phi$ relatively prime with p breaks down when $p = 2$.

giving the desired subgroup inside $\mathcal{E}_t$. Now consider the ring $\mathbb{Z}[t, t_1, \ldots, t_n]$ where the t_i 's are the j-invariants of elliptic curves admitting an isogeny from $\mathcal{E}_t$ of degree n. Note that by construction $\mathcal{E}_t/H$, which exists by Theorem 1.22, is such a curve. Thus, setting $\mathcal{E}_s = \mathcal{E}_t/H$, we may suppose that $j(\mathcal{E}_s) = s$ is one of the t_i's. By Theorem 3.5 the aforementioned ring is an integral extension of $\mathbb{Z}[t]$. We will denote by R its integral closure in some finite extension K of $\mathbb{Q}(t)$ to be specified later. The integrality of R over $\mathbb{Z}[t]$ allows us to extend ρ to a homomorphism $R \to \overline{\mathbb{F}}_p$ which we will, with a slight abuse of notation, still call ρ. We may assume, possibly after selecting different models, that $\mathcal{E}_t$ reduces to E since they have the same j-invariant as curves over $\overline{\mathbb{F}}_p$. Thus

$$\widetilde{\mathcal{E}}_s \simeq \widetilde{\mathcal{E}_t/H} \simeq \widetilde{\mathcal{E}}_t \simeq E$$

or, put differently,

$$(p, t - s) \subset \ker \rho.$$

Let $\mathfrak{q}$ be a minimal prime of R above $t - s$. By Krull's Hauptidealsatz the height of $\mathfrak{q}$ is 1. In particular $\mathfrak{q} \cap \mathbb{Z} = \emptyset$, otherwise there would be a chain $(0) \subset (q) \subset \mathfrak{q}$ for some prime number q. It follows immediately that $R/\mathfrak{q}$ is a number field and the curves $\mathcal{E}_t, \mathcal{E}_s$, when reduced modulo $\mathfrak{q}$, become isomorphic after possibly passing to some finite extension of said field. Quotienting out by H therefore descends to an endomorphism

$$\varphi : \mathcal{E} \longrightarrow \mathcal{E}$$

whose kernel further reduces to $\ker \phi$ modulo (the class of) $\ker \rho$. This almost proves the theorem, for $\widetilde{\varphi}$ and ϕ can only differ by an automorphism of E. If $\mathrm{Aut}(E) = \{[\pm 1]\}$ we are done. The only curves which are left out are those of invariant 0 or 1728, where the existence of lifts must be checked by way of direct computation. □

Equipped with theorem Theorem 3.6, we are finally ready to complete the proof begun in the previous chapter.

Proposition 3.7. *If E_1 and E_2 have the same number of $\mathbb{F}_q$-rational points then they are $\overline{\mathbb{F}}_q$-isogenous.*

Proof. Suppose E_1, E_2 are elliptic curves defined over $\mathbb{F}_q$ having the same number of $\mathbb{F}_q$-points. By the results of Chapter 2 this implies that the Frobenius endomorphisms $\mathfrak{f}_1$, $\mathfrak{f}_2$ have the same characteristic polynomial. In particular, either both curves are ordinary or they are supersingular. In the ordinary case we apply the lifting theorem to the pairs $(E_i, \mathfrak{f}_i)$ to obtain elliptic curves $\mathcal{E}_i$ with endomorphisms f_i, all defined over two number fields which we may assume to coincide after taking composites. Since we have supposed the original curves to be ordinary, the $\mathfrak{f}_i$ and therefore also the f_i are not integers. Thus the $\mathrm{End}(\mathcal{E}_i)$ are both orders in the imaginary quadratic field $K = \mathbb{Q}(f_1) = \mathbb{Q}(f_2)$. We now pass to the complex point of view. Under the correspondence of theorem Theorem 3.1 the curves E_i are represented by the homothety classes of certain lattices $\Lambda_i \subset \mathbb{C}$. But we have just shown that these lattices admit complex multiplication by some order in K, so up to a different choice of representative the Λ_i are themselves spanned by elements of K. This implies that there exists a matrix with integer coefficients and nonzero determinant mapping, say, Λ_1 onto Λ_2, which corresponds to a nonconstant isogeny $\mathcal{E}_1 \longrightarrow \mathcal{E}_2$. Reducing this isogeny modulo the prime above p shows that E_1 and E_2 are isogenous. The above argument must be modified in the case when the $\mathfrak{f}_i$ are integers. This can only happen if the curves are supersingular (though the converse is false, e.g. when q is an odd power of p). To find suitable elements to lift, we resort to the classification of quaternion algebras over number fields ([Ser73]). Recall that such an algebra is determined, up to isomorphism, by the set of places at which it is ramified (i.e.

nonsplit), which is known to be finite and of even cardinality. We already know, by virtue of Proposition 2.7, that for any supersingular elliptic curve E in characteristic p the quaternion algebra $\mathrm{End}(E)\otimes\mathbb{Q}$ splits at all primes $\ell \neq p$ .If it were also split at p or infinity, then it would be forced to be split everywhere and so be isomorphic to the matrix algebra $M_2(\mathbb{Q})$. But this is impossible, for the latter has zero-divisors which $\mathrm{End}(E)$ lacks. We infer that $\mathrm{End}(E)$ is an order in the unique quaternion algebra over $\mathbb{Q}$ which ramifies precisely at p and infinity. Now, since $\mathrm{End}(E_1)$ and $\mathrm{End}(E_2)$ are orders in the same quaternion algebra, it is possible to find a pair of nontrivial elements which are scalar multiples of one another and run through the previous argument. □

This is not quite Tate's result - the isogeny we have constructed is only guaranteed to be defined over some finite extension of the base field. The final step in the proof is to bootstrap the previous result to yield the full strength of Theorem 2.8

Proof of surjectivity in Theorem 2.8. We shall prove a slightly weaker statement, namely, that the isomorphism holds upon tensoring with $\mathbb{Q}_\ell$. This will suffice to establish Theorem 2.9. Because we are now dealing with Galois modules which are vector spaces over a field, it is enough to prove the statement for a finite extension of the base field and invoke Galois descent to transport the isomorphism back down. To see this, recall ([GS17]) that for a finite Galois extension L/K and an L-vector space V the main theorem of Galois descent for vector spaces establishes an equivalence of categories between the K-forms of V and the semilinear actions of the Galois group G on L. Therefore, if we prove the theorem for L we may apply the functor $\cdot^G$ and still obtain isomorphic vector spaces by the descent theorem. If $\mathrm{Hom}_K(T_\ell(E_1), T_\ell(E_2)) = 0$ for all finite extensions $K/\mathbb{F}_q$ there is nothing to prove (and in particular the two curves are not isogenous). Else, there is a nonzero map f between the Tate modules, defined over some K with $[K : \mathbb{F}_q] = r$. Observe that, since the absolute Galois group Γ_K of K is topologically generated by the q^r-power Frobenius, a linear map is Γ_K-equivariant if and only if it commutes with the Frobenius. We now show that the traces of $(\mathfrak{f}_1)_\ell$ and $(\mathfrak{f}_2)_\ell$ are the same. If f is nonsingular this follows immediately from the definition of equivariance. Otherwise, its kernel and image are both one dimensional Galois-invariant subspaces. In other words there exist nonzero $v_1 \in \ker f$, $w \in \mathrm{im}\, f$ such that $(\mathfrak{f}_1)_\ell(v_1) = \lambda v$ and $(\mathfrak{f}_2)_\ell(w) = \mu w$ for some nonzero $\lambda, \mu \in \mathbb{Z}_\ell$. Adjoin to v_1 a linearly independent vector v_2, scaled so that $f(v_2) = w$. The matrix of $(\mathfrak{f}_1)_\ell$ with respect to this basis takes the form $\begin{pmatrix} \lambda_1 & a \\ 0 & \lambda_2 \end{pmatrix}$. Computing $\mu w = (\mathfrak{f}_2)_\ell(f(v_2)) = f((\mathfrak{f}_1)_\ell(v_2)) = af(v_1) + \lambda_2 f(v_2) = \lambda_2 w$ we deduce that $(\mathfrak{f}_1)_\ell$ and $(\mathfrak{f}_2)_\ell$ share an eigenvalue. Since they have the same determinant q^r, the other eigenvalues and hence the traces also coincide.We may then apply Proposition 3.7 to find a nonzero isogeny $\phi : E_1 \longrightarrow E_2$ defined over some finite extension of K which we will, without loss of generality, suppose to be K again. By the same argument we have already used, the endomorphism rings of the two curves are isomorphic when tensored with $\mathbb{Q}$. If they are both supersingular we are immediately done, for $\dim \mathrm{Hom}_K(T_\ell(E_1), T_\ell(E_2)) \leq \dim \mathrm{Hom}(T_\ell(E_1), T_\ell(E_2)) = 4$. If both curves are ordinary the linear Frobenius is semisimple and using Jordan theory it is easy to see that its commutator in $M_2(\mathbb{Q}_\ell)$ is two-dimensional. On the other hand we can produce two linearly independent isogenies in the left hand side: ϕ and $\phi \circ \mathfrak{f}_2$ cannot be scalar multiples of one another. Indeed, if they were $\mathbb{Q}_\ell$-linearly dependent, then because of the injection in Proposition 2.7 they would have to be already $\mathbb{Q}$- and hence $\mathbb{Z}$-linearly dependent. The identity $\phi \circ (m - n\mathfrak{f}_2) = 0$ with m, $n \in \mathbb{Z}$ implies, by comparing degrees, that $n\mathfrak{f}_2 = m$, hence the Frobenius is rational. On the other hand its minimal polynomial has integral coefficients, therefore we would in fact have $\mathfrak{f}_2 \in \mathbb{Z}$ which is impossible since E_2 is ordinary. □

Bibliography

[Deu41] Max Deuring. Die Typen der Multiplikatorenringe elliptischer Funktionenkörper. *Abh. Math. Sem. Hansischen Univ.*, 14:197–272, 1941.

[GS17] Philippe Gille and Tamás Szamuely. *Central simple algebras and Galois cohomology*, volume 165 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2017.

[KZ98] M. Kaneko and D. Zagier. Supersingular j-invariants, hypergeometric series, and Atkin's orthogonal polynomials. In *Computational perspectives on number theory (Chicago, IL, 1995)*, volume 7 of *AMS/IP Stud. Adv. Math.*, pages 97–126. Amer. Math. Soc., Providence, RI, 1998.

[Lan87] Serge Lang. *Elliptic functions*, volume 112 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1987. With an appendix by J. Tate.

[Mum99] David Mumford. *The red book of varieties and schemes*, volume 1358 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, expanded edition, 1999. Includes the Michigan lectures (1974) on curves and their Jacobians, With contributions by Enrico Arbarello.

[Mum08] David Mumford. *Abelian varieties*, volume 5 of *Tata Institute of Fundamental Research Studies in Mathematics*. Tata Institute of Fundamental Research, Bombay; by Hindustan Book Agency, New Delhi, 2008. With appendices by C. P. Ramanujam and Yuri Manin, Corrected reprint of the second (1974) edition.

[Neu99] Jürgen Neukirch. *Algebraic number theory*, volume 322 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999. Translated from the 1992 German original and with a note by Norbert Schappacher, With a foreword by G. Harder.

[Ser73] J.-P. Serre. *A course in arithmetic*, volume No. 7 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Heidelberg, 1973. Translated from the French.

[Shi94] Goro Shimura. *Introduction to the arithmetic theory of automorphic functions*, volume 11 of *Publications of the Mathematical Society of Japan*. Princeton University Press, Princeton, NJ, 1994. Reprint of the 1971 original, Kanô Memorial Lectures, 1.

[Sil94] Joseph H. Silverman. *Advanced topics in the arithmetic of elliptic curves*, volume 151 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994.

[Sil09] Joseph H. Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, Dordrecht, second edition, 2009.

[Tat66] John Tate. Endomorphisms of abelian varieties over finite fields. *Invent. Math.*, 2:134–144, 1966.

[Tat68] John Tate. Residues of differentials on curves. *Ann. Sci. École Norm. Sup. (4)*, 1:149–159, 1968.

Acknowledgments

If this thesis has seen the light of day, I have far more people to thank than the allotted space allows. However measly its mathematical content may be, it is not merely the exposition of the proof of some (albeit important) theorem. It embodies the three years I have spent at the University of Pisa, maturing as a person and as a mathematician. Between the lines lie hidden the struggles, joys, thrills and hardships that anyone who has undertaken university studies will be familiar with – new friendships, duties, opportunities and a dizzyingly vast pool of knowledge unfolding before one's eyes. Nevertheless, I wish to attempt to give back what little I may to those who have accompanied me on this journey.

The first and most heartfelt thanks cannot but go to my parents. Without the unconditional support, both material and emotional, it would have been impossible for me to fully devote myself to the pursuit of this degree. Let it be known that it is only because they nurtured and encouraged my passion from a young age that I can write these words at the present time.

A special word of thanks goes to my advisor, professor Lombardo, for opening my eyes to the power of algebraic and geometric techniques in number theory; for having very kindly invited me to join his reading group; for suggesting and mentoring this thesis; and for making me realise that "Hensel's Lemma is really just the Implicit Function Theorem". Naturally, I extend my gratitude to all the professors of this institution who have helped shape and expand my mathematical baggage.

To all the friends, old and new, that have enriched my stay in Pisa throughout these years – I could thank you a thousand times, and still it would not suffice. My attempt at summarising this experience in a handful of words will no doubt leave out many of you and do a poor job at conveying my gratitude to those I will mention. If this is indeed the case, please accept my apologies.

To Francesco, for being an overly patient listener to my numerous ramblings; for all your help, academic and otherwise; and for the good times shared together across eight years and counting.

To Giuseppe, Clementina, and the rest of the 'Canteen Debating Team', for giving me something to look forward to at every lunch- and dinnertime; for embracing and making the best out of our collective quirkiness; for the fits of laughter during our nights out and D&D sessions; and for supporting me when I needed it the most.

To Francesco "Fra' Minnocci" Minnocci, for welcoming me in the fold of the 'PHC'; for always playing along to my tongue-in-cheek dialogues; for our cultural dinners together; and for making me understand why truffle is like Sibelius.

To Lorenzo "Lorenzo Contorni", for being one of the first people I befriended in university; for our sometimes-serious-and-sometimes-not symposia; and for introducing me to the work of J. L. Borges.

To "Fenu", Davide Pierrat, Davide Ferri, Enrico, Tommaso, Oscar, and all the denizens of the Department of Mathematics and 'Aula Stud'[2], for the many fruitful conversations held in those halls during the day and the leisurely evenings spent therein after office hours.

[2] untranslatable

Finally, I must express my gratitude to a number of people who do not fit neatly in any category, but whose contribution has nevertheless been crucial.
To professor Tozzi, for teaching me perhaps the most important lesson I have ever learnt – that if one truly does love Mathematics, She shall never abandon him.
To professor Prina, for doing his best to keep the flame burning when times were dire.
To Matilda, whose greatest wish would have been to witness this day. This thesis is dedicated to her.