MINI COURSE: ELLIPTIC CURVES

ALEKSANDER HORAWA

These are notes for a summer mini course on Elliptic Curves at the Mathematics Department of the University of Michigan.

The newest version of the notes can be found on

https://sites.google.com/site/aleksanderhorawa/

If you find any typos or corrections, please let me know at ahorawa@umich.edu.

The main reference for the course is [Sil09] and we follow it closely throughout, but specific references are also provided where relevant.

Abstract. Elliptic curves are certain algebraic curves that arise naturally when studying Diophantine equations. Their surprising abelian group structure makes them prolific in number theory and cryptography, with applications to famous conjectures such as Fermat's Last Theorem and the ABC conjecture, and encryption standards employed by the biggest computer companies such as Google and Facebook.

This mini course will focus on studying elliptic curves over number fields. While the main goal will be the proof of the famous Mordell-Weil theorem, generally useful methods such as Galois cohomology, the theory of heights, and infinite descent will be introduced along the way. The last part of the mini-course will present one of the millennium prize problems about the rank of an elliptic curve and some of its invariants, the Birch–Swinnerton-Dyer conjecture.

1. Introduction to elliptic curves

One of the goals of number theory that goes back to Ancient Greeks is to understand Diophantine equations: polynomial equations over $\mathbb{Q}$. The degree 1 case is linear equations, resolved by linear algebra. The degree 2 case is conics, resolved by methods such as the *line trick*, which gives a rational parametrization of a conic. For degree 3 equations, the *line trick* does not give a parametrization, but this time it gives an *addition law* on the set of solutions. This makes the study of these curves extremely fruitful and interesting.

We first present some basic definitions. For a detailed introduction, see [ST92, Chap. 1] or [Sil09, Chap. III].

Definition 1.1. Let K be a field and $\mathrm{char}(K) \neq 2, 3$. An *elliptic curve* over K is the set of solutions $(x, y) \in \bar{K}^2$ of

$$E : y^2 = x^3 + ax + b, \qquad \text{for } a, b \in K \text{ such that } \Delta = 4a^3 + 27b^2 \neq 0,$$

together with a point O called the *point at infinity*.

Date: June 12–16, 2017.

We assume that $\Delta \neq 0$ to guarantee that the curve is smooth. This means that $x^3 + ax + b$ has no repeated roots, so we can define tangents at every point of the curve.

Remark 1.2. Equivalently, we could have defined an elliptic curve as a smooth projective curve of genus one together with any point $O \in E$. Then one can prove [Sil09, Prop. III.3.1.] that any such curve is of this form. The equation of an elliptic curve, however, is not unique.

Examples 1.3. The following curves are examples of elliptic curves over $\mathbb{R}$. Note that the graphs are smooth everywhere.

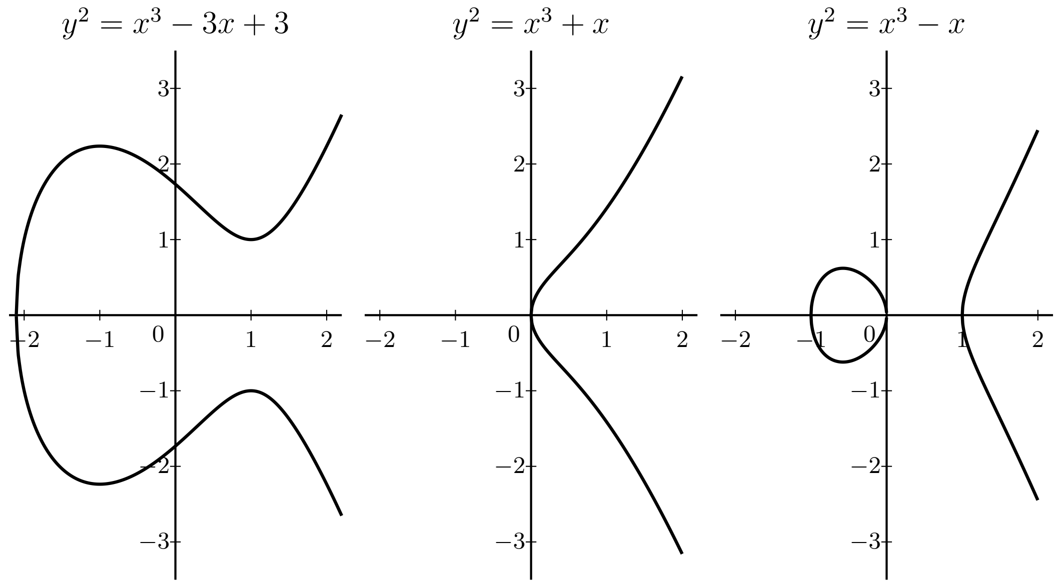
$y^2 = x^3 - 3x + 3$ $y^2 = x^3 + x$ $y^2 = x^3 - x$

However, the following curves are not elliptic curves. Clearly, for both of them $27b^2 + 4a^3 = 0$.

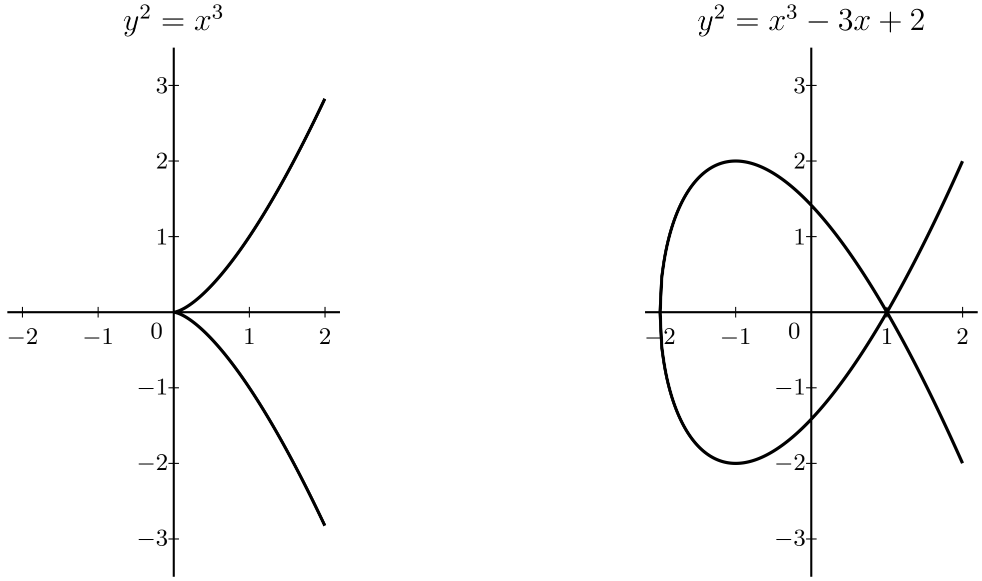
$y^2 = x^3$ $y^2 = x^3 - 3x + 2$

The first curve is singular at $(0, 0)$, but there is one tangent direction: we call it a *cusp*. The second one is singular at $(1, 0)$, but there are two distinct tangent directions: we call it a *node*.

However, these kind of examples will still be relevant in the study of elliptic curves. We will be interested in studying elliptic curves over $\mathbb{Q}$ (or some number field), and it will be useful to sometimes reduce them modulo some number. Note that the first example, $y^2 = x^3 - 3x + 3$, reduces to the first non-example, $y^2 = x^3$, modulo 3.

What is the point at infinity, O? This point does not belong to the plane but we think of it as the *direction upwards*. That is, if we wish to draw a line through O and any given point P on the plane, we would simply draw a vertical line through P. This is because elliptic curves are in fact projective curves on $\mathbb{P}^2(\bar{K}) = \mathbb{P}^2$ and choosing the embedding $(x, y) \mapsto [x : y : 1]$ of $\bar{K}^2$ into $\mathbb{P}^2$, O is the point $[0 : 1 : 0]$.

Why are elliptic curves so important and have so many applications? We can define a non-trivial *addition* on them.

Let E be an elliptic curve. Let us think of how a line can intersect with the cubic. Using Bézout's theorem (i.e. counting the intersection multiplicity) [Kir92, Th. 3.1], we can show that:

- any non-tangent line through two points on E intersects it at exactly one more point (this may be O);
- the tangent at O to E does not intersect E at any other point;
- any other line tangent to an elliptic curve intersects the curve at exactly one more point.

This allows us to naturally define the addition on E.

(1) The point at infinity O is defined to be the identity (i.e. $-O = O$ and $P + O = O + P = P$ for any point P).
(2) The negative $-P$ of $P = (x, y)$ is defined to be $(x, -y)$.
(3) If $P \neq Q$, then the line through P and Q intersects the curve at another point, say R. We then define $P + Q = -R$:

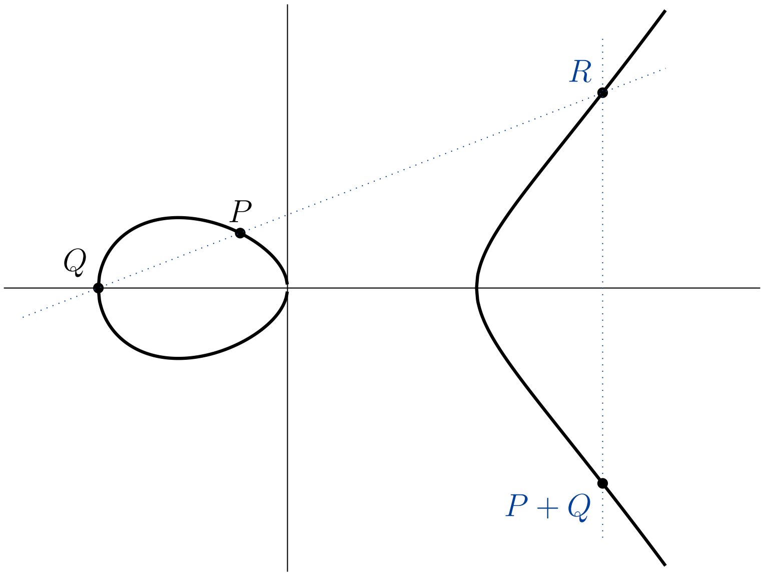

(4) If the line tangent to P intersects the curve at point R, then $2P = -R$:

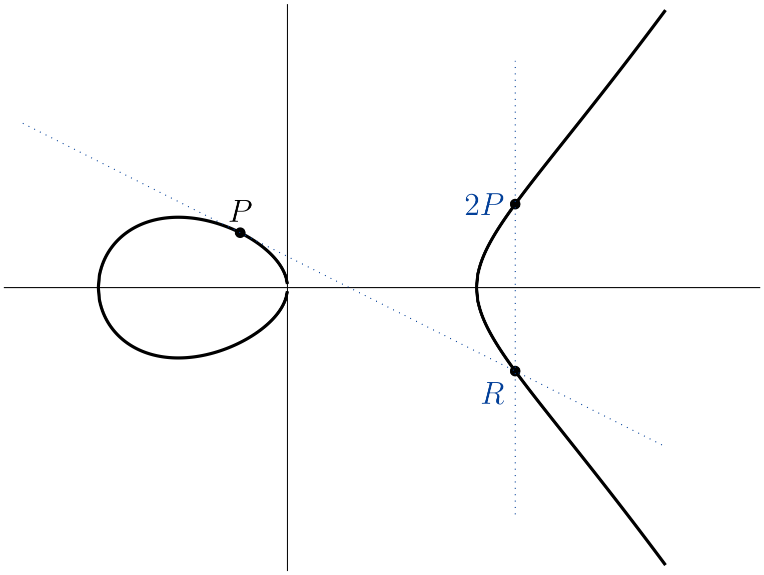

Why do we not define $P+Q$ equal to R, the third point of intersection? There are several reasons for this. To name one, we want O to be the identity of the group, i.e. $P+0=P$. Since the line through P and O is the line through P pointing upwards, it intersects the cubic at $R=-P$. Therefore, we need $P+O=-R=P$.

One can check that this makes E into an abelian group. The only group axiom which is not obvious from the definition is associativity, which can be shown using projective geometry or Abel's Theorem (see [Kir92, Ch. 3]).

Remark 1.4. The *geometric group law* described here agrees with the *algebraic group law* induced from the degree zero Picard group on E. See [Sil09, Prop. III.3.4].

The above definition is geometric in its nature, making in rather involved computationally. Fortunately, the addition law can be expressed by explicit formulas. Suppose we have $P=(x_1,y_1)$ and $Q=(x_2,y_2)$ on the curve, and we wish to find $P+Q=(x_3,y_3)$ and $2P=(x_4,y_4)$. By writing down the equation of the line passing through two points checking where it intersects the curve, one verifies that (see [Kob94, Ch. VI.1] for details):

$$\begin{aligned} x_3 &= \left(\tfrac{y_2-y_1}{x_2-x_1}\right)^2 - x_1 - x_2, & y_3 &= -y_1 + \left(\tfrac{y_2-y_1}{x_2-x_1}\right)(x_1-x_3), \\ x_4 &= \left(\tfrac{3x_1^2+a}{2y_1}\right)^2 - 2x_1, & y_4 &= -y_1 + \left(\tfrac{3x_1^2+a}{2y_1}\right)(x_1-x_4). \end{aligned} \tag{1}$$

While these formulas may seem complicated, they are very easy to implement in an algorithm. Also, one can now prove that E is an abelian group by verifying all the axioms algebraically.

1.1. **Outline of the course.**

The goal of the course is to study elliptic curves over $\mathbb{Q}$ and, more generally, over finite extensions of $\mathbb{Q}$, number fields.

Definition 1.5. An elliptic curves E is *defined over* K, written E/K, if there is an equation

$$E : y^2 = x^3 + ax + b, \qquad \text{with } \Delta = 4a^3 + 27b^2 \neq 0$$

such that $a, b \in K$ defining E. The set of *K-rational points* of E is then defined as

$$E(K) = \{(x, y) \in K^2 : y^2 = x^3 + ax + b\} \cup \{O\}.$$

The main goal of the course is to prove the following theorem.

Theorem 1.6 (Mordell–Weil)**.** *Let K be a number field. Then $E(K)$ is a finitely-generated abelian group, and hence is of the form*

$$E(K) \cong \mathbb{Z}^r \oplus \mathbb{Z}/n_1\mathbb{Z} \oplus \cdots \oplus \mathbb{Z}/n_k\mathbb{Z}.$$

We will write

$$E(K)_{\text{tors}} = \mathbb{Z}/n_1\mathbb{Z} \oplus \cdots \oplus \mathbb{Z}/n_k\mathbb{Z}$$

and call it the *torsion of* $E(K)$. We call r the *rank of* $E(K)$.

Outline of the proof:

(1) If an abelian group A is finitely-generated, then for any $m \geq 2$, A/mA is finite. Hence, along the way, we will have to show that

$$E(K)/mE(K) \text{ is finite,}$$

which is known as the *Weak* Mordell–Weil Theorem 3.1. We will start by proving this result.

  (a) We are interested in the cokernel of the multiplication by m map, which fits in the short exact sequence

$$0 \longrightarrow E[m] \longrightarrow E \xrightarrow{m} E \longrightarrow 0,$$

  where $E[m]$ is the kernel by definition; in other words $A[m] = \{a \in A : ma = 0\}$ for an abelian group A. We will hence introduce group cohomology in Section 2, and show that this short exact sequence gives a long exact sequence is group cohomology.

  (b) In Section 3, we extract from the long exact sequence the *Kummer sequence* and *Kummer pairing*, which reduces the finiteness statement to a result about finiteness of a certain extension of number fields. This can be proved using Kummer Theory and classical algebraic number theory.

(2) In Section 4, we develop a theory of heights on projective spaces, and apply it to elliptic curves. We prove the Descent Theorem 4.1: using the generators for $E(K)/mE(K)$, we can get any point on the elliptic curve below a certain height. But there are only finitely many points below this height, and hence the group is finitely-generated.

Methods such as cohomology of groups in Section 2, heights and descent in Section 4 are used in number theoretic problems and even other areas of math, and hence should be useful to the general audience.

The torsion of an ellptic curve, at least over $\mathbb{Q}$, is well-understood, due to the following theorem.

Theorem 1.7 (Mazur)**.** *The torsion of $E(\mathbb{Q})$ is one of the following 15 groups:*

$$\begin{aligned} &\mathbb{Z}/N\mathbb{Z} && \textit{for } 1 \leq N \leq 10,\ N = 12, \\ &\mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/2N\mathbb{Z} && \textit{for } 1 \leq N \leq 4. \end{aligned}$$

There are also efficient algorithms to compute $E(K)_{\text{tors}}$. However, little is known about the rank of an elliptic curve. For example, it is now know if there exist elliptic curves of arbitrarily large rank—the highest known example was found in 2006 by Elkies and has rank at least 28 (for a historical summary, see https://web.math.pmf.unizg.hr/~duje/tors/rankhist.html). The last section 5 discusses the famous Birch–Swinnerton-Dyer conjecture about the rank of an elliptic curve, together with a summary of all the factors involved.

Remark 1.8. The Mordell–Weil Theorem 1.6 for elliptic curves was proved by Mordell and then generalized to *abelian varieties* by Weil. Many of the results and techniques in this course apply in this more general setting. Especially, the cohomological arguments and the Tate–Shafarevich group defined in Section 5.1 feature in the study of abelian varieties.

2. Group cohomology

In this section, we introduce Galois cohomology, which is necessary for the proof of the Weak Mordell–Weil Theorem 3.1. We focus mostly on the first two cohomology groups: H^0 and H^1, following [Sil09, Appendix B].

2.1. **Cohomology of finite groups.** Let G be a finite group.

Definition 2.1. A *(right) G-module* with an abelian group M together with an action G, for $\sigma \in G$, $m \mapsto m^\sigma$, such that

$$m^1 = m, \quad (m+m')^\sigma = m^\sigma + (m')^\sigma, \quad (m^\sigma)^\tau = m^{\sigma\tau}.$$

Definition 2.2. The 0*th cohomology group* of M is the submodule

$$H^0(G, M) = M^G = \{m \in M \mid m^\sigma = m \text{ for all } \sigma \in G\}$$

of elements of M invariant under the action of G.

The main example we will be interested in is the Galois group action on the points of an elliptic curve.

Example 2.3. Let E/K be an elliptic curve and L/K be a finite Galois extension. The abelian group $E(L)$ is a $G_{L/K}$-module, with the action given by $(x,y)^\sigma = (x^\sigma, y^\sigma)$, $O^\sigma = O$. Then clearly:

$$H^0(G_{L/K}, E(L)) = E(K).$$

Definition 2.4. Let M and N be G-modules. A *G-module homomorphism* is a group homomorphism

$$\varphi\colon M \to N,$$

commuting with the action of G, i.e. for any $\sigma \in G$ the diagram

$$\begin{tikzcd} M \arrow[r, "\varphi"] \arrow[d, "\sigma"] & N \arrow[d, "\sigma"] \\ M \arrow[r, "\varphi"] & N \end{tikzcd}$$

commutes: for any $m \in M$, $\varphi(m^\sigma) = \varphi(m)^\sigma$.

Note that the functor F mapping $M \mapsto M^G$ is only left exact. Given an exact sequence of G-modules

$$0 \longrightarrow P \xrightarrow{\varphi} M \xrightarrow{\psi} N \longrightarrow 0,$$

the sequence

$$0 \longrightarrow P^G \xrightarrow{\varphi} M^G \xrightarrow{\psi} N^G$$

is exact, but the induced map ψ need not be surjective. To measure how far ψ is from being surjective, we use group cohomology. One way to define group cohomology is to consider F as a functor of $\mathbb{Z}[G]$-modules. Its right derived functor gives the group cohomology of M:

$$R^i(F)(M) = H^i(G, M),$$

or, in other words,

$$H^i(G, M) = \operatorname{Ext}^i_{\mathbb{Z}[G]}(\mathbb{Z}, M).$$

For the reader unfamiliar with homological algebra, we give a very explicit definition of the 1st homology group.

Definition 2.5. We define

$$\begin{aligned} \textit{1-cochains}: &\quad C^1(G, M) = \{\xi \colon G \to M\}, \\ \textit{1-cocycles}: &\quad Z^1(G, M) = \{\xi \in C^1(G, M) \mid \xi(\sigma\tau) = \xi(\sigma)^\tau + \xi(\tau)\}, \\ \textit{1-coboundaries}: &\quad B^1(G, M) = \left\{ \xi \in C^1(G, M) \;\middle|\; \begin{array}{l} \text{there exists } m \in M \text{ such that} \\ \xi(\sigma) = m^\sigma - m \text{ for all } \sigma \in G \end{array} \right\}. \end{aligned}$$

Then note that for $\xi \in B^1(G, M)$ we have

$$\xi(\sigma\tau) = m^{\sigma\tau} - m = (m^\sigma)^\tau - m^\tau + m^\tau - m = \xi(\sigma)^\tau + \xi(\tau),$$

so $B^1(G, M) \subseteq Z^1(G, M)$, and we define the 1*st cohomology group* as

$$H^1(G, M) = \frac{Z^1(G, M)}{B^1(G, M)},$$

the 1-cocycles modulo the relation defining the 1-coboundries.

Remark 2.6. Note that if the action of G on M is trivial, then

$$H^0(G, M) = M, \quad H^1(G, M) = \operatorname{Hom}(G, M).$$

We now formalize the claim that H^1 measures how far ψ is from being surjective.

Proposition 2.7 (Long exact sequence for cohomology)**.** *Let*

$$0 \longrightarrow P \xrightarrow{\varphi} M \xrightarrow{\psi} N \longrightarrow 0$$

be an exact sequence of G-modules. Then there is a long exact sequence

$$\begin{tikzcd}
0 \arrow[r] & H^0(G,P) \arrow[r, "\varphi"] & H^0(G,M) \arrow[r, "\psi"] & H^0(G,N) \arrow[dll, "\delta", rounded corners, to path={ -- ([xshift=2ex]\tikztostart.east) |- ([yshift=-2ex]$(\tikztostart)!0.5!(\tikztotarget)$) -| ([xshift=-2ex]\tikztotarget.west) -- (\tikztotarget)}] \\
 & H^1(G,P) \arrow[r, "\varphi_*"] & H^1(G,M) \arrow[r, "\psi_*"] & H^1(G,N).
\end{tikzcd}$$

The reader familiar with homological algebra will recognize this as the long exact sequence for right derived functors (or Ext). One can prove it by a simple application of Snake Lemma, but we present a direct proof.

Proof. We can assume $P \subseteq M$ is a submodule by replacing it with its image in under φ. We first explicitly define δ. Let $n \in H^0(G,N) = N^G$. By surjectivity of ψ, choose $m \in M$ such that $\psi(m) = n$, and define $\xi \in C^1(G,M)$ by

$$\xi(\sigma) = m^\sigma - m.$$

Since

$$\psi(\xi(\sigma)) = \psi(m^\sigma) - \psi(m) = \psi(m)^\sigma - \psi(m) = n^\sigma - n = 0,$$

$\xi(\sigma) \in \ker\psi$, and by exactness, $\xi(\sigma) \in P$, so $\xi \in C^1(G,P)$. Finally,

$$\xi(\sigma\tau) = m^{\sigma\tau} - m = (m^\sigma)^\tau - m^\tau + m^\tau - m = \xi(\sigma)^\tau + \xi(\tau),$$

so $\xi \in Z^1(G,P)$, and define

$$\delta(n) = \text{class of } \xi \text{ in } H^1(G,P).$$

The only non-trivial part of the proof is showing that

$$\operatorname{im}\delta = \ker\varphi_*, \tag{1}$$

$$\operatorname{im}\psi = \ker\delta. \tag{2}$$

We have that

$$\varphi_*(\delta(n))(\sigma) = \varphi(\xi(\sigma)) = m^\sigma - m,$$

so $\delta(n) \in Z^1(G,M)$, showing that $\operatorname{im}(\delta) \subseteq \ker\varphi_*$. Conversely, if $\xi \in H^1(G,P)$ satisfies $\varphi_*(\xi) = 0$, then

$$\varphi(\xi(\sigma)) = m^\sigma - m$$

for some $m \in M$, and hence the class of ξ is the same as the class of $\delta(\psi(m))$. This shows equation (1).

We proceed similarly to show equation (2). First, note that

$$\delta(\psi(m)) = m^\sigma - m = 0,$$

if $m \in H^0(G,M) = M^G$, showing that $\operatorname{im}\psi \subseteq \ker\delta$. Conversely, let $n \in N^G$ be such that $\delta(n)$ is 0 in $H^1(G,P)$. Then the class of $\xi(\sigma) = m^\sigma - m$ is represented by $\xi(\sigma) = p^\sigma - p$ for some $p \in P$. This shows that for any $\sigma \in G$, $p^\sigma - p = m^\sigma - m$. Hence $(m-p)^\sigma = m - p$, so $m - p \in M^G = H^0(G,M)$, and

$$\psi(m-p) = \psi(m) - \psi(p) = \psi(m) = n.$$

This shows equation (2), and completes the proof. □

Suppose H is a subgroup of G. Then any G-module is an H-module, and if $\xi\colon G \to M$ is a 1-cochain, $\xi_{|H}\colon H \to M$ is a 1-cochain too, and we can define the *restriction homomorphism*

$$\mathrm{Res}\colon H^1(G,M) \to H^1(H,M).$$

If H is moreover a normal subgroup of G, then M^H is a G/H-module. Let $\xi\colon G/H \to M^H$ be a 1-cochain. Then the composition

$$\begin{tikzcd} G \arrow[r, two heads] & G/H \arrow[r, "\xi"] & M^H \arrow[r, hook] & M \end{tikzcd}$$

defines the *inflation homomorphism*

$$\mathrm{Inf}\colon H^1(G/H,M^H) \to H^1(G,M).$$

Proposition 2.8 (Inflation–restriction sequence)**.** *Let M be a G-module and H be a normal subgroup of G. Then there is an exact sequence*

$$\begin{tikzcd} 0 \arrow[r] & H^1(G/H,M^H) \arrow[r, "\mathrm{Inf}"] & H^1(G,M) \arrow[r, "\mathrm{Res}"] & H^1(H,M). \end{tikzcd}$$

Proof. It is clear that $\mathrm{Res}\circ\mathrm{Inf} = 0$. To show Inf is injective, suppose $\xi\colon G/H \to M^H$ is a 1-cocycle such that $\mathrm{Inf}([\xi])$ is 0 in $H^1(G,M)$. Then there is an $m \in M$ such that $\xi(\sigma) = m^\sigma - m$ for all $\sigma \in G$. But $\xi(\sigma)$ only depends on the class of σ in G/H, so $m^\sigma - m = m^{\sigma\tau} - m$ for any $\tau \in H$, showing that $m^\tau = m$, i.e. $m \in M^H$. This shows that ξ is a 1-coboundary, i.e. $[\xi] = 0$ in $H^1(G/H,M^H)$.

Finally, to show exactness at $H^1(G,M)$, take any 1-cocycle $\xi\colon G \to M$ such that $\mathrm{Res}([\xi]) = 0$. Then there exists $m \in M$ such that

$$\xi(\tau) = m^\tau - m \quad \text{for all } \tau \in H.$$

Subtracting $\sigma \mapsto m^\sigma - m$ from ξ does not change the class $[\xi] \in H^1(G,M)$, so we may assume $\xi(\tau) = 0$ for all $\tau \in H$. But then

$$\xi(\tau\sigma) = \xi(\tau)^\sigma + \xi(\sigma) = \xi(\sigma),$$

so ξ only depends on the class of σ in G/H. Finally, since H is normal, there exists $\tau' \in H$ such that $\sigma\tau = \tau'\sigma$, and hence

$$\xi(\sigma)^\tau = \xi(\sigma)^\tau + \xi(\tau) = \xi(\sigma\tau) = \xi(\tau'\sigma) = \xi(\sigma),$$

showing that ξ induces a 1-cocycle $G/H \to M^H$, whose image under the inflation map is $[\xi]$. □

2.2. **Galois cohomology.** In the study of elliptic curves, group cohomology will prove useful for Galois groups. In the previous section, we assumed that the group G is finite, so the theory only applies to finite extensions. Recall that to study infinite Galois extensions, we note that for an extension $\bar{K}/K$:

$$G_{\bar{K}/K} = \varprojlim G_{L/K},$$

where the limit is taken over all finite Galois extensions L of K contained in $\bar{K}$. Giving $G_{L/K}$ the discrete topology, this gives $G_{\bar{K}/K}$ the profinite topology. Explicitly, a basis of open sets around the identity consists of the sets of normal subgroups having finite index in $G_{\bar{K}/K}$, i.e. kernels of maps $G_{\bar{K}/K} \to G_{L/K}$ for finite Galois extensions L/K.

This topology plays a crucial role in the study of infinite extensions; for example, the fundamental theorem of infinite Galois theory gives a correspondence between intermediate extensions and **closed** subgroups of the Galois group. We hence adapt the definition of the first homology group to account for the topology.

Definition 2.9. Let $\bar{K}/K$ be a Galois extension and $G = G_{\bar{K}/K}$. A map $\xi\colon G \to M$ is *continuous* if it is continuous with respect to the profinite topology on $G_{\bar{K}/K}$ and discrete topology on M. We then define for

$$\begin{aligned} \textit{1-cochains}: \quad & C^1_{\text{cont}}(G, M) = \{\text{continuous maps } \xi\colon G \to M\}, \\ \textit{1-cocycles}: \quad & Z^1_{\text{cont}}(G, M) = \{\xi \in C^1_{\text{cont}}(G, M) \mid \xi(\sigma\tau) = \xi(\sigma)^\tau + \xi(\tau)\} \end{aligned}$$

Any map of the form $\sigma \mapsto m^\sigma - m$ is automatically continuous, and hence we can define the 1*st cohomology group of the G-module M* to be

$$H^1(G, M) = \frac{Z^1_{\text{cont}}(G, M)}{B^1(G, M)}.$$

Remark 2.10. As before, if the action of G on M is trivial, then $H^0(G, M) = M$, and $H^1(G, M) = \operatorname{Hom}_{\text{cont}}(G, M)$.

Both the long exact sequence for cohomology 2.7 and the inflation–restriction sequence 2.8 also exist for Galois cohomology.

3. Weak Mordell–Weil Theorem

We will now apply the results from Section 2 to prove the Weak Mordell–Weil Theorem.

Theorem 3.1 (Weak Mordell–Weil)**.** *Let K be a number field and E/K be an elliptic curve,. Then for any $m \geq 2$, the quotient group*

$$E(K)/mE(K)$$

is finite.

3.1. **Kummer sequence and pairing.** In other words, we want to prove that the cokernel of the multiplication by m map is finite. Therefore, we look at the following exact sequence of $G_{\bar{K}/K}$-modules:

$$\begin{tikzcd} 0 \arrow[r] & E[m] \arrow[r] & E(\bar{K}) \arrow[r, "m"] & E(\bar{K}) \arrow[r] & 0. \end{tikzcd}$$

The long exact sequence for Galois cohomology 2.7 then yields

$$\begin{tikzcd} 0 \arrow[r] & E(K)[m] \arrow[r] & E(K) \arrow[r, "m"] & E(K) \arrow[dll, rounded corners, "\delta"] \\ & H^1(G_{\bar{K}/K}, E[m]) \arrow[r] & H^1(G_{\bar{K}/K}, E(\bar{K})) \arrow[r, "m"] & H^1(G_{\bar{K}/K}, E(\bar{K})). \end{tikzcd}$$

From the middle of this sequence, we extract the following short exact sequence, called the *Kummer sequence* for E/K:

$$(*) \qquad 0 \longrightarrow \tfrac{E(K)}{mE(K)} \xrightarrow{\delta} H^1(G_{\bar{K}/K}, E[m]) \longrightarrow H^1(G_{\bar{K}/K}, E(\bar{K}))[m] \longrightarrow 0$$

Using the general construction of δ from the proof of Proposition 2.7, we can describe it explicitly: for $P \in E(K)$, we choose some $Q \in E(\bar{K})$ such that $mQ = P$, and represent $\delta(P)$ by the 1-cocycle

$$c\colon G_{\bar{K}/K} \to E[m], \quad c(\sigma) = Q^\sigma - Q.$$

This gives a pairing

$$\kappa\colon E(K) \times G_{\bar{K}/K} \to E[m], \quad \kappa(P, \sigma) = c(\sigma) = Q^\sigma - Q,$$

called the *Kummer pairing.* By our construction, it is clearly independent of the choice of Q.

We now use the inflation-restriction sequence 2.8 to show that we may assume $E[m] \subseteq E(K)$, by replacing K with a finite extension.

Lemma 3.2. *Let L/K be a finite Galois extension. If $E(L)/mE(L)$ is finite, then $E(K)/mE(K)$ is also finite.*

Proof. We have a natural map

$$\varphi\colon \frac{E(K)}{mE(K)} \longrightarrow \frac{E(L)}{mE(L)}$$

and we show that its kernel is finite. By the inflation-restriction sequence 2.8 together with the Kummer sequence $(*)$, we get the following diagram with exact rows

$$\begin{tikzcd}
0 \arrow[r] & \ker(\varphi) \arrow[r] & \frac{E(K)}{mE(K)} \arrow[r, "\varphi"] \arrow[d, hook, "\delta_K"] & \frac{E(L)}{mE(L)} \arrow[d, hook, "\delta_L"] \\
0 \arrow[r] & H^1(G_{L/K}, E(L)[m]) \arrow[r, "\mathrm{Inf}"] & H^1(G_{\bar{K}/K}, E[m]) \arrow[r, "\mathrm{Res}"] & H^1(G_{\bar{L}/L}, E[m])
\end{tikzcd}$$

where the square commutes. The rest of the proof follows from the five lemma, but we prove it directly: there exists an injective map $\lambda\colon \ker(\varphi) \to H^1(G_{L/K}, E[m])$. Since both $G_{L/K}$ and $E[m]$ are finite, this indeed shows $\frac{E(K)}{mE(K)}$ is finite.

Take $P \in \ker(\varphi)$, and note that $\mathrm{Res}(\delta_K(P)) = \delta_L(\varphi(P)) = 0$, so $\delta_K(P) \in \ker \mathrm{Res} = \operatorname{im} \mathrm{Inf}$, and hence there exists $[\xi] \in H^1(G_{L/K}, E[m])$ such that $\mathrm{Inf}([\xi]) = \delta_K(P)$. Set

$$\lambda(P) = [\xi].$$

To show that λ is injective, suppose $\lambda(P) = 0$. Then

$$\delta_K(P) = \mathrm{Inf}(0) = 0,$$

so $P = O$, since δ_K is injective. □

From now on, suppose (possibly replacing K by a finite extension L) that $E[m] \subseteq E(K)$. Then the action of $G_{\bar{K}/K}$ on $E[m]$ is trivial and hence

$$H^1(G_{\bar{K}/K}, E[m]) = \mathrm{Hom}(G_{\bar{K}/K}, E[m]),$$

and the Kummer sequence $(*)$ gives an injective homomorphism

$$\delta\colon \frac{E(K)}{mE(K)} \hookrightarrow \operatorname{Hom}(G_{\bar{K}/K}, E[m]), \qquad P \mapsto \kappa(P,-).$$

Proposition 3.3. *For a fixed P, the kernel of the map $\sigma \mapsto \kappa(P,\sigma)$ is $G_{\bar{K}/L}$, where*

$$L = K(m^{-1}E(K))$$

is the compositum of all fields $K(Q)$ over all $Q \in E(\bar{K})$ such that $mQ \in E(K)$.

In other words, L is obtained by adjoining the *mth roots* of elements of $E(K)$. This is a similar process to what is done in classical Kummer theory.

Proof. If $\sigma \in G_{\bar{K}/L}$, then $\kappa(P,\sigma) = Q^\sigma - Q = O$, since $Q \in E(L)$ by definition of L. Conversely, suppose that $\sigma \in G_{\bar{K}/K}$ satisfies $\kappa(P,\sigma) = O$ for all $P \in E(K)$. Then for every $Q \in E(\bar{K})$ such that $mQ \in E(K)$ we have

$$O = \kappa(mQ,\sigma) = Q^\sigma - Q,$$

and hence σ fixes any such Q. This shows that σ fixes the compositum, L, and hence $\sigma \in G_{\bar{K}/L}$. □

Having described the kernel, recalling that $G_{\bar{K}/K}/G_{\bar{K}/L} \cong G_{L/K}$, we obtain the following corollary.

Corollary 3.4. *The Kummer pairing induces a perfect bilinear pairing*

$$E(K)/mE(K) \times G_{L/K} \to E[m],$$

where L is the field defined in Proposition 3.3, and hence there are injective homomorphisms

$$E(K)/mE(K) \hookrightarrow \operatorname{Hom}(G_{L/K}, E[m]),$$
$$P \mapsto \kappa(P,-).$$

Note that $E[m]$ is always finite. This is clear from the addition equations, and Proposition 3.7 gives a formal proof. Therefore, to complete the proof of weak Mordell–Weil Theorem 3.1, it suffices to show that L/K is finite. We do this in two steps: first show L/K has certain number-theoretic properties, and then use classical number theoretic results to show that any such field extension is finite.

3.2. Local fields and reduction of elliptic curves.

For conics, the set of rational solutions is fully determined by the set of *local* solutions.

Theorem 3.5 (Hasse–Minkowski, Local–Global principle, [Cas91, Chap. 3])**.** *Suppose f is a quadratic form over $\mathbb{Q}$. Then*

$$f(x,y)=0 \text{ for some } (x,y)\in\mathbb{Q} \quad \text{if and only if} \quad \begin{cases} f(x_p,y_p)=0 & \text{for some } (x_p,y_p)\in\mathbb{Q}_p \\ & \text{for any prime } p, \\ f(x_\infty,y_\infty)=0 & \text{for some } (x_\infty,y_\infty)\in\mathbb{R}. \end{cases}$$

More generally, if K is a number field and f is a quadratic form over K, then

$$f(x,y)=0 \text{ for some } (x,y)\in K \quad \text{if and only if} \quad \begin{array}{l} f(x_v,y_v)=0 \text{ for some } (x_v,y_v)\in K_v \\ \text{for all valuations } v \text{ on } K. \end{array}$$

We will explain the precise meaning of a valuation v and K_v soon. One reason this is extremely powerful is that we have good analytic methods, such as Hensel's lemma, for solving equations in a local setting.

For higher degree equations (or higher genus curves), this principle fails. However, looking at the *local picture* still allows us to extract a lot of information about elliptic curves.

In order to study the torsion of elliptic curves, we consider the *reduction* of elliptic curves modulo a prime. Over $\mathbb{Q}$, this would mean, we take a prime $p \in \mathbb{Z}$ and look at the solutions of

$$\tilde{E}_p : y^2 \equiv x^3 + \tilde{a}x + \tilde{b} \mod p.$$

As long as this remains an ellpitic curve, so $\Delta \not\equiv 0 \mod p$, the torsion of $E(\mathbb{Q})$ embeds into

$$\widetilde{E}_p(\mathbb{F}_p) = \{(x,y)\in \mathbb{F}_p : y^2 = x^3 + \tilde{a}x + \tilde{b}\} \cup \{\tilde{O}\}.$$

Example 3.6. Consider

$$E : y^2 = x^3 + x \text{ with } \Delta = 4$$

over $\mathbb{Q}$. We can easily check that

$$\tilde{E}(\mathbb{F}_3) = \{\tilde{O}, (0,0), (2,1), (2,2)\} \cong \mathbb{Z}/4\mathbb{Z},$$

$$\tilde{E}(\mathbb{F}_5) = \{\tilde{O}, (0,0), (2,0), (3,0)\} \cong \mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z},$$

since a point (x,y) has order 2 if and only if $y=0$. We know that $E(\mathbb{Q})_{\text{tors}}$ embeds into both of these groups, since 3, 5 do not divide $\Delta = 4$. Hence is has order 1 or 2. But it contains the points $(0,0)$ of order 2, and thus

$$E(\mathbb{Q})_{\text{tors}} = \{O, (0,0)\} \cong \mathbb{Z}/2\mathbb{Z}.$$

See [ST92, Chap. IV.3] for more details on the $K=\mathbb{Q}$ case.

Over a general number field, we need to find the appropriate replacement for reduction modulo a prime p. Once again, we have to understand what happens *locally*, by focusing at one prime ideal at a time. Let

$$M_K = \{\text{inequivalent absolute values on } K\}$$

which are classified by Ostrowski's Theorem into two categories.

- The *infinite places* $v \in M_K^\infty$: $|-|_v$ such that $|x|_v = |x|$ for $x \in \mathbb{Q}$, the standard absolute value on $\mathbb{Q}$. The completion K_v of K with respect to these absolute values is either $\mathbb{R}$ or $\mathbb{C}$.
- The *finite places* $v \in M_K^0$ corresponding to prime ideals $\mathfrak{p}$ of $\mathcal{O}_K$, the ring of integers of K:

$$|x|_v = N\mathfrak{p}^{-\operatorname{ord}_{\mathfrak{p}}(x)} \qquad \text{for } x \in K,$$

where $\operatorname{ord}_{\mathfrak{p}}(x)$ is the highest power of the prime ideal $\mathfrak{p}$ dividing the ideal (x). These restrict to the p-adic absolute values on $\mathbb{Q}$, where $\mathfrak{p} \cap \mathbb{Q} = (p)$, and the completion of K with respect to them is a local field $K_{\mathfrak{p}}$, a finite extension of $\mathbb{Q}_p$. This means that the ring of integers $\mathcal{O}_\mathfrak{p}$ of $K_\mathfrak{p}$ has a unique maximal ideal $\mathfrak{m}_\mathfrak{p}$. The *residue field* $k_\mathfrak{p} = k_v$ of $\mathcal{O}_\mathfrak{p}$ is

$$k_\mathfrak{p} = \mathcal{O}_\mathfrak{p}/\mathfrak{m}_\mathfrak{p} \cong \mathbb{F}_q \supseteq \mathbb{F}_p = \mathbb{Z}_p/p\mathbb{Z}_p,$$

a finite extension of $\mathbb{F}_p$.

For a number field K, the analogue of *reduction mod p* for $\mathbb{Q}$ will be to consider the curve over the residue field $k_\mathfrak{p}$. Fix a finite place $v \in M_K^0$ corresponding to a prime ideal $\mathfrak{p}$. After a change of coordinates, we may assume that the equation for E has $a, b \in \mathcal{O}_K$, and minimal possible $\operatorname{ord}_\mathfrak{p}(\Delta)$ (see [Sil09, Chap. VII.1]).

There is a natural reduction map $\mathcal{O}_\mathfrak{p} \to k_\mathfrak{p}$, $x \mapsto \tilde{x}$, and we write

$$\tilde{E} : y^2 = x^3 + \tilde{a}x + \tilde{b}.$$

Then we define the *reduction map* as

$$E(K) \to \tilde{E}(k_\mathfrak{p}), \qquad [x : y : z] \mapsto [\tilde{x} : \tilde{y} : \tilde{z}]$$

by choosing coordinates such that $x, y, z \in \mathcal{O}_\mathfrak{p}$. If $v(\Delta) = 0$, then $\tilde{E}$ defines an elliptic curves over k_v, and we say E has *good reduction* at v. In this case, the above map is a group homomorphism, and we have the following result.

Proposition 3.7 ([Sil09, Prop. VII.3.1])**.** *Let $v \in M_K^0$ be a discrete valuation such that $v(m) = 0$ and E has good reduction at v. Then the reduction map*

$$E(K)[m] \to \tilde{E}_v(k_v)$$

is injective.

Remark 3.8. If E does not have good reduction at v, we say it has *bad reduction.* In that case, we have a group structure on E_0, the points that reduce to non-singular points on $\tilde{E}$. Its reduction $\tilde{E}_{\text{non-sing}}(k_v)$ is a group isomorphic to $\bar{k_v}^*$ or $\bar{k_v}^+$, depending on whether the singularity is a cusp or a node.

For full details of elliptic curves over finite fields, the reader is encouraged to consult [Sil09, Chap. VII].

3.3. **The extension L/K is finite.** To show the finiteness of L/K, we will first show that L/K is *unramified* outside S. To understand what this means, in this context, it is most natural to explain this notion as follows. For a prime $\mathfrak{p}$ of K, choose a prime $\mathfrak{P}$ in L such that $\mathfrak{P} \cap K = \mathfrak{p}$. The *decomposition group of $\mathfrak{P}$ over $\mathfrak{p}$* is

$$G_{\mathfrak{P}/\mathfrak{p}} = \{\sigma \in G_{L/K} \mid \sigma(\mathfrak{P}) = \mathfrak{P}\}.$$

Note that $G_{\mathfrak{P}/\mathfrak{p}}$ acts naturally on the completion $L_\mathfrak{P}/K_\mathfrak{p}$. There is a natural map

$$G_{\mathfrak{P}/\mathfrak{p}} \to G_{l_\mathfrak{P}/k_\mathfrak{p}},$$

and the *inertia group $I_{\mathfrak{P}/\mathfrak{p}}$ of $\mathfrak{P}$ over $\mathfrak{p}$* is the kernel of this map. The prime $\mathfrak{P}/\mathfrak{p}$ is *unramified* if and only if $I_{\mathfrak{P}/\mathfrak{p}}$ is trivial.

A more detailed discussion of ramification can be found in [Lan94, Chap. I].

Proposition 3.9. *Let $L = K(m^{-1}E(K))$ be as in Proposition 3.3. We show that:*

(1) *The extension L/K is abelian of exponent m, i.e. the Galois group G is abelian and every element of G has order dividing m.*

(2) *Let*

$$S = \{v \in M_K^0 : E \textit{ has bad reduction at } v\} \cup \{v \in M_K^0 : v(m) \neq 0\} \cup M_K^\infty.$$

Then L/K is unramified outside S.

Proof. By Proposition 3.3, we see that there is an injection

$$G_{L/K} \hookrightarrow \mathrm{Hom}(E(K), E[m]), \qquad \sigma \mapsto \kappa(-, \sigma),$$

and (1) follows. For (2), take any $v \in M_K$, $v \notin S$. Since L is the compositum of $K' = K(Q)$ for $Q \in E(\bar{K})$ such that $[m]Q \in E(K)$, it is enough to show that K'/K is unramified at v. Choose $v' \in M_{K'}$ above v and let k'_v/k_v be the corresponding extension of residue fields. Since E has good reduction at v, it has good reduction at v', and we have the reduction map

$$E(K') \to \tilde{E}(k'_{v'}).$$

Let $I_{v'/v} \subseteq G_{K'/K}$ be the inertia group for v'/v and take any $\sigma \in I_{v'/v}$. By definition of the inertia group, σ acts trivially on $\tilde{E}(k'_{v'})$, so

$$\widetilde{Q^\sigma - Q} = \widetilde{Q}^\sigma - \widetilde{Q} = \widetilde{O}.$$

On the other hand, $mQ \in E(K)$ implies that

$$m(Q^\sigma - Q) = (mQ)^\sigma - mQ = O,$$

so $Q^\sigma - Q$ is a point of order dividing m that is in the kernel of the reduction modulo v' map. Thus, by Proposition 3.7, we have that $Q^\sigma - Q = O$, and hence $Q^\sigma = Q$. This shows that $I_{v'/v}$ is trivial, and hence K' is unramified over K at v'. This completes the proof that L/K is unramified outside of S. □

Finally, we show that any extension L/K satisfying these properties is finite.

Theorem 3.10. *Let K be a number field, $S \subseteq M_K$ be a finite set of places that contains M_K^∞, and let $m \geq 2$ be an integer. Let L/K be the maximal abelian extension of K having exponent m that is unramified outside S. Then L/K is a finite extension.*

Proof. By replacing K with a finite extension, we may assume without loss of generality that $\boldsymbol{\mu}_m \subseteq K$.

Note that increasing the set S makes the field L larger. If the class group of K is

$$C_K = \{\mathfrak{a}_1, \dots, \mathfrak{a}_n\},$$

then for any $\mathfrak{p}$ such that $\mathrm{ord}_{\mathfrak{p}}(\mathfrak{a}_i) \neq 0$ for some i, add the valuation corresponding to $\mathfrak{p}$ to the set S. Then the ring

$$R_S = \{a \in K \mid v(a) \geq 0 \text{ for any } v \in M_K \setminus S\}$$

is a principal ideal domain. By adding the valuations for which $v(m) \neq 0$, we can also guarantee that $v(m) = 0$ for $v \in M_K \setminus S$.

By Kummer theory, L is the largest subfield of

$$K(\sqrt[m]{a} : a \in K)$$

which is unramified outside S. Take $v \in M_K \backslash S$. One can show (Lemma A.2) that $K(\sqrt[m]{a})/K$ is unramified at v if and only if

$$\operatorname{ord}_v(a) \equiv 0 \mod m.$$

Since $K(\sqrt[m]{a})$ only depends on the class of a in $K^*/(K^*)^m$, if we write

$$T_S = \{a \in K^*/(K^*)^m : \operatorname{ord}_v(a) \equiv 0 \mod m \text{ for all } v \notin S\},$$

then

$$L = K(\sqrt[m]{a} : a \in T_S).$$

We will show that T_S is finite, thus completing the proof. This is a simple application of the Dirichlet S-units Theorem A.3, which says that R_S^* is finitely-generated. We consider the obvious map

$$\varphi\colon R_S^* \to T_S$$

and note that $(R_S^*)^m$ is contained in the kernel, so φ factors through a map

$$R_S^*/(R_S^*)^m \to T_S.$$

Since R_S^* is finitely-generated, we see that the codomain is finite, and hence it is enough to prove that φ is surjective. Take any $a \in K^*$, representing an element of T_S. Then (a) as an ideal of R_S factors into primes

$$(a) = \mathfrak{p}_1^{\alpha_1} \ldots \mathfrak{p}_n^{\alpha_n}.$$

Since $v(a) \geq 0$ for $v \notin S$, the primes $\mathfrak{p}_i$ correspond to valuations $v_i \notin S$, for which

$$\alpha_i = \operatorname{ord}_{v_i}(a) \equiv 0(m).$$

This shows that (a) is an mth power of an ideal of R_S. Since R_S is a principal ideal domain, there exists $b \in K^*$ such that

$$(a) = (b^m).$$

Thus $a = ub^m$ for some $u \in R_S^*$. But then a and u define the same element of T_S, which shows that $\varphi(u) = a$. Thus φ is surjective, and hence T_S is finite, which shows that $L = K(\sqrt[m]{a} : a \in T_S)$ is a finite extension of K. □

Remark 3.11. An alternative way to prove Theorem 3.10 is to use Hermite's Theorem [Lan94, pp. 122] that says that there are only a finite number of fields L of bounded degree over K unramified outside S. One shows that the extensions $K(Q)$ for $Q \in m^{-1}E(K)$ have bounded degree over K, and hence there are only finitely many of them.

4. Heights on projective spaces and proof of Mordell–Weil Theorem

4.1. **Descent.** We have now shown (Theorem 3.1) that $E(K)/mE(K)$ is finite. We now present an *infinite descent* argument, which shows that, with an approporiate theory of *heights*, this will be enough to conclude that $E(K)$ is finitely-generated.

Theorem 4.1 (Descent Theorem, [Sil09, Th. VIII.3.1])**.** *Let A be an abelian group and $h\colon A \to \mathbb{R}$ be a function such that*

(1) *for $Q \in A$, there exists $C_1 = C_1(A, Q)$ such that*

$$h(P+Q) \leq 2h(P) + C_1 \qquad \textit{for any } P \in A,$$

(2) *there exists $m \geq 2$ and $C_2 = C_2(A)$ such that*
$$h(mP) \geq m^2 h(P) - C_2,$$

(3) *for any C_3, the set*
$$\{P \in A : h(P) \leq C_3\}$$
is finite. If for the number m from (2), A/mA is finite, then A is finitely-generated.

Proof. Let $Q_1, \ldots, Q_r \in A$ represent the cosets in A/mA. The idea of the proof is to show that for any $P \in A$, there exist $a_i \in \mathbb{Z}$ such that
$$h\left(P - \sum_{i=1}^{r} a_i Q_i\right) \leq C_3$$
for some constant C_3 independent of P. Then A will be generated by
$$\{Q_1, \ldots, Q_r\} \cup \{Q \in A : h(Q) \leq C_3\}$$
which is a finite set by (3).

Since $Q_1, \ldots, Q_r$ represent the cosets in A/mA, there exists $1 \leq i_1 \leq r$ and $P_1 \in A$ such that
$$P = mP_1 + Q_{i_1}.$$
Similarly, there exists $1 \leq i_2 \leq r$ and $P_2 \in A$ such that
$$P_1 = mP_2 + Q_{i_2},$$
and continuing this way, we obtain elements $P_1, \ldots, P_n \in A$ and numbers $1 \leq i_1, \ldots, i_n \leq r$ such that
$$\begin{aligned} P &= mP_1 + Q_{i_1}, \\ P_1 &= mP_2 + Q_{i_2}, \\ P_2 &= mP_3 + Q_{i_3}, \\ &\vdots \\ P_{n-1} &= mP_n + Q_{i_n}. \end{aligned}$$
Now, for any $1 \leq j \leq n$:
$$\begin{aligned} h(P_j) &\leq \tfrac{1}{m^2}(h(mP_j) + C_2) && \text{by (2)} \\ &= \tfrac{1}{m^2}(h(P_{j-1} - Q_{i_j}) + C_2) && \text{by definition of } P_j \\ &\leq \tfrac{1}{m^2}(2h(P_{j-1}) + C_1' + C_2) && \text{by (1),} \end{aligned}$$
where $C_1' = \max\{C_1(A, -Q_i) \mid 1 \leq i \leq r\}$. Then we have that
$$\begin{aligned} h(P_n) &\leq \left(\tfrac{2}{m^2}\right)^n h(P) + \left(\tfrac{1}{m^2} + \tfrac{2}{m^2} + \cdots + \tfrac{2^{n-1}}{m^{2n}}\right)(C_1' + C_2) \\ &< \left(\tfrac{2}{m^2}\right)^n h(P) + \tfrac{C_1' + C_2}{m^2 - 2} \\ &\leq \tfrac{1}{2^n} h(P) + \tfrac{1}{2}(C_1' + C_2) && \text{since } m \geq 2. \end{aligned}$$
Therefore, for a given P, we may choose n large enough so that
$$h(P_n) \leq 1 + \frac{1}{2}(C_1' + C_2) = C_3,$$

and the constant C_3 is independent of P. Thus we have indeed expressed

$$P = m^n P_n + \sum_{j=1}^{n} m^{j-1} Q_{i_j},$$

a linear combination of elements from the finite set

$$\{Q_1, \ldots, Q_r\} \cup \{Q \in A : h(Q) \leq C_3\}.$$

This completes the proof. □

To complete the proof of the Mordell–Weil Theorem 1.6, we just have to find a function $h \colon E(K) \to \mathbb{R}$ that satisfies the above three properties. We provide a summary, omitting some of the proofs—the full details can, as always, be found in [Sil09, Chap. VIII.5,6].

4.2. **Heights on projective spaces.** This section follows [Sil09, Chap. VIII.5].

Over $\mathbb{Q}$, it is easy to write down such a function: for $P \in \mathbb{P}^N(\mathbb{Q})$, write $P = [x_0 : \ldots : x_N]$ such that $x_i \in \mathbb{Z}$ and $\gcd(x_0, \ldots, x_N) = 1$, and define

$$H(P) = \max\{|x_0|, \ldots, |x_N|\},$$
$$h(P) = \log H(P) \text{ for } P \in E(\mathbb{Q}).$$

This method works because $\mathbb{Z}$ is a principal ideal domain, and hence we can write a point in that form. For a general number field K, the ring of integers $\mathcal{O}_K$ may not be a principal ideal domain, and instead of trying to write a point in a specific form, we simply take **all** the absolute values on K into account.

Definition 4.2. Let M_K be the set of absolute values on K. We then define the *local degree* of K at $v \in M_K$ to be

$$n_v = [K_v : \mathbb{Q}_v].$$

For $P = [x_0 : \ldots : x_N] \in \mathbb{P}^N(K)$, we let

$$H_K(P) = \prod_{v \in M_K} \max\{|x_0|_v, \ldots, |x_n|_v\}^{n_v}.$$

Using a result about the behavior of local degrees in towers A.4, and the Product Formula A.5, we obtain the following proposition.

Proposition 4.3. *Let $P \in \mathbb{P}^N(K)$. Then*

(1) *$H_K(P)$ is well-defined, i.e. independent of the choice of coordinates,*
(2) *$H_K(P) \geq 1$,*
(3) *if L/K is a finite extension, then $H_L(P) = H_K(P)^{[L:K]}$.*

Part (3) prompts the following definition.

Definition 4.4. Let $P \in \mathbb{P}^N(\bar{\mathbb{Q}})$. The *height of P* is defined as follows: choose a number field K such that $P \in \mathbb{P}^N(K)$, and let

$$H(P) = H_K(P)^{1/[K:\mathbb{Q}]}.$$

By Proposition 4.3, the height is well-defined and at least 1.

4.3. **Heights on elliptic curves.** We now specialize the previous subsection to elliptic curves, following [Sil09, Chap. VIII.6]. Fix an elliptic curve E/K and a function on it $f \in \bar{K}(E)$. We can then identify f with a function

$$f\colon E \to \mathbb{P}^1, \qquad P \mapsto \begin{cases} [1:0] & \text{if } P \text{ is a pole,} \\ [f(P):1] & \text{otherwise.} \end{cases}$$

For example, we can take $f(P) = x(P)$, where $x(P) = x$ for $P = (x, y)$.

Definition 4.5. The *logarithmic height* on $\mathbb{P}^N$ is $h\colon \mathbb{P}^N(\bar{\mathbb{Q}}) \to \mathbb{R}$ given by $h(P) = \log(H(P))$. The *height on E relative to f* is

$$h_f\colon E(\bar{K}) \to \mathbb{R}, \qquad h_f(P) = h(f(P)).$$

Theorem 4.6. *For any non-trivial even function $f \in \bar{K}(E)$, the height function h_f satisfies the assumptions (1)–(3) of the Descent Theorem 4.1. Hence the Mordell–Weil Theorem 1.6 follows.*

Proof. The proof of this is long and technical, so we omit it here. It can be found, as a series of lemmas and propositions, in [Sil09, Chap. VIII.5,6]. An example of such an even function is $f(P) = x(P)$, mapping P to the x-coordinate of P, and hence the Mordell–Weil Theorem 1.6 follows. □

4.4. **The canonical height.** One of the key theorems involved in the proof of Theorem 4.6 shows ([Sil09, Theorem 6.2]) that for an even f and P, Q on E:

$$h_f(P+Q) + h_f(P-Q) = 2h_f(P) + 2h_f(Q) + C$$

for some constant C dependent on E and f, but independent of P and Q. This shows that h_f is *almost* a quadratic form. Moreover, h_f is not canonical—dependent on the choice of f. André Néron and John Tate independently showed that there is a canonical height, independent of choices, and it actually is a quadratic form.

Definition 4.7. The *canonical* (or *Néron–Tate*) *height on E/K* is $\hat{h}\colon E(\bar{K}) \to \mathbb{R}$ given by

$$\hat{h}(P) = \frac{1}{\deg f} \lim_{N \to \infty} 4^{-N} h_j([2^N]P)$$

for a non-constant even function $f \in K(E)$.

Proposition 4.8 (Néron, Tate, [Sil09, Prop. 9.1, Th. 9.3])**.** *Let E/K be an elliptic curve, $f \in K(E)$ be a non-constant even function. Then the canonical height $\hat{h}$ exists and is independent of the choice of f. Moreover, it obeys*

$$\hat{h}(P+Q) + \hat{h}(P-Q) = 2\hat{h}(P) + 2\hat{h}(Q),$$

and hence the pairing

$$\langle -, - \rangle\colon E(\bar{K}) \times E(\bar{K}) \to \mathbb{R}$$
$$\langle P, Q \rangle = \hat{h}(P+Q) - \hat{h}(P) - \hat{h}(Q),$$

is bilinear.

We now define an invariant of an elliptic curve related to the free part of $E(K)$.

Definition 4.9. The *elliptic regulator* of E/K is

$$E_{E/K} = \det(\langle P_i, P_j \rangle)_{1 \leq i,j \leq r}$$

for generators $P_1, \dots, P_r$ of the free part $E(K)/E_{\text{tors}}(K)$ of $E(K)$.

The regulator E/K is the volume of the fundamental domain for $E(K)/E_{\text{tors}}(K)$ with respect to the quadratic form $\hat{h}$.

5. The rank and the Birch–Swinnerton-Dyer Conjecture

We have shown that the group $E(K)$ is finitely-generated, and hence of the form

$$E(K) \cong \mathbb{Z}^r \oplus \mathbb{Z}/n_1\mathbb{Z} \oplus \cdots \oplus \mathbb{Z}/n_k\mathbb{Z}.$$

As mentioned, there are effective methods of computing the torsion for $E(K)$ (see Example 3.6). In this section, we discuss the rank of $E(K)$, about which little is known so far. We present the famous Birch–Swinnerton-Dyer conjecture, following [Sil09, Appendix C.16] and [Wil06].

5.1. **The Tate–Shafarevich group.** The problem of computing the rank is equivalent to computing the size of the group $E(K)/mE(K)$. Indeed, suppose we know the torsion of $E(K)$, i.e.

$$E(K) \cong F \oplus \mathbb{Z}/n_1\mathbb{Z} \oplus \cdots \oplus \mathbb{Z}/n_k\mathbb{Z}$$

for a free abelian group F. Supposing, for simplicity, that m is coprime to all the n_i, we obtain

$$E(K)/mE(K) \cong (\mathbb{Z}/m\mathbb{Z})^r,$$

where $r = \text{rank}(F) = \text{rank}(E)$. Moreover, if it was possible to compute the generators of $E(K)/mE(K)$ effectively, then it would be possible to compute the generators of $E(K)$ effectively.

Let us hence investigate the proof of the Weak Mordell–Weil Theorem 3.1 and see whether we could make it explicit. We considered the short exact sequence

$$\begin{tikzcd} 0 \arrow[r] & E[m] \arrow[r] & E(\bar{K}) \arrow[r, "m"] & E(\bar{K}) \arrow[r] & 0, \end{tikzcd}$$

which gave the long exact sequence for Galois cohomology 2.7

$$\begin{tikzcd} 0 \arrow[r] & E(K)[m] \arrow[r] & E(K) \arrow[r, "m"] & E(K) \arrow[dll, "\delta"] \\ & H^1(G_{\bar{K}/K}, E[m]) \arrow[r] & H^1(G_{\bar{K}/K}, E(\bar{K})) \arrow[r, "m"] & H^1(G_{\bar{K}/K}, E(\bar{K})), \end{tikzcd}$$

from which we extracted the Kummer sequence for E/K:

$$(*) \qquad \begin{tikzcd} 0 \arrow[r] & \frac{E(K)}{mE(K)} \arrow[r] & H^1(G_{\bar{K}/K}, E[m]) \arrow[r] & H^1(G_{\bar{K}/K}, E(\bar{K}))[m] \arrow[r] & 0. \end{tikzcd}$$

Computing the group $E(K)/mE(K)$ is hence the same as computing the kernel of the latter map.

Remark 5.1. The cohomology group $H^1(G_{\bar{K}/K}, E(\bar{K}))$ that appears above has a geometric interpretation as (classes of) *principal homogeneous spaces*, known as the Weil–Châtelet group, $WC(E/K)$. This interpretation can be found in [Cas91, Chap. 22, 23].

Since these groups are not easily accessible, we can, just like in Section 3, resort to the local picture instead. For $v \in M_K$, let $G_v = G_{\bar{K}_{\bar{v}}/K_v} \hookrightarrow G_{\bar{K}/K}$ be the decomposition group of $\bar{v}|v$ in $\bar{K}/K$. Then, similarly to $(*)$, we obtain a local Kummer sequence:

$$(**) \qquad 0 \longrightarrow \frac{E(K_v)}{mE(K_v)} \longrightarrow H^1(G_v, E[m]) \longrightarrow H^1(G_v, E)[m] \longrightarrow 0.$$

The difference now is that computing these groups is easier in the local setting by analytic methods such as Hensel's lemma, because the fields are complete. This is especially visible in the interpretation mentioned in Remark 5.1.

We can combine the sequence $(*)$ and the sequences $(**)$ for all v to get the commutative diagram:

$$\begin{tikzcd}
0 \arrow[r] & \frac{E(K)}{mE(K)} \arrow[r] \arrow[d] & H^1(G_{\bar{K}/K}, E[m]) \arrow[r] \arrow[d, "\varphi_1"] & H^1(G_{\bar{K}/K}, E(\bar{K}))[m] \arrow[r] \arrow[d, "\varphi_2"] & 0 \\
0 \arrow[r] & \prod_{v\in M_K} \frac{E(K_v)}{mE(K_v)} \arrow[r] & \prod_{v\in M_K} H^1(G_v, E[m]) \arrow[r] & \prod_{v\in M_K} H^1(G_v, E)[m] \arrow[r] & 0
\end{tikzcd}$$

with exact rows. In order to understand how the functions φ_1 and φ_2 work, we define the following groups.

Definition 5.2. The *m-Selmer group* of E/K is

$$S^{(m)}(E/K) = \ker\left(H^1(G_{\bar{K}/K}, E[m]) \to \prod_{v\in M_K} H^1(G_v, E)\right),$$

and the *Tate-Shafarevich group* is

$$Ш(E/K) = \ker\left(H^1(G_{\bar{K}/K}, E) \to \prod_{v\in M_K} H^1(G_v, E)\right).$$

Then the commutative diagram above shows that there is an exact sequence

$$0 \longrightarrow \frac{E(K)}{mE(K)} \longrightarrow S^{(m)}(E/K) \longrightarrow Ш(E/K)[m] \longrightarrow 0.$$

The group $S^{(m)}(E/K)$ is finite (the proof is essentially what we did to prove Weak Mordell–Weil, see [Sil09, Th. 4.2]) and possible to compute effectively ([Sil09, Rem. X.4.5]). However, the group $Ш(E/K)[m]$ is not easily accessible. It is not even known whether it is finite or not.

Conjecture 5.3. *Let E/K be an elliptic curve. Then $Ш(E/K)$ is finite.*

Remark 5.4. One way to interpret the group $Ш(E/K)$ is a way of measuring the *obstruction* to the local-global principle 3.5.

Remark 5.5. While this group is not known to be finite, Cassels showed that that exists an alternating bilinear pairing $Ш \times Ш \to \mathbb{Q}/\mathbb{Z}$. He concluded that, if the order is finite and there are no infinitely-divisible elements in Ш, it must be a perfect square. Surprisingly, this (see [Cas91, pp. 110]) played a role in the statement of the Birch–Swinnerton-Dyer Conjecture 5.9.

5.2. L-series and the conjecture.

One way to study arithmetic properties of elliptic curves is through their L-series. It is defined by an Euler product of *local factors*

$$L_{E/K}(s) = \prod_{v \in M_K^0} L_v(s). \tag{1}$$

If $v \in M_K^0$, the local factor L_v has a simple definition. We set $q_v = \# k_v$ and count how many points we expect to be on the elliptic curve $\widetilde{E}_v(k_v)$. The points are O and the solutions to the equation

$$y^2 = x^3 + ax + b.$$

For each $x \in k_v$, we expect around half of the numbers $x^3 + ax + b$ to be squares, and for each of the squares, we get two possible values of y. Thus, the heuristic estimate for the number of points in $\widetilde{E}_v(k_v)$ is $q_v + 1$. In fact, one can show that the absolute value of

$$a_v = q_v + 1 - \#\widetilde{E}_v(k_v)$$

is bounded by $2\sqrt{q_v}$, which is called the *Hasse bound* [Sil09, Th. V.1.]. We then define

$$L_v(s) = (1 - a_v q_v^{-s} + q_v^{-2s+1})^{-1}.$$

Remark 5.6. One may recognize this factor from the Euler product for the *zeta function* of an elliptic curve [Sil09, Chap. V.2].

If the curve has bad reduction at v, the definition depends on the kind of reduction (see Remark 3.8). We only say that

$$L_v(s) \in \left\{(1 + q_v^{-s})^{-1}, (1 - q_v^{-s})^{-1},\ 1\right\}$$

in this case.

It is easy to show using the Hasse bound above that the product (1) defines an analytic function for $\mathrm{Re}(s) > \frac{3}{2}$.

Conjecture 5.7. *There exists an analytic continuation of $L_{E/K}$ to the whole complex plane with a functional equation relating its values at s and $2 - s$.*

This conjecture is a theorem for elliptic curves with *complex mulitplication*, and for all elliptic curves over $\mathbb{Q}$.

Birch and Swinnerton-Dyer made the following conjecture relating the rank of the elliptic curve (known as the *algebraic rank*) and the order of vanishing of the L-function at $s = 1$ (known as the *analytic rank*).

Conjecture 5.8 (Birch–Swinnerton-Dyer, part I)**.** *Let $E/\mathbb{Q}$ be an elliptic curve. Then the order of vanishing of $L_{E/\mathbb{Q}}$ at $s = 1$ is the rank of $E(\mathbb{Q})$.*

The refined version of the conjecture relates the coefficient of the expansion of $L_{E/\mathbb{Q}}$ at $s = 1$ to certain algebraic properties of the elliptic curve. We first define for $v \in M_{\mathbb{Q}}$,

$$w_v = \begin{cases} \int\limits_{E(\mathbb{R})} \left|\frac{dx}{y}\right| & \text{for } v = \infty, \\ 1 & \text{if } E \text{ has good reduction at } v \in M_{\mathbb{Q}}^0, \\ \frac{\#E(\mathbb{Q}_p)}{\#E_0(\mathbb{Q}_p)} & \text{if } E \text{ has bad reduction at } v \in M_{\mathbb{Q}}^0. \end{cases}$$

Conjecture 5.9 (Birch–Swinnerton-Dyer, part II)**.** *Let $E/\mathbb{Q}$ be an elliptic curve and r be the rank of $E(\mathbb{Q})$. Then*

$$\lim_{s\to 1} \frac{L_E(s)}{(s-1)^r} = \frac{2^r R(E/\mathbb{Q})}{\#E_{\text{tors}}(\mathbb{Q})^2} \cdot \#Ш(E/\mathbb{Q}) \cdot \prod_{v\in M_{\mathbb{Q}}} w_v.$$

The factor of 2^r comes from the way we defined the absolute height; if we used $2\hat{h}$ instead, this number would disappear. The value

$$\frac{2^r R(E/\mathbb{Q})}{\#E_{\text{tors}}(\mathbb{Q})^2}$$

is hence a certain *normalized* way of writing the regulator.

The product

$$\prod_{v\in M_{\mathbb{Q}}} w_v$$

was called *fudge factors* by Birch and Swinnerton-Dyer and their presence was explained by Tate ([Sil09, pp. 451]).

The fact that Ш appears in this product is mysterious. According to Cassels [Cas91, pp. 110], Birch and Swinnerton-Dyer did not expect it to appear originally, but according to their data, their estimate was wrong by a factor of a perfect square. As mentioned in Remark 5.5, while the group Ш was not (and still is not) known to be finite, the existence of an alternating bilinear pairing on Ш leads to believe that the order of the group is a perfect square. (It is now known that this is not always the case.) Hence, it was conjectured that #Ш was the missing factor, and it has now been confirmed by new numerical data.

Appendix A. Algebraic number theory

We recall some fundamental results from algebraic number theory. For a detailed introduction, see [Neu99] or [Lan94].

In the proof of Theorem 3.10, we needed the following results.

Theorem A.1 ([Neu99, Th. I.6.3])**.** *Let I_K be the set of fractional ideals of K, and P_K be the subset of principal ideals. Then the class group I_K/P_K is finite.*

Lemma A.2. *Let K be a number field containing the mth roots of unity, $a \in K$. Let $b = \sqrt[m]{a}$ be a root of $x^m - a$, and $L = K(b)$. Then for a prime $\mathfrak{p}$ such that $\operatorname{ord}_{\mathfrak{p}}(m) = 0$*

$$L/K \text{ is unramified at } \mathfrak{p} \text{ if and only if } \operatorname{ord}_{\mathfrak{p}}(a) \equiv 0 \mod m.$$

Proof. Replacing the fields K by $K_{\mathfrak{p}}$ and L by $K_{\mathfrak{p}}(b)$, we may assume that K is a local field. To show the *if* implication, suppose that $\mathfrak{p}$ ramifies. Without changing the extension, we may replace a to assume that $\mathrm{ord}_{\mathfrak{p}}(a) \in \{0, 1, \dots, m-1\}$. The discriminant of the polynomial $x^m - a$ is $\pm m^m a^{m-1}$, and hence

$$\Delta_{L/K} | m^m a^{m-1},$$

because $\mathbb{Z}[b] \subseteq \mathcal{O}_L$. This shows that $\mathfrak{p} | m^m a^{m-1}$, and hence $\mathfrak{p}|a$. Therefore, $\mathrm{ord}_{\mathfrak{p}}(a) \neq 0$. For the *only if* implication, suppose L/K is unramified. Choose uniformizers π_L of L and π_K of K, and if v is the valuation on K corresponding to $\mathfrak{p}$, choose a valuation v' on L above v, normalized so that $v(\pi_L) = 1$. Then $v(\pi_K) = 1$, since L/K is unramififed. We may write

$$a = \pi_K^{n_1} u_1, \quad \text{for some } u_1 \in \mathcal{O}_K^*,$$
$$b = \pi_L^{n_2} u_2, \quad \text{for some } u_2 \in \mathcal{O}_L^*.$$

Then

$$\pi_K^{n_1} u_1 = b^m = \pi_L^{mn_2} u_2^m,$$

and taking v' of both sides, we get

$$n_1 = mn_2.$$

This shows $n_1 = \mathrm{ord}_{\mathfrak{p}}(a) \equiv 0 \mod m$. □

Theorem A.3 (Dirichlet S-units Theorem, [Lan94, Chap. V.1])**.** *Let K be a number field containing the roots of unity, and S be a finite set of places containing M_K^∞. The group R_S^* of units of the ring*

$$R_S = \{a \in K \mid v(a) \geq 0 \textit{ for any } v \in M_K \setminus S\}$$

is finitely-generated.

In Section 4 on heights, we needed the following results.

Proposition A.4 (Local Degree in Towers, [Lan94, Chap. II.1])**.** *Suppose $L/K/\mathbb{Q}$ is a tower of number fields and $v \in M_K$. Then*

$$\sum_{w|v} n_w = [L:K] n_v,$$

where the sum is over all $w \in M_L$ such that w restricted to K is v.

Theorem A.5 (Product Formula, [Lan94, Chap. V.1])**.** *If $x \in K^*$, then* $\prod_{v \in M_K} |x|_v^{n_v} = 1.$

References

[Cas91] J. W. S. Cassels, *Lectures on elliptic curves*, London Mathematical Society Student Texts, vol. 24, Cambridge University Press, Cambridge, 1991, doi:10.1017/CBO9781139172530, http://dx.doi.org.proxy.lib.umich.edu/10.1017/CBO9781139172530. MR 1144763

[Kir92] Frances Kirwan, *Complex algebraic curves*, London Mathematical Society Student Texts, vol. 23, Cambridge University Press, Cambridge, 1992, doi:10.1017/CBO9780511623929. MR 1159092 (93j:14025)

[Kob94] Neal Koblitz, *A course in number theory and cryptography*, second ed., Graduate Texts in Mathematics, vol. 114, Springer-Verlag, New York, 1994, doi:10.1007/978-1-4419-8592-7. MR 1302169 (95h:94023)

[Lan94] S. Lang, *Algebraic number theory*, Applied Mathematical Sciences, Springer, 1994, https://books.google.com/books?id=u5eGtA0YalgC.

[Neu99] Jürgen Neukirch, *Algebraic number theory*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 322, Springer-Verlag, Berlin, 1999, Translated from the 1992 German original and with a note by Norbert Schappacher, With a foreword by G. Harder, doi:10.1007/978-3-662-03983-0, http://dx.doi.org.proxy.lib.umich.edu/10.1007/978-3-662-03983-0. MR 1697859

[Sil09] Joseph H. Silverman, *The arithmetic of elliptic curves*, second ed., Graduate Texts in Mathematics, vol. 106, Springer, Dordrecht, 2009, doi:10.1007/978-0-387-09494-6. MR 2514094

[ST92] Joseph H. Silverman and John Tate, *Rational points on elliptic curves*, Undergraduate Texts in Mathematics, Springer-Verlag, New York, 1992, doi:10.1007/978-1-4757-4252-7. MR 1171452 (93g:11003)

[Wil06] Andrew Wiles, *The birch and swinnerton-dyer conjecture*, Official Problem Description from the Clay Mathematics Insitute, http://www.claymath.org/sites/default/files/birchswin.pdf.