Galois Cohomology and the Weak Mordell-Weil Theorem

David Grabovsky

July 26, 2017

Abstract

It was an ancient problem posed by the Greeks to find integer solutions to polynomial equations; or in more modern terminology, to find rational points on algebraic curves. To that end, we will study elliptic curves over the rational numbers and endeavor to prove a weak version of the Mordell-Weil Theorem: over a number field, an elliptic curve forms a finitely generated abelian group. Our weapons of choice will be the cohomology of Galois groups and the algebra of elliptic curves over the p-adic field. Time permitting, I will also mention some of the famous open problems facing modern mathematics, such as the conjecture of Birch and Swinnerton-Dyer and the question of computing the rank of an elliptic curve.

Contents

1 Introduction: Rational Points and Mordell-Weil

Early hunter-gatherers noticed a peculiar property of the bison they hunted: they always came in integer numbers. This observation later developed into Diaphontus's fascination with integral solutions to polynomial equations. The Greeks solved linear equations in x and y, and by the Rennaissance a complete theory existed for quadratic forms. Integer solutions to cubic (and higher-order) equations (e.g. $x^n + y^n = z^n$) proved much more difficult, and mathematicians hoped that they could solve them by doodling. It was by now understood that the zero-sets of polynomials described smooth curves and that integer solutons corresponded to rational points on those curves. The problem of finding rational points on *elliptic curves*, equations of the form $y^2 = ax^3 + bx + c$, caught the eye of number theorists, and they quickly began uncovering unexpected and beautiful algebraic structure.

For instance, they noticed that the points on an elliptic curve form an abelian group, and that considered over $\mathbb{C}$, this group is actually isomorphic to that of a complex torus. Yet over the rational numbers, elliptic curves eluded an easy description. Getting frustrated with $\mathbb{Q}$, people hoped that the problem would be easier modulo p: here the world is finite, and we only look at rational numbers "locally" (i.e. with respect to their degree of divisibility by p, and refusing to distinguish between numbers with the same remainder modulo p^n). If this p-adic analysis could be carried out for *every* p, then (hopefully) this should be enough to piece together a global solution for $\mathbb{Q}$. This *local-global* principle of Hasse will be our main line of attack, and quantifying its failure will lead us partway to a surprising result:

Theorem 1.1 (Mordell-Weil)**.** *For any number field*[1] *k and E an elliptic curve, the group $E(k)$ is finitely generated.*

The Mordell-Weil theorem is deep; it states that a finite set of points on an elliptic curve suffice to generate the entire thing. Its proof is difficult, so we will focus on a weaker result:

Theorem 1.2 (Weak Mordell-Weil)**.** *For any $n \in \mathbb{Z}$, the group $E(k)/nE(k)$ is finite.*

If an abelian group M is finitely generated, then its quotient M/nM is clearly finite for $n > 1$, but the converse is false in general: $\mathbb{Q} = n\mathbb{Q}$ is not finitely generated.

Thm. 1.1 follows from Thm. 1.2 by an argument by infinite descent originating with Fermat: we will not prove this, but we can give a cute demonstration for $n = 2$ and $k = \mathbb{Q}$.

Proposition 1.3. *Assume that $E(\mathbb{Q})/2E(\mathbb{Q})$ is finite. Then $E(\mathbb{Q})$ is finitely generated.*

Proof (sketch). Consider a point $P = [a : b : c] \in \mathbb{P}^2(\mathbb{Q})$ with $a, b, c \in \mathbb{Z}$ mutually coprime. Define the *height* of P to be $H(P) := \max(|a|, |b|, |c|)$; we can use H to define a canonical height function $h\colon E(\mathbb{Q}) \to \mathbb{R}$ on $E(\mathbb{Q})$. Let $P_1, ..., P_n$ be a set of representativesof $E(\mathbb{Q})/2E(\mathbb{Q})$: that is, any point $Q \in E(\mathbb{Q})$ can be written $Q = P_i + 2Q'$ for some representative P_i and some $Q' \in E(\mathbb{Q})$.

Next, claim that $h(Q') < h(Q)$ provided that $h(Q) > C$ for some fixed constant C. To prove this, argue by descent. If $h(Q') \geq C$, then we're done. Otherwise, $h(Q') > C$ and we repeat by writing $Q' = P_j + 2Q''$ and so on, inductively decomposing $Q^{(n)}$. We obtain

$$Q = P_i + 2Q' = P_i + 2(P_j + 2Q'') = \cdots . \tag{1.1}$$

[1] A *number field* is any field extension of $\mathbb{Q}$ of finite degree.

That is, the set $\{Q_1, ..., Q_m\} := \{Q \in E(\mathbb{Q}) \mid h(Q) \leq C\}$ is finite. We see that Q is a linear combination of the P_i's and the $Q^{(j)}$'s, so $\{P_i, Q^{(j)}\}$ generates $E(\mathbb{Q})$. ■

2 Cohomology of Galois Groups

We now train our sights on the Weak Mordell-Weil theorem, whose proof makes use of some of the machinery of Galois cohomology.

2.1 Modules and Cohomology

Definition 2.1. Let G be a group. A *G-module* is a pair (M, G) consisting of an abelian group M together with a G-action that respects its abelian structure, i.e. for all $\sigma \in G$ and all $m, m' \in M$, $\sigma(m + m') = \sigma m + \sigma m'$.

Note that $\sigma \in G$ defines an automorphism of M, so that the G-action defines a homomorphism $G \to \mathrm{Aut}(M)$, generalizing the notion of a representation and agreeing with the definition of a module over the group ring $\mathbb{Z}[G]$. For example, if L/K is a finite Galois extension with $G = \mathrm{Gal}(L/K)$, an elliptic curve $E(L)$ is naturally a G-module, with the action given by applying a field automorphism.

Definition 2.2. For a G-module M, the 0^{th} *cohomology group* is the fixed set by G:

$$H^0(G, M) := M^G = \{m \in M \mid \forall \sigma \in G,\ gm = m\}. \tag{2.1}$$

Definition 2.3. A *crossed homomorphism* is a map $f \colon G \to M$ such that

$$\forall \sigma, \tau \in G, \quad f(\sigma\tau) = f(\sigma) + \sigma f(\tau). \tag{2.2}$$

Definition 2.4. A crossed homomorphism is called *principal* if it is of the form

$$\forall \sigma \in G, \quad f(\sigma) = \sigma m - m. \tag{2.3}$$

Definition 2.5. For a G-module M, the *first cohomology group* is (roughly) the set of crossed homomorphisms that aren't principal. More precisely,

$$H^1(G, M) := \frac{\{\text{crossed homomorphisms}\}}{\{\text{principal crossed homos}\}}. \tag{2.4}$$

A few remarks are in order. First of all, any crossed homomorphism satisfies $f(1) = f(1 \cdot 1) = f(1) + f(1) \implies f(1) = 0$. Sums and differences of crossed (resp. principal) homomorphisms are crossed (resp. principal), and one easily verifies that the quotient $H^1(G, M)$ is an abelian group. There are higher cohomology groups $H^n(G, M)$, and for $n \geq 1$ they are all torsion (i.e. every element has finite order), but we won't need them.

Example 2.6. If G acts on M trivially, then $H^0(G, M) = M^G = M$. The crossed homomorphisms "uncross:" $f(\sigma\tau) = f(\sigma) + f(\tau)$, and principal crossed homomorphisms vanish: $f(\sigma) = \sigma m - m = 0$. Therefore $H^1(G, M) = \mathrm{Hom}(G, M)$.

Theorem 2.7. *Given a short exact sequence* $0 \longrightarrow M \longrightarrow N \longrightarrow P \longrightarrow 0$ *of G-modules, there is a corresponding (long) exact cohomology sequence*

$$\begin{array}{l} 0 \longrightarrow H^0(G,M) \longrightarrow H^0(G,N) \longrightarrow H^0(G,P) \xrightarrow{\ \delta\ } \\ \qquad H^1(G,M) \longrightarrow H^1(G,N) \longrightarrow H^1(G,P). \end{array} \tag{2.5}$$

Proof. First, recall that the above asks us to see M as a submodule of N (with f acting as the embedding), and $P \cong N/M$ as the factor module by the First Isomorphism Theorem.

We adorn our picture with inclusions from the H^0 groups to their corresponding modules:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & M & \xrightarrow{\ f\ } & N & \xrightarrow{\ g\ } & P & \longrightarrow & 0 \\ & & \uparrow & & \uparrow & & \uparrow & & \\ 0 & \longrightarrow & M^G & \xrightarrow{\ f_0\ } & N^G & \xrightarrow{\ g_0\ } & P^G & \xrightarrow{\ \delta\ } & \\ & & H^1(G,M) & \xrightarrow{\ f_1\ } & H^1(G,N) & \xrightarrow{\ g_1\ } & H^1(G,P). & & \end{array} \tag{2.6}$$

The maps f_0 and g_0 are the restrictions of f and g to M^G and N^G, respectively. To define f_1, let ϕ be the equivalence class of a crossed homomorphism $\tilde{\phi}\colon G \to M$; define $f_1(\phi)$ to be the equivalence class of $f \circ \tilde{\phi}$. Similarly, for $\psi := [\tilde{\psi}] \in H^1(G,N)$, define $g_1(\psi) := [g \circ \tilde{\psi}]$. That these maps are well-defined is easy to check if one has the patience.

Next, we need to construct the map δ; this step is tricky. Let $p \in P^G \subset P$, i.e. $\sigma p = p$ for all $\sigma \in G$. By surjectivity of g, $g(n) = p$ for some $n \in N$. Next, consider $\sigma n - n \in N$:

$$g(\sigma n - n) = \sigma g(n) - g(n) = p - p = 0, \tag{2.7}$$

so $\sigma n - n \in \ker(g) = \mathrm{im}(f)$. This means that we can consider $m := \sigma n - n$ as an element of M. Define the map $\delta_p\colon G \to M$ by $\delta_p(\sigma) = f^{-1}(\sigma n - n)$. This is a crossed homomorphism, but it need not be principal: indeed, one easily checks that

$$\delta_p(\sigma\tau) \text{ “}=\text{” } \sigma\tau n - n = \delta_p(\sigma) + \sigma\delta_p(\tau). \tag{2.8}$$

Finally, define $\delta\colon H^0(G,P) \to H^1(G,M)$ by $\delta(p) := [\delta_p]$, the class of δ_p up to principality. We need to check that δ is well-defined: indeed, another $n' \neq n$ mapping (by g) to p gives rise to the crossed homomorphism $\delta_{p'}\colon \sigma \mapsto \sigma n' - n'$, but $(\delta_p - \delta_{p'})(\sigma) = \sigma(n - n') - (n - n')$ is principal, so $\delta(p) = \delta(p')$ and δ is well-defined.

Having constructed the maps, we must show that the sequence is exact. For the most part, this is routine: exactness in f_0 and g_0 falls like manna from f and g by restriction, and exactness for g_1 is similar. Exactness for g_0 and f_1 involves the image and kernel of δ: $\ker(\delta)$ contains those p giving rise to principal δ_p, while $\mathrm{im}(\delta)$ is essentially the crossed homomorphisms arising as maps δ_p from some p fixed by G. Verifying that $\mathrm{im}(g_0) = \ker(\delta)$ and that $\mathrm{im}(\delta) = \ker(f_1)$ is tedious but standard, so we omit it. ∎

2.2 The Krull Topology

Let k be a perfect field, $\overline{k}$ its algebraic closure, and let $G := \mathrm{Gal}(\overline{k}/k) = G_k$ denote its absolute Galois group. We can dress G in the *Krull topology*: a subgroup is open if and only if it fixes a finite extension of k. In this topology, "more transcendental" elements are fixed by smaller subgroups of G; another way to say this is that the open subgroup $H < G$ form a neighborhood base for 1_G. G is a compact topological group, and open subgroups H are closed because their cosets, whose union form a complement of H, are open.

Definition 2.8. A G-module is *discrete* if the action $G \times M \xrightarrow{\mu} M$ is continuous with respect to the Krull topology on G and the discrete topology on M, with the product topology taken on $G \times M$).

The discrete topology on M severely restricts the form a continuous G-action can assume: since each $\{m\} \subset M$ is open, any two $\sigma, \tau \in H < G$ fixing the same field extension must also act the same way on M. It turns out that

Proposition 2.9. *Every principal crossed homomorphism $f\colon G \to M$ is continuous.*

Proof. Exercise with hints: first show that a crossed homomorphism $f\colon G \to M$ is continuous iff f is constant on the cosets of some open $H \lhd G$. Next, prove that G is compact, and use this to show that the orbit Gm of any $m \in M$ is finite. Then, show that the stabilizer H_m of any $m \in M$ is open in G: intersect the H_m of every $m \in Gm$, and obtain a normal open subgroup $J \lhd G$ contained in every H_m. Finally, assume that f is principal, and compute $f(J) = 0$ to find that f is constant on the cosets gJ and therefore continuous. ■

Definition 2.10. For a discrete G-module M, $H^1(G, M) := \dfrac{\{\text{cont. cr. homos}\}}{\{\text{pr. cr. homos}\}}$.

3 Way Too Many Exact Sequences

Having set up the necessary machinery, we can attack the weak Mordell-Weil theorem.

3.1 Surjectivity of $[n]$

Theorem 3.1. *For any elliptic curve E over an algebraically closed field k of characteristic 0 and any $n \in \mathbb{Z}$, the homomorphism $[n]\colon E(k) \to E(k)$ sending $P \mapsto nP$ is surjective.*

More generally, any morphism of smooth projective varieties (over any field) is either constant or surjective. For proof, consult your local (or global) algebraic geometry wizard.

Lemma 3.2. *The map $[n]\colon E(\mathbb{C}) \to E(\mathbb{C})$ is surjective.*

Proof. Let $\omega_1, \omega_2 \in \mathbb{C}$ be linearly independent over $\mathbb{R}$, and let $\Lambda = \omega_1\mathbb{Z} + \omega_2\mathbb{Z} \subset \mathbb{C}$ be a lattice. We proved that every elliptic curve $E(\mathbb{C})$ is isomorphic as a group to $E(\Lambda)$, and clearly $[n]$ maps $\Lambda \mapsto \Lambda$. Hence for any $P \in E(\mathbb{C})$, $nP \in E(\mathbb{C})$ as well, so $[n]$ is surjective. □

Lemma 3.3 (World's ugliest lemma)**.** *Let E be the curve $Y^2Z = X^3 + aXZ^2 + bZ^3$. Then for $P = [x : y : 1] \in E$ and any $n \in \mathbb{Z}$, $nP = [X\psi_n^4 - \psi_{n-1}\psi_n^2\psi_{n+1} : \frac{1}{2}\psi_{2n} : \psi_n^4]$, where*

$$\begin{aligned}
\psi_1 &= 1, \qquad \psi_2 = 2Y, \qquad \psi_3 = 3X^4 + 6aX^2 + 12bX - a^2, \\
\psi_4 &= 4Y(X^6 + 5aX^4 + 20bX^3 = 5a^2X^2 - 4abX - 8b^2 - a^3), \\
&\vdots \\
Y\psi_{2n} &= \psi_n(\psi_{n-1}^2\psi_{n+2} = \psi_{n+1}^2\psi_{n-2}), \\
\psi_{2n+1} &= \psi_n^3\psi_{n+2} - \psi_n^3\psi_{n-1}.
\end{aligned}$$

Proof. Laborious; by induction. See Cassels 1966 (7.2) and 1991 (p. 133). □

Lemma 3.4 (Hilbert's Nullstellensatz)**.** *Let K/k be an algebraically closed field extension; consider the polynomial ring $R = k[X_1, ..., X_n]$, and let $I \subset R$ be an ideal. Define $V(I) := \{\mathbf{x} = (x_1, ..., x_n) \in K^n \mid \forall f \in I,\ f(\mathbf{x}) = 0\}$ to be the common zero-set of all polynomials f in the ideal. Then I is proper (i.e. doesn't contain 1) if and only if $V(I)$ is nonempty.*

Proof. Omitted: see literally any book on algebraic geometry. □

Proof (theorem). Consider the subfield of k generated over $\mathbb{Q}$ by the coefficients of the polynomial defining E, and denote its algebraic closure in k by k_0. Then any subset of k_0 algebraically independent over $\mathbb{Q}$ is finite, i.e. k_0 has finite transcendence degree over $\mathbb{Q}$, and can therefore be embedded into $\mathbb{C}$. Moreover, $E(k)$ arises naturally as a curve $E_0(k_0)$, so WLOG we assume that $k \subset \mathbb{C}$. Next, recall from Lemma 3.3 that given a point $P \in E(k)$, nP has coordinates given by polynomials in X and Y; therefore, finding a point $Q \in E(k)$ such that $nQ = P$ (which is what we want!) amounts to solving these polynomials simultaneously. Now by Lemma 3.2, $[n]$ is surjective over $\mathbb{C}$, so these polynomials have a common solution in $\mathbb{C}$. They therefore generate a proper ideal in $k[X, Y]$, so by Hilbert's Nullstellensatz, they have a common solution in k. Hence, Q exists, so $[n]\colon E(k) \to E(k)$ is surjective. ■

3.2 The First Line of Attack

Thm. 3.1 furnishes the surjectivity of $[n]$, so we obtain the exact sequence

$$0 \longrightarrow E_n(\overline{\mathbb{Q}}) \xrightarrow{\iota} E(\overline{\mathbb{Q}}) \xrightarrow{n} E(\overline{\mathbb{Q}}) \longrightarrow 0, \tag{3.1}$$

where $E_n(k) := \ker([n])$ is the n-torsion subgroup of E consisting of all points on E of order n, and ι is just the inclusion. We now view $E(\mathbb{Q})$ as a $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$-module, and henceforth we abbreviate $H^i(k, E) := H^i(\mathrm{Gal}(\overline{k}/k), E(\overline{k}))$. From this perspective, Eq. 3.1 is a short exact sequence of G-modules, and we may apply the long exact cohomology sequence of Eq. 2.7:

$$\begin{array}{l}
0 \longrightarrow E_n(\mathbb{Q}) \xrightarrow{\iota} E(\mathbb{Q}) \xrightarrow{n} E(\mathbb{Q}) \xrightarrow{\delta} \\
\quad H^1(\mathbb{Q}, E_n) \xrightarrow{\tilde{\iota}} H^1(\mathbb{Q}, E) \xrightarrow{\tilde{n}} H^1(\mathbb{Q}, E).
\end{array} \tag{3.2}$$

Note that from the exactness of (3.2), we have

$$\begin{aligned} nE(\mathbb{Q}) &= \mathrm{im}([n]) = \ker(\delta); \\ H^1(\mathbb{Q}, E)_n &= \ker([\tilde{n}]) = \mathrm{im}(\iota). \end{aligned} \tag{3.3}$$

From here we extract the fundamental exact sequence

$$0 \longrightarrow E(\mathbb{Q})/nE(\mathbb{Q}) \xrightarrow{f} H^1(\mathbb{Q}, E_n) \xrightarrow{g} H^1(\mathbb{Q}, E)_n \longrightarrow 0, \tag{3.4}$$

Exactness follows from the above: the map $f := \delta \circ \pi \colon E(\mathbb{Q}/nE(\mathbb{Q}) \to H^1(\mathbb{Q}, E_n)$ induced on the factor group is injective by construction, and $g = \tilde{\iota}$ certainly surjects onto its image.

Now if $H^1(\mathbb{Q}, E_n)$ were finite, then $E(\mathbb{Q})/nE(\mathbb{Q})$ would also be finite, and we'd be finished with Weak Mordell-Weil. Unfortunately, this need not be the case. Instead, we will try to replace $H^1(\mathbb{Q}, E_n)$ with a subset that contains $\mathrm{im}(f)$, and which we can prove to be finite.

4 Elliptic Curves over $\mathbb{Q}_p$

Our next move is to consider elliptic curves over the p-adic field, which in some sense is easier because our analysis will be blind to anything besides the p-divisibility of numbers. Thus "local" algebra proves more tractible than "global" algebra, and if a local analysis succeeds for every p, we might be able to piece together a global solution.

4.1 The p-adic Field and the Localization Maps

The idea of p-adic numbers is to replace the usual notion of the distance between two numbers and replace it with a divisibility criterion: two numbers are declared close if a prime p divides them with the same remainder, closer if p^2 leaves the same remainder, and so on. The more divisible by p their difference is, the closer they are with respect to the p-adic metric.

Definition 4.1. Let p be prime; for $n \in \mathbb{Z}$, let $v_p(n)$ denote the largest number for which $p^{v_p(n)} \mid n$. For $r = \frac{a}{b} \in \mathbb{Q}$, define its *$p$-adic valuation* to be $v_p(r) := v_p(a) - v_p(b)$. Endow $\mathbb{Q}$ with the *p-adic norm*: $\|r\|_p := p^{-v_p(r)}$. Then the *p-adic field*, denoted $\mathbb{Q}_p$, is the metric completion of $\mathbb{Q}$ with respect to $\|\cdot\|_p$.

Now, let's choose an algebraic closure $\overline{\mathbb{Q}} \subset \mathbb{C}$ for $\mathbb{Q}$, and $\overline{\mathbb{Q}}_p$ for $\mathbb{Q}_p$. The embedding $\mathbb{Q} \hookrightarrow \mathbb{Q}_p$ extends to an embedding $\overline{\mathbb{Q}} \hookrightarrow \overline{\mathbb{Q}}_p$ (see below). Moreover, the action of $\mathrm{Gal}(\overline{\mathbb{Q}}_p/\mathbb{Q}_p)$ on $\overline{\mathbb{Q}} \subset \overline{\mathbb{Q}}_p$ defines a homomorphism $\psi \colon G_{\mathbb{Q}_p} \to G_{\mathbb{Q}}$ by rectriction of the Galois action.

$$\begin{array}{ccc} \overline{\mathbb{Q}} & \hookrightarrow & \overline{\mathbb{Q}}_p \\ \uparrow & & \uparrow \\ \mathbb{Q} & \hookrightarrow & \mathbb{Q}_p \end{array} \tag{4.1}$$

A crossed homomorphism $f \colon G_{\mathbb{Q}} \to E(\overline{\mathbb{Q}})$ therefore defines a crossed homomorphism $\tilde{f} \colon G_{\mathbb{Q}_p} \to E(\overline{\mathbb{Q}}_p)$ by the composition $\tilde{f} = f \circ \psi$. Taking $f \mapsto \tilde{f}$ yields the *localization homomorphism* $\phi \colon H^1(\mathbb{Q}, E) \to H^1(\mathbb{Q}_p, E)$. Checking that $\tilde{f}$ and ϕ are well-defined is easy, so we omit it. Note that this construction is independent of the embedding $\overline{\mathbb{Q}} \hookrightarrow \overline{\mathbb{Q}}_p$.

We now return to the exact sequence (3.4) and note that its construction is equally valid using $\mathbb{Q}_p$ instead of $\mathbb{Q}$. We therefore have the commutative diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & E(\mathbb{Q})/nE(\mathbb{Q}) & \xrightarrow{f} & H^1(\mathbb{Q}, E_n) & \xrightarrow{g} & H^1(\mathbb{Q}, E)_n & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & E(\mathbb{Q}_p)/nE(\mathbb{Q}_p) & \xrightarrow{f_p} & H^1(\mathbb{Q}_p, E_n) & \xrightarrow{g_p} & H^1(\mathbb{Q}_p, E)_n & \longrightarrow & 0, \end{array} \tag{4.2}$$

where the top and bottom rows are exact, and the vertical maps are embeddings. The first is an inclusion, while the latter two restrict ϕ on H^1 and its n-torsion. Commutativity is clear, as one may travel along the top row before "localizing" to the bottom row at any point.

4.2 The Selmer and Tate-Shafarevich Groups

We now reach a crucial argument. If some $\gamma \in H^1(\mathbb{Q}, E_n)$ comes from the class of an element of $E(\mathbb{Q})$, then its image $\gamma_p \in H^1(\mathbb{Q}_p, E_n)$ arises from an element of $E(\mathbb{Q}_p)$. We want to quantify (a) all those γ whose local versions γ_p come from $E(\mathbb{Q}_p)$, and (b) the extent to which the local-global principle fails, i.e. all those $\gamma \in H^1(\mathbb{Q}, E)$ that vanish locally.

Definition 4.2. The *n-Selmer group* is defined by

$$\begin{aligned} S^{(n)}(E/\mathbb{Q}) &:= \{\gamma \in H^1(\mathbb{Q}, E_n) \mid \forall p, \gamma_p \text{ comes from } E(\mathbb{Q}_p)\} \\ &= \ker\left(H^1(\mathbb{Q}, E_n) \longrightarrow \prod_{p \text{ prime}} H^1(\mathbb{Q}_p, E) \right). \end{aligned} \tag{4.3}$$

Definition 4.3. The *Tate-Shafarevich*[2] *group* is defined by

$$Ш(E/\mathbb{Q}) := \ker\left(H^1(\mathbb{Q}, E) \longrightarrow \prod_{p \text{ prime}} H^1(\mathbb{Q}_p, E) \right). \tag{4.4}$$

We are now ready to pin down E/nE. The following lemma is the nail in the coffin:

Lemma 4.4. *For any chain of modules* $A \xrightarrow{\alpha} B \xrightarrow{\beta} C$ *, there is an exact sequence*

$$0 \longrightarrow \ker(\alpha) \longrightarrow \ker(\beta\alpha) \longrightarrow \ker(\beta) \longrightarrow \operatorname{coker}(\alpha) \longrightarrow \operatorname{coker}(\beta\alpha) \longrightarrow \operatorname{coker}(\beta) \longrightarrow 0. \tag{4.5}$$

[2]The group $H^1(k, E)$ is sometimes called the Weil-Châtelet group and is often written $WC(E/k)$. Cassels admits responsibility for denoting the Tate-Shafarevich group bythe Cyrillic Ш. He comments: "This is the author's most lasting contribution to the subject. The original notation was TS, which, Tate tells me, was intended to continue the lavatorial allusion of WC. The Americanism "tough shit" indicates the part that is difficult to eliminate" (1990, p. 109). Although some authors write "Shafarevich-Tate" for the group's name, the original is the correct alphabetical order in Cyrillic.

All of the maps are natural, and we decorate the diagram to make exactness obvious: note that $\beta \circ \alpha = 0$ for all chains, and that $[\beta]\colon B/V \to C$ is defined iff $\beta(V) = \{0\}$.

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \overset{\subset A}{\ker(\alpha)} & \xrightarrow{\iota} & \overset{=A}{\ker(\beta\circ\alpha)} & \xrightarrow{\alpha} & \overset{\subset B}{\ker(\beta)} & \xrightarrow{\pi} & \\ & & \underset{=B/\mathrm{im}(\alpha)}{\mathrm{coker}(\alpha)} & \xrightarrow{[\beta]} & \underset{=C}{\mathrm{coker}(\beta\circ\alpha)} & \xrightarrow{\pi} & \underset{=C/\mathrm{im}(\beta)}{\mathrm{coker}(\beta)} & \longrightarrow & 0. \end{array} \tag{4.6}$$

Finally, we consider the module homomorphisms

$$H^1(\mathbb{Q}, E_n) \xrightarrow{\alpha} H^1(\mathbb{Q}, E)_n \xrightarrow{\beta} \prod_{p \text{ prime}} H^1(\mathbb{Q}_p, E)_n,$$

obtaining the kernel-cokernel exact sequence of the lemma. Note that the map $\beta \circ \alpha$ sends $\gamma \in H^1(\mathbb{Q}, E_n)$ to the product of all $\gamma_p \in H^1(\mathbb{Q}_p, E)_n$: at last, the kernel-cokernel sequence yields the fundamental exact sequence

$$0 \longrightarrow E(\mathbb{Q})/nE(\mathbb{Q}) \longrightarrow S^{(n)}(E/\mathbb{Q}) \longrightarrow Ш(E/\mathbb{Q})_n \longrightarrow 0. \tag{4.7}$$

As it turns out, $S^{(n)}(E/\mathbb{Q})$ is finite. Proofs may be found in Silverman (pp. 190-196) and Milne (IV.3), and involve some algebraic number theory. On the other hand, $Ш(E/\mathbb{Q})$ is conjectured (and widely believed) to be finite, but this is not known.

Theorem 4.5. *The weak Mordell-Weil theorem is true.*

Proof. As shown above, $E(\mathbb{Q})/nE(\mathbb{Q})$ injects into $S^{(n)}(E/\mathbb{Q})$, which is finite; therefore, $E(\mathbb{Q})/nE(\mathbb{Q})$ is finite. With some algebraic number theory, $\mathbb{Q}$ can be generalized to any number field k, and p is replaced by the more general notion of a *place*. In any case, the core of the argument is the same, so excepting the proof of the Selmer group's finiteness, we have the weak Mordell-Weil theorem. ■

Corollary 4.6. *The Mordell-Weil group takes the form $E(k) \cong E(k)_{tors} \oplus \mathbb{Z}^r$, where $r \geq 0$ is called the* rank *of the elliptic curve, and gives the number of free generators for E.*

There is no proven algorithm for computing the rank of an elliptic curve. However, the conjectural finiteness of $Ш(E/\mathbb{Q})$ would yield such a procedure; this is described in more detail by Milne (IV.5). It is widely believed that there is no maximum rank for elliptic curves, and Noam Elkies has shown that elliptic curves of rank at least 28 exist. It is conjectured that the average rank of all elliptic curves should be 1/2, and an upper bound of 2 (assuming both the Birch and Swinnerton-Dyer conjecture and the Generalized Riemann Hypothesis) was recently beaten by 1.17 (without assuming either conjecture).