Big Book™

by
Rawley Harrison & Aaryan Sukhadia

Stanford University

Started 2025

Contents

0 Introduction and Outlines **6**

0.1 Motivations and Goals of this Project . . . 6
0.2 Format . . . 7
0.3 Conventions and Assumed Background . . . 8
0.4 Notational Index . . . 11

I Abelian Varieties and Tate's Conjecture **12**

1 Abelian Varieties over a Field **13**

1.1 Definitions, First Properties, and Line Bundles . . . 14
1.1.1 Rigidity and Commutativity . . . 14
1.1.2 The Picard Scheme . . . 15
1.1.3 Line Bundles on Abelian Varieties . . . 17
1.1.4 Torsion . . . 24
1.2 The Dual Abelian Variety . . . 31
1.2.1 The Poincaré Bundle . . . 31
1.2.2 Smoothness of $\mathrm{Pic}^0_{A/k}$ . . . 32
1.2.3 A Characterization of Line Bundles in $\mathrm{Pic}^0_{A/k}$ . . . 34
1.2.4 Dual Morphisms . . . 37
1.2.5 The Néron-Severi Group . . . 38
1.2.6 More on $\mathscr{P}_A$ and the Duality Hypothesis . . . 40
1.3 Applications of Duality . . . 50
1.3.1 Cartier Duality and The Weil Pairing . . . 50
1.3.2 The p-rank of an Abelian Variety . . . 56
1.3.3 Cohomology of Line Bundles . . . 60

2 Tate's Conjecture for Finite Fields & Endomorphisms of Abelian Varieties **70**

2.1 Algebraic Structure of $\mathrm{Hom}(A, B)$ . . . 71
2.1.1 Étale Coverings over $\overline{k}$ . . . 71
2.1.2 Structure of $\mathrm{Hom}(A, B)$ . . . 76
2.1.3 Riemann Forms . . . 84
2.1.4 The Rosati Involution . . . 89
2.1.5 The Structure of $\mathrm{End}^0 A$ for Simple A . . . 92

2.1.6 The Riemann Hypothesis . . . 94
2.1.7 The Structure of $\mathrm{NS}^0(A)$ . . . 97
2.1.8 CM Abelian Varieties of characteristic 0 . . . 98
2.1.9 Revisiting Elliptic Curves in Characteristic $p > 0$ . . . 98
2.2 Tate's Isogeny Theorem . . . 103
2.2.1 The Proof of Tate's Conjecture for Finite Fields . . . 103
2.2.2 Applications and Tate's Isogeny Theorem . . . 109

3 Abelian Schemes 114

4 p-divisible Groups 115

II Moduli Spaces 116

5 Algebraic Spaces and Stacks: A Primer 117
5.1 Moduli Functors . . . 118
5.2 Sheaves on Sites . . . 120
5.3 Definitions and Key Properties . . . 120

6 Moduli of Stable Curves 121

7 Moduli of Principally Polarized Abelian Varieties 122

III Some Key Intermediary Constructions and Theories 123

8 Néron Models 124
8.1 Models and their Uses . . . 126
8.1.1 Motivation and Definitions . . . 126
8.1.2 Elementary Properties . . . 130
8.1.3 Local Existence Prelude: Boundedness . . . 134
8.1.4 (An Outline of) Local Existence . . . 138
8.1.5 Local Existence Postlude: Global Existence . . . 142
8.2 Weak Néron Models (Construction Steps I-II) . . . 144
8.2.1 Separatedness . . . 144
8.2.2 Smoothening: Theorem Statement . . . 149
8.2.3 Defect of Smoothness . . . 154
8.2.4 S-Rational Maps and the Weak Néron Mapping Property . . . 160
8.3 Birational Groups (Construction Steps III-IV) . . . 167
8.3.1 Invariant Differential Forms . . . 167
8.3.2 Local Néron Models Exist (assuming Weil's Theorem) . . . 171
8.3.3 Solving Birational Group Laws (proving Weil's Theorem) . . . 177
8.4 Semistable Reduction of Abelian Varieties . . . 190

9 The Theory of Heights 191
9.1 Reaching new Heights . . . 191
9.1.1 Review of Absolute Values and Norms . . . 191

9.1.2 Heights on Projective Space . . . 193
9.1.3 Heights on Projective Varieties . . . 197
9.1.4 Canonical Heights . . . 203
9.1.5 Heights on Abelian Varieties . . . 206
9.2 The Mordell-Weil Theorem . . . 216
9.2.1 Mordell-Weil from Weak Mordell-Weil . . . 216
9.2.2 Weak Mordell-Weil . . . 218
9.2.3 Bounds on the Rank of $A(K)$ . . . 222

IV Appendices 223

A Reviewing "Nice" Properties of Schemes and Morphisms 225
A.1 Smoothness . . . 226
A.1.1 Flatness . . . 226
A.1.2 Differential Forms . . . 227
A.1.3 Unramified, Smooth and Étale Morphisms . . . 231
A.2 Henselian Schemes and Rings . . . 234
A.3 Blowups . . . 235
A.3.1 Basic construction . . . 235
A.3.2 Computing blow-ups . . . 237
A.3.3 Strict transform . . . 241
A.3.4 Functoriality and base change for blow-ups . . . 243

B Galois Cohomology 246
B.1 Primer – Profinite Group Cohomology . . . 247
B.1.1 Defininining Cohomology . . . 247
B.1.2 Cohomological Dimension . . . 252
B.2 Commutative Galois Cohomology . . . 256
B.2.1 First Definitions and Examples . . . 256
B.2.2 Brauer Groups . . . 257
B.2.3 Fields of dimension ≤ 1 . . . 257
B.2.4 Cohomology of p-adic Fields . . . 260
B.3 Non-abelian Galois Cohomology . . . 261

C Miscellaneous Useful Theories 262
C.1 Descent . . . 262
C.1.1 Definitions and Main Theorem . . . 262
C.1.2 The Theorem of the Square . . . 267
C.1.3 Torsors . . . 267
C.2 Valuative Spaces . . . 268
C.2.1 Constructibly Supported Rings . . . 268
C.2.2 Sheaves of Φ-Rings . . . 271
C.3 Some Useful Algebraic Number Theory . . . 273
C.3.1 Ramification . . . 273
C.3.2 Kummer Theory . . . 273

References **275**

Chapter 0

Introduction and Outlines

Section 0.1

Motivations and Goals of this Project

After a talk given by Professor Ravi Vakil in Fall 2022 in which he closes by discussing the statement (and absurdity) of the Mordell Conjecture, which says the following.

Theorem 0.1.0.1. Let K be a number field, and let X be an algebraic curve over K of genus $g \geq 2$. Then $|X(K)| < \infty$, i.e. there are only finitely many K-rational points.

I (Aaryan) went up to him and asked what it would take to build up to Falting's proof of the conjecture. Ravi's response was something along the following:

> *"First probably the whole grad sequence in algebraic aeometry here at Stanford. Then maybe a year-long course in some harder algebraic geometry. And then maybe another year of even harder algebraic geometry."*

I (Aaryan) decided to take that on. Together with Rawley, we decided to build up all the background foundational knowledge to understand Faltings' proof of 0.1.0.1. This book is a collaborative attempt at documenting and expositing the shared processes and learnings of the authors towards that goal.

There are certainly refinments of Faltings' original paper that simplify a lot of the arguments [cite some here], but for the authors the end goal isn't simply a proof, but rather to build up as much background in algebraic geometry as possible in the process of understanding that proof. Faltings' original paper is simply scaffolding around which the authors chose courses, directed readings and self-propelled readings in various domains of algebraic geometry. It is possible that in future drafts we will build up and exposit even more background to realy alternate proofs (e.g. that of Venkatesh making use of p-adic Hodge theory [cite]) or generalizations (e.g [find some later]).

Section 0.2

Format

One of the benefits to working in the modern age is the lack of constraints due to printing considerations. As such, we will use a variety of colors and boxes to make this text more readable and easy to identify important results. This template is from Aaryan Sukhadia. There will be three main types of colored boxes used frequently in this template.

Definition 0.2.0.1. A bright, almost-blaring red-backgrounded, red-striped box to house a vital Definition.

Example 0.2.0.2. An anber-striped environment to indicate an important or prototypical Example of a concept is contained in its confines.

Theorem 0.2.0.3. A neutrally-toned gray-shaded box to house a Result, whether that be a theorem (such as this particular box), a proposition, etc.

Big, major theorems throughout this text are in more distinct boxes, as in the following.

Theorem 0.2.0.4. A BIG, MAJOR result.

Section 0.3

Conventions and Assumed Background

Writers of mathematics are always challenged with balancing mathematical acessibility with the limits of time and space. As a litmus test we suggest the reader go through the following subsection on Conventions and Notations. The goal of this text is that it should be broadly understandable to a reader comfortable with some dense subset of the conventions described below.

Miscellaneous Notation Conventions

- $\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}$ and $\mathbb{C}$ will denote the sets of Naturals, Integers, Rationals, Reals and Complex Numbers respectively. Naturals include 0.
- S^1 denotes the circle, i.e the group of norm-1 complex numbers.
- $\mathbb{F}_q$ denotes the (unique, up to isomorphism!) finite field with q elements. In general by $\text{char}\,(K)$ we denote the **characteristic** of a field K.
- When we have an infinite product or a vector with infinitely many components we say **"almost everywhere"** to mean "in all but finitely many components". If the components of the infinite product are indexed by $i \in I$ then we write $(x_i)_i$ to mean the "vector" of a single element in the infinite product, and x_i to mean the i-th coordinate.
- For a subset A of a set X, we denote by A^c, or $X \setminus A$, or $X - A$ the complement of A in X.
- Given a function f defined on a space X and a subset $Y \subseteq X$, we write $f|_Y$ for the restriction of f to Y.

Abstract Algebra

- Given subgroups H, K of a commutative group G with operation denoted by addition, we denote by $H + K$ the subgroup of all elements of the form $h + k$ with $h \in H$ and $k \in K$.
- All rings are assumed to be unital unless stated otherwise; $R^\times$ will denote the set of invertible elements of a ring R.
- If H, G are groups we use $H \leq G$, $H \trianglelefteq G$, to denote "H is a subgroup of the group G", "H is a normal subgroup of the group G" respectively. If R is a ring we write $I \trianglelefteq R$ to denote "I is an ideal of the ring R".
- If $r_1, r_2, \ldots$ is a set of elements in a group G (resp. a ring R), we denote by $\langle r_1, r_2, \ldots \rangle$ the subgroup (resp. ideal) generated by the r_i.
- The symbol $\mathfrak{p}$ *always* refers to a prime ideal. When necessary we add subscripts to indicate the relevant field we are working in. When working in a Dedekind domain we denote by $N\mathfrak{a}$ the **norm** of a nonzero ideal $\mathfrak{a}$, i.e the cardinality of $R/\mathfrak{a}$.

- When working a **discrete valuation ring** (a DVR), unless otherwise stated, we denote by π its uniformizer, by $\mathfrak{m}$ its unique maximal ideal and by k its residue field.

Number Theory

- By a **local field** we mean a (non-trivially!) valued field whose additive group is locally compact with respect to the topology induced by the valuation. Note this convention includes $\mathbb{R}$ and $\mathbb{C}$, though they are sometimes excluded from the definition in other texts.

- We say a local field F is **Archimedean** if it is isomorphic to $\mathbb{R}$ or $\mathbb{C}$, and **non-Archimedean** otherwise. For F non-Archimedean with absolute value $|\cdot|_F$ we denote by $\mathcal{O}_F$ its **valuation ring** (i.e the ring of elements with absolute value ≤ 1).

- By a **global field** K we mean a finite extension of $\mathbb{Q}$ (in which case we call K a **number field**) or of $\mathbb{F}_p(t)$ (in which case we call K a **function field**). For any (non-trivial!) valuation v on K we denote by K_v the **completion** of K with respect to v; this is always a local field. When K_v is non-Archimedean we write $\mathcal{O}_v$ to denote the valuation ring of K_v. Note this convention particularly in the case of $K = \mathbb{Q}$; when we write $\mathbb{Z}_p$ we are *always* referring to the p-adic integers! We assume the reader is familiar with the standard fact that every local field arises as the completion of a global field.

- We denote by $\mathscr{P}_K$ the set of all **places** (i.e topological equivalence classes of absolute values) of a global field K. We may speak of a place v being Archimedean, non-Archimededan, real, complex, if K_v is Archimedean, non-Archimededan, real, complex, respectively. In particular when $\mathrm{char}\,(K) > 0$ it has no Archimedean places.

- We recall that if K'/K is an algebraic extension of fields than any nontrivial absolute value $|\cdot|_K$ on K extends uniquely to K'. Thus if $v \in \mathscr{P}_K$ and $w \in \mathscr{P}_{K'}$ are such that there exists an absolute value on K corresponding to v whose unique extension corresponds to $w \in \mathscr{P}_{K'}$, then we say w **lies over** v, denoted $w \mid v$.

- We say a place v of a global field K is **infinite** if it is Archimedean, and **finite** otherwise. In particular, all places of a global function field are finite.

Algebraic Geometry

- By an **affine scheme** we mean the spectrum of a commutative ring R along with the data of its canonically associated structure sheaf. A **scheme** is any locally ringed space $(X, \mathcal{O}_X)$ – we will always use $\mathcal{O}_X$ to denote the structure sheaf of a ringed space X – that is locally isomorphic to an affine scheme.

- If $f : X \to S$ is a morphism of schemes we say X is an S**-scheme** with **structure morphism** f. If $S' \to S$ is another morphism of schemes we denote by $X_{S'}$ the **base change** to S', i.e. the fiber product $X \times_S S'$. In the case that $S' = \mathrm{Spec}(R)$ is an affine scheme we abuse notation write X_R; this should cause no confusion.

- If X and T are schemes by $X(T)$ we denote the T**-valued points** of X, i.e. the set of scheme morphisms $T \to X$.

- An **algebraic variety** X is an integral, separated, finite-type affine scheme over a field k. An **algebraic curve** is a 1-dimensional algebraic variety. The **function field** $k(X)$ is the stalk of the (unique!) generic point of X, or equivalently the field of fractions of $\mathcal{O}_X(U)$ for any non-empty open U.
- By a **line bundle** $\mathscr{L}$ on a scheme X we mean an invertible sheaf, i.e. there exists some other $\mathcal{O}_X$-module $\mathscr{M}$ such that $\mathscr{L} \otimes_{\mathcal{O}_X} \mathscr{M} \cong \mathcal{O}_X$. We recall this condition amounts to $\mathscr{L}$ being locally free of rank one (hence the term line bundle).
- [divisor-line bundle correspondence. Talk about Cartier vs Weil]
-

Section 0.4

Notational Index

Part I

Abelian Varieties and Tate's Conjecture

Chapter 1

Abelian Varieties over a Field

"*To write a thorough treatise on abelian varieties would be a formidable job...However, [we] have attempted to keep the discussion as simple as was compatible with the amount of material to be covered.*"

Introduction to *Abelian Varieties*
David Mumford

Abelian varieties are one of the most ubiquitous objects in algebraic geometry and number theory. Much like many of the most beautiful objects in algebra, they are defined simply and yet give rise to such rigid, nuanced, and deep structures. Beginning with elliptic curves, the one-dimensional case, study of abelian varieties goes as far back as Diophantus in the 3rd century AD. Abelian varieties were first generalized to higher dimensions via the advent of the Abel-Jacobi map by Abel in the early 19th century. This map gave the first examples of abelian varieties as Jacobians, the kernel of these maps. The modern basic theory was developed by Weil in the 1940s, in which the term "abelian variety" was coined. However, the modern formulation of abelian varieties really only took off after their modern reformulation in terms of schemes, which allowed for the theory to be developed in the case of characteristic $p > 0$.

This chapter largely follows [Mum74], but with adaptations to the case of non-algebraically-closed fields assisted by the notes of Brian Conrad [Con15]. Other good resources on this topic include [Mil08] and [CS86]. Additionally, for the sake of examples, throughout this text we will occasionally refer to the analytic theory over $\mathbb{C}$, which interested readers can read about in [Mum74, Ch. 1], [CS86, Ch. 4], and [BL04a].

This first chapter builds up all the basic theory covered in Chapters II and III of [Mum74], but over general fields, which will prepare us to prove theorems such as Tate's Isogeny Theorem, the Riemann Hypothesis for abelian varieties, and Tate's Conjecture for finite fields.

Section 1.1

Definitions, First Properties, and Line Bundles

We begin this chapter with a discussion of the basic theory of abelian varieties.

Definition 1.1.0.1. An **abelian variety** over a field k is a proper, smooth, connected group scheme over k.

§1.1.1 Rigidity and Commutativity

Our first goal is to prove that abelian varieties are actually abelian (commutative). To do this we will need the following result which will be very useful going forward.

Theorem 1.1.1.1 (Rigidity). Let X, Y be geometrically integral schemes of finite type over k, and Z a separated k-scheme. Furthermore let X be proper. Let $f : X \times Y \to Z$ be a k-morphism such that for some algebraically closed K/k, there exists $y_0 \in Y(K)$ such that the restriction $f|_{X_K \times \{y_0\}} : X_K \to Z_K$ is a constant map to some $z_0 \in Z(K)$. Then f is independent of X, i.e., there exists a unique k-morphism $g : Y \to Z$ such that the following diagram commutes:

$$\begin{array}{ccc} X \times Y & & \\ \downarrow & \searrow^{f} & \\ Y & \xrightarrow[g]{\exists!} & Z \end{array}$$

In other words, f being constant on one fiber of $X \times Y$ over Y, properness of X forces it to be constant over *every* such fiber.

Proof. Uniqueness follows since Y is reduced and $X \times Y \to Y$ is surjective. We next argue by Galois descent that if such a g exists for $f_{k'}$ over some finite Galois extension k'/k, then g is defined over k. This is clear since for any $\sigma \in \mathrm{Gal}(k'/k)$, $\sigma^* f_{k'} = f_{k'}$, so by uniqueness, $\sigma^* g = g$ which by Galois descent implies that g is defined over k.

We will now prove existence. Firstly, for some finite Galois extension k'/k, there exists a rational point $x_0 \in X(k')$ (see Lemma 1.1.1.2 below). Thus, by the descent argument above, we may assume $X(k) \neq \emptyset$. Then for some $x_0 \in X(k)$, define $g : Y \to Z$ as the composition $Y = \mathrm{Spec}(k) \times Y \xrightarrow{x_0 \times \mathrm{id}} X \times Y \xrightarrow{f} Z$, i.e., on rational points define $g(y) = f(x_0, y)$. We want to show that $f = g \circ \mathrm{pr}_2$. To check this, we may extend the base field and assume that that $y_0 \in Y(k)$ and k is algebraically closed. By assumption, $f(X \times \{y_0\})$ is the constant map to z_0. Let U be an affine neighborhood of z_0 in Z, $F = Z - U$, and $G = \mathrm{pr}_2(f^{-1}(F))$. Then G is closed since X is proper and hence pr_2 is a closed morphism. Further, if we let $V = Y - G$, it is a nonempty open subset of Y (as it contains y_0). For all $y \in V(k)$, we have that $X \times \{y\} \to U$ is constant since X is proper and U is affine. Thus, on $(X \times V)(k)$, we have that $f(x, y) = f(x_0, y) = g(y)$. Since Z is separated, the scheme $(f, g \circ \mathrm{pr}_2)^{-1}(\Delta_{Z/k}) \subset X \times Y$ classifying equality for f and $g \circ \mathrm{pr}_2$ is a closed, reduced subscheme of the scheme $X \times Y$. But this subscheme contains the subset $(X \times V)(k)$ which is Zariski-dense since $k = \overline{k}$ and Y is irreducible. Hence the scheme of equality is $X \times Y$, i.e., $f = g \circ \mathrm{pr}_2$. □

To complete the proof, we need the claim:

LEMMA 1.1.1.2. If X is a geometrically reduced scheme of finite type over $k = k_s$ then $X(k) \neq \emptyset$.

Proof. Consider $x \in X$ a closed point. Then since X is finite type over k, we must have that $\kappa(x)/k$ is finite and thus purely inseparable (as $k = k_s$. If this is nontrivial, then $\kappa(x) \otimes_k \overline{k}$ is not reduced (since $\kappa(x) \supset k[y]/(y^p - t)$ for some t without a pth root in k and thus $\kappa(x) \otimes_k \overline{k} \supset \overline{k}[y]/(y - t^{1/p})^p$ which contains nilpotents). Thus, $\kappa(x)/k$ must be trivial and so $x \in X(k)$. □

A very important corollary of this, which will come up many times throughout the proofs in this paper, is:

COROLLARY 1.1.1.3. Let A, A' be abelian varieties over k. Any morphism $f : A \to A'$ is of the form $T_{a'} \circ h$ for $T_{a'} : A' \to A'$ translation by some $a' \in A'(k)$ and $h : A \to A'$ a homomorphism, i.e., any morphism between abelian varieties differs from a homomorphism by translation.

Proof. We may replace f with $T_{f(e)^{-1}} \circ f$ so that $f(e) = e'$ (the identity elements in A, A' respectively). We claim that this map is a homomorphism.

Consider the map $h : A \times A \to A'$ given by

$$(a_1, a_2) \mapsto f(a_2 a_2) f(a_1)^{-1} f(a_1)^{-1}$$

for S-points $a_1, a_2 \in A(S)$. But we have $f(A \times \{e\}) = f(\{e\} \times A) = e'$, hence by rigidity, f factors through both the first and second projection. But this implies that $h(A \times A) = e'$, i.e., f is a homomorphism. □

COROLLARY 1.1.1.4. Abelian varieties are commutative.

Proof. Apply the above corollary to the inverse morphism $i : A \to A$. Since $i(e) = e$, i is a homomorphism and so we have $(xy)^{-1} = x^{-1}y^{-1} = y^{-1}x^{-1}$. □

§1.1.2 The Picard Scheme

To study line bundles on abelian varieties, we will introduce the Picard scheme. For this, we follow the construction in [Fan05, Ch. 9]. The ultimate goal of this is to construct the dual abelian variety, which we will discuss in the next section. Note that in [Mum74], the Picard group is not discussed scheme theoretically, and the dual abelian variety is constructed independently. We choose to instead construct it using the Picard scheme for two reasons. Firstly, it allows a construction of the dual independent of a choice of line bundle. Secondly, it allows us to more easily adapt to our case of non-algebraically-closed fields.

And secondly, this will allow us to easily generalize to abelian schemes in Chapter 3. final draft

DEFINITION 1.1.2.1. The **Picard group** $\mathrm{Pic}(X)$ of a scheme X is the group of isomorphism classes of invertible sheaves (i.e. line bundles) on X with the group action given by the tensor product.

There are multiple notions of the Picard functor given in [Fan05], which when representable are all represented by the same scheme. We call this scheme the **Picard scheme**. We will use the "strongest" definition (in the sense that if it is representable, then every other Picard functor is representable), which will turn out to be representable for our purposes.

DEFINITION 1.1.2.2. For any S-scheme X, we define the **Picard functor** by

$$\underline{\mathrm{Pic}}_{X/S}(T) := \mathrm{Pic}(X_T) \Big/ f_T^*(\mathrm{Pic}(T))$$

where $f_T : X_T \to T$ is the base change of the structure morphism.

REMARK 1.1.2.3. Note that for each equivalence class $[\mathscr{L}] \in \mathrm{Pic}(X_T)/f_T^*(\mathrm{Pic}(T))$, if there exists a section, i.e., a T-point $e_T : T \to X_T$ (e.g. the identity element of a group scheme),

$$\begin{array}{c} X_T \\ {\scriptstyle e_T}\uparrow \big\downarrow {\scriptstyle f_T} \\ T \end{array}$$

then there is a unique representative of $[\mathscr{L}]$ such that there exists a *trivialization* $e_T^*(\mathscr{L}) \simeq \mathscr{O}_T$. This is because if $e_T^*(\mathscr{L}) = \mathscr{M} \in \mathrm{Pic}(T)$, then $e_T^*(\mathscr{L} \otimes f_T^*\mathscr{M}^{-1}) = e_T^*\mathscr{L} \otimes \mathscr{M}^{-1} = \mathscr{O}_T$ (since $f_T \circ e_T = \mathrm{id}$). Thus, the unique element with trivialization in $[\mathscr{L}]$ is $\mathscr{L} \otimes f_T^*\mathscr{M}$. We will occasionally use this characterization as pinning down a particular representative will be helpful.

THEOREM 1.1.2.4. Let X be a geometrically reduced, geometrically connected, proper k-scheme. If $X(k) \neq \emptyset$, then $\underline{\mathrm{Pic}}_{X/k}$ is representable by a locally finite type k-scheme $\mathrm{Pic}_{X/k}$.

Proof. See [Fan05, Thm. 9.4.8]. □

Since the functor is represented in groups, we note that $\mathrm{Pic}_{X/k}$ is a group scheme. In addition, the connected component of the identity $\mathrm{Pic}^0_{X/k}$ is geometrically irreducible and finite type (in particular quasi-compact). It is separated by the following general (and easy) lemma:

LEMMA 1.1.2.5. A group scheme G over a field k is separated.

Proof. Consider the diagram

$$\begin{array}{ccc} G & \xrightarrow{\Delta} & G \times G \\ {\scriptstyle \mathrm{id}}\downarrow & & \downarrow{\scriptstyle (x,y)\mapsto xy^{-1}} \\ \mathrm{Spec}\, k & \xrightarrow{e} & G \end{array}$$

which expresses Δ as a base change of e. For k a field e is a closed immersion (which is generally not true over arbitrary schemes), hence so is Δ. □

Our goal is to show that for an abelian variety A, $\mathrm{Pic}^0_{A/k}$ is also an abelian variety. Thus, what we have remaining is to show it is proper and smooth. We will first show it is proper.

PROPOSITION 1.1.2.6. If X is smooth and projective over k, then $\mathrm{Pic}^0_{X/k}$ is proper.

Proof. We will show that $\mathrm{Pic}_{X/k}$ satisfies the valuative criterion for properness, then since $\mathrm{Pic}^0_{X/k}$ is finite type (and Noetherian), this will imply the proposition. So we need to show that the structure map $f : \mathrm{Pic}_{X/k} \to \operatorname{Spec} k$ satisfies the following: for any valuation ring A with field of fractions K, fitting into commutative diagram

$$\begin{array}{ccc} \operatorname{Spec} K & \longrightarrow & \mathrm{Pic}_{X/k} \\ \downarrow & \nearrow & \downarrow \\ \operatorname{Spec} A & \longrightarrow & \operatorname{Spec} k \end{array}$$

there is a unique map $\operatorname{Spec} A \to \mathrm{Pic}_{X/k}$ making the diagram commute. Since all of these are k-morphisms, this is equivalent to the map $\mathrm{Pic}_{X/k}(A) \to \mathrm{Pic}_{X/k}(K)$ being bijective. So we have reduced to showing the map

$$\mathrm{Pic}(X_A) \Big/ f_A^*(\mathrm{Pic}(A)) \to \mathrm{Pic}(X_K) \Big/ f_A^*(\mathrm{Pic}(K))$$

is bijective. We already know from Lemma 1.1.2.5 that this is injective (via the valuative criterion for separatedness). So, we only need to prove surjectivity. Because X is smooth, each of its connected components is integral (a smooth connected scheme over a field is integral). Surjectivity on each component implies surjectivity, so we may further assume X is connected and so line bundles on these components correspond to Cartier divisors [Har77, Ch. 2, §6, Prop. 6.15]. Any Cartier divisor may be written as a difference of *effective* Cartier divisors which correspond to codimension 1 closed subschemes. Our problem now reduces to extending a codimension 1 subscheme of the generic fiber of a projective scheme to that of a projective scheme over a valuation ring. But this may be done by taking the scheme-theoretic closure in X_A. □

In order to show smoothness of $\mathrm{Pic}^0_{A/k}$ for A an abelian variety, we need to develop the theory of line bundles more carefully in the specific case of abelian varieties. We will start this theory in the following section, but the proof of smoothness is given in §1.2.2.

§1.1.3 Line Bundles on Abelian Varieties

THEOREM 1.1.3.1 (Seesaw Theorem). Let X be a geometrically reduced, geometrically connected proper scheme over k, Y a k-scheme, and $\mathscr{L}$ a line bundle on $X \times Y$. There exists a closed subscheme $Y_1 \hookrightarrow Y$ such that for any k-scheme S, and diagram

$$\begin{array}{ccc} X_S & \xrightarrow{1_X \times f} & X \times Y \\ \downarrow & & \downarrow \\ S & \xrightarrow{f} & Y, \end{array}$$

the line bundle $\mathscr{N} := (1_X \times f)^* \mathscr{L}$ on X_S comes from a line bundle on S if and only if f factors through Y_1.

Note that this theorem can be proved in more generality for general locally free coherent sheaves (see [GW23, Ch. 24, §24.12, Thm. 24.66]), but for our purposes the case of line bundles will suffice.

LEMMA 1.1.3.2. Let $f : X \to S$ be a proper flat surjective morphism with geometrically reduced and connected fibers, with S locally noetherian. For a line bundle $\mathscr{L}$ on X, $\mathscr{L} = f^*\mathscr{M}$ for some line bundle $\mathscr{M}$ if and only if $f_*\mathscr{L}$ is invertible and the natural map $f^*f_*\mathscr{L} \to \mathscr{L}$ is an isomorphism (and thus $\mathscr{M} = f_*\mathscr{L}$). Additionally, in such cases, the formation of $f_*\mathscr{L}$ commutes with arbitrary base change on S.

Proof. By [GW23, Ch. 24, §24.12, Lem. 24.67], it is sufficient to prove that $\mathscr{O}_S = f_*\mathscr{O}_X$. To see this, we clearly have a map $f_*\mathscr{O}_X \to \mathscr{O}_S$ induced by $X \to S$ and it suffices to prove that this is an isomorphism on stalks. But, because f is flat, we have that for any $s \in S$, $(f_*\mathscr{O}_X)_s = \Gamma(\mathscr{O}_{f^{-1}(s)}, f^{-1}(s))$. Further, since f is surjective, these are all non-empty. In addition, if $\mathscr{O}_S = k$ and the induced map is an isomorphism i.e. $\Gamma(\mathscr{O}_{f^{-1}(s)_{\overline{k}}}, f^{-1}(s)_{\overline{k}}) \simeq \mathscr{O}_{S_{\overline{k}},s} = \mathscr{O}_{S,s} \otimes_k \overline{k} = \overline{k}$, then we must also have an isomorphism $\Gamma(\mathscr{O}_{f^{-1}(s)}, f^{-1}(s)) \simeq \mathscr{O}_{S,s}$ for dimension reasons. Thus, we have reduced to showing that for Z a proper reduced connected k-scheme for k an algebraically closed field, $\Gamma(\mathscr{O}_Z, Z) = \overline{k}$. But, we know that the global sections $s \in \Gamma(\mathscr{O}_Z, Z)$ correspond to maps $f^{-1}(s) \to \mathbf{A}_k$. But by our assumptions on Z the image of this must be proper, closed, and connected and hence a closed point. Therefore we have an isomorphism $\Gamma(\mathscr{O}_Z, Z) \simeq k$.

As for $\mathscr{M} = f_*\mathscr{L}$ commuting with arbitrary base change, let $g : T \to S$ be arbitrary and

$$\begin{array}{ccc} X_T & \xrightarrow{g_X} & X \\ {\scriptstyle f_T}\downarrow & & \downarrow{\scriptstyle f} \\ T & \xrightarrow{g} & S \end{array}$$

be the associated cartesian diagram. We want $g^*f_*\mathscr{L} \stackrel{?}{=} (f_T)_*g_X^*\mathscr{L}$. Note that we have $g_X^*\mathscr{L} = g_X^*f^*\mathscr{M} = f_T^*g^*\mathscr{M}$. But this implies that $g_X^*\mathscr{L} = f_T^*(\mathscr{M}')$ where $\mathscr{M}' = g^*\mathscr{M}$ is the pullback of a line bundle and thus a line bundle on T. This implies that $g^*\mathscr{M} = (f_T)_*g_X^*\mathscr{L}$ and since $\mathscr{M} = f_*\mathscr{L}$, this is our desired result. □

We can now prove the Theorem.

Proof of 1.1.3.1. Let $\phi : X_S \to S$ be the proper flat structure morphism. By the above Lemma 1.1.3.2, the line bundle $\mathscr{N} := (1_X \times f)^*\mathscr{L}$ comes from a line bundle on S if and only if $\phi_*(\mathscr{N})$ is invertible and $\phi^*\phi_*\mathscr{N} \to \mathscr{N}$ is an isomorphism. As these conditions may be checked after a finite Galois extension K/k, by Galois descent for closed subschemes, it suffices to build Y_1 after this finite Galois extension. Thus, we may assume $X(k) \neq \emptyset$.

Note that the question of whether $\mathscr{L}$ comes from S is equivalent to $\mathscr{L} \otimes p_2^*(\mathscr{M})$ coming from S for any $\mathscr{M} \in \mathrm{Pic}(Y)$. Thus, the question depends only on the coset of $\mathscr{L}$ in $\mathrm{Pic}(X \times Y)/p_2^*(\mathrm{Pic}(Y))$. But this coset is equivalent to a k-morphism $Y \xrightarrow{j} \mathrm{Pic}_{X/k}$. By functoriality, the pullback of $\mathscr{L}$ to X_S corresponds to the map $S \to Y \to \mathrm{Pic}_{X/k}$, so we see that this line bundle comes from S if and only if the composite map is zero. I.e., $S \to Y$ factors through the scheme-theoretic preimage $Y_1 := j^{-1}(0)$. □

THEOREM 1.1.3.3 (Theorem of the Cube). Let Z be separated and finite type over k, and X, Y be proper schemes over k. Further, assume that X, Z are geometrically integral, and Y is geometrically reduced and connected. Let $x_0 \in X(k)$, $y_0 \in Y(k)$, and $z_0 \in Z(k)$ and suppose $\mathscr{L}$ is a line bundle on

$X \times Y \times Z$ such that $\mathscr{L}_{x_0} := \mathscr{L}|_{\{x_0\}\times Y\times Z} \simeq \mathscr{O}_{Y\times Z}$, and $\mathscr{L}_{y_0}, \mathscr{L}_{z_0}$ are also trivial. Then $\mathscr{L} \simeq \mathscr{O}_{X\times Y\times Z}$.

Proof. Since $\mathscr{L}_{y_0}$ is trivial, $\mathscr{L}$ is classified by a map $\phi : X \times Z \to \mathrm{Pic}_{Y/k}$ (by the Remark following Definition 1.1.2.2). We want to show this is zero (note that our hypotheses on Y make sure the Picard scheme exists). The hypotheses on $X \times Y$ lets us apply the rigidity theorem, and we have that $X \times \{z_0\} \to 0$ because $\mathscr{L}_{z_0} \simeq \mathscr{O}_{X\times Y}$. Thus, $\phi(x, z) = \phi(x_0, z)$ for all z. But, $\phi(x_0, z) = 0$ since $\mathscr{L}_{x_0}$ is trivial on $Y \times Z$. □

We can now deduce the following, which will be useful for studying how multiplication interacts with line bundles:

COROLLARY 1.1.3.4 (Cubical Structure on $\mathscr{L}$). Let A be an abelian variety, $\mathscr{L} \in \mathrm{Pic}(A)$, S a k-scheme, and $a_1, a_2, a_3 \in A(S)$. Then the line bundle

$$(a_1+a_2+a_3)^*\mathscr{L} \otimes (a_1+a_2)^*\mathscr{L}^{-1} \otimes (a_1+a_3)^*\mathscr{L}^{-1} \otimes (a_2+a_3)^*\mathscr{L}^{-1} \otimes a_1^*\mathscr{L} \otimes a_2^*\mathscr{L} \otimes a_3^*\mathscr{L} \otimes (e^*\mathscr{L})_S^{-1}$$

on S is canonically trivial.

Proof. We may reduce to the universal case $A \times A \times A$ and $\alpha_i = \mathrm{pr}_i : A \times A \times A \to A$ since each a_i is canonically a pullback of α_i via

$$\begin{array}{ccc} S & \xrightarrow{(a_1,a_2,a_3)} & A\times A\times A \\ & \searrow^{a_i} & \downarrow{\scriptstyle \alpha_i} \\ & & A. \end{array}$$

In this case, we have

$$\begin{aligned} \alpha_1 + \alpha_2 + \alpha_3 = m : A^3 \to A \\ \alpha_i + \alpha_j = m_{ij} : A^2 \to A, i \neq j \\ \alpha_i = \mathrm{pr}_i : A^3 \to A, \end{aligned}$$

so we have reduced to showing the line bundle

$$\mathscr{M} = m^*\mathscr{L} \otimes m_{12}^*\mathscr{L}^{-1} \otimes m_{13}^*\mathscr{L}^{-1} \otimes m_{23}^*\mathscr{L}^{-1} \otimes \mathrm{pr}_1^*(\mathscr{L}) \otimes \mathrm{pr}_2^*(\mathscr{L}) \otimes \mathrm{pr}_3^*(\mathscr{L}) \otimes (0^*\mathscr{L})_A^{-1}$$

on A^3 is trivial. To apply the Theorem of the Cube (1.1.3.3), we first study the restriction

$$\mathscr{M}|_{\{0\}\times A\times A} = \mu^*\mathscr{L} \otimes q_1^*\mathscr{L}^{-1} \otimes q_2^*\mathscr{L}^{-1} \otimes \mu^*\mathscr{L}^{-1} \otimes (0^*\mathscr{L}) \otimes q_1^*\mathscr{L} \otimes q_2^*\mathscr{L} \otimes (0^*\mathscr{L})^{-1}$$

where $\mu, q_1, q_2 : A \times A \to A$ are the multiplication map, and projections respectively. This is clearly isomorphic to $\mathscr{O}_{A\times A}$. By symmetry, both $\mathscr{M}|_{A\times\{0\}\times A}$ and $\mathscr{M}|_{A\times A\times\{0\}}$ are trivial. Thus, by the Theorem of the Cube (1.1.3.3), $\mathscr{M}$ is trivial. □

The $\phi_{\mathscr{L}}$ construction

For $x \in A(S)$, let $t_x : A_S \to A_S$ be the translation map $y \mapsto x + y$.

DEFINITION 1.1.3.5. Given a line bundle $\mathscr{L}$ on A, we define the map of k-schemes

$$\phi_{\mathscr{L}} : A \to \mathrm{Pic}^0_{A/k}$$

functorially by $x \mapsto t_x^*\mathscr{L} \otimes \mathscr{L}^{-1}$.

Note that this sends $0 \mapsto \mathscr{O}_A$ (which is why this map is to the connected component of the identity). We will soon show that $\phi_{\mathscr{L}}$ is a homomorphism.

REMARK 1.1.3.6. This homomorphism will be very important throughout this paper as it will allow us to relate an abelian variety A to its dual $\widehat{A} := \mathrm{Pic}^0_{A/k}$. For example, if $\mathscr{L}$ is ample (or anti-ample by 1.1.3.8), then $\phi_{\mathscr{L}} : A \to \widehat{A}$ is an isogeny (surjective, finite, flat homomorphism), which we will call a **polarization**.

COROLLARY 1.1.3.7 (Theorem of the Square). The map $\phi_{\mathscr{L}}$ is a homomorphism; i.e. for any k-scheme T and $x, y \in A(T)$, we have canonical isomorphisms of line bundles on A_T

$$t_{x+y}^*(\mathscr{L}_T) \otimes \mathscr{L}_T \simeq t_x^*(\mathscr{L}_T) \otimes t_y^*(\mathscr{L}_T) \otimes [(x,y)^*(\Lambda(\mathscr{L}))]_{A_T}$$

where $\Lambda(\mathscr{L}) := m^*\mathscr{L} \otimes p_1^*\mathscr{L}^{-1} \otimes p_2^*\mathscr{L}^{-1} \otimes (e^*\mathscr{L})_A$ on $A \times A$ (which we call the "*Mumford construction*").

Note that in the special case $T = \operatorname{Spec} K$, for K/k a field extension, this asserts that

$$t_{x+y}^*(\mathscr{L}_T) \otimes \mathscr{L}_T \simeq t_x^*(\mathscr{L}_T) \otimes t_y^*(\mathscr{L}_T)$$

which is the version which appears in [Mum74]. Note that this is sufficient to show that $\phi_{\mathscr{L}}$ is a homomorphism.

Proof. Define $[x] = x \circ \mathrm{pr}_2 : A_T \to T \xrightarrow{x} A$. Apply Cor 1.1.3.4 with $S = A_T$, $a_1 = \mathrm{pr}_1$, $a_2 = [x]$, and $a_3 = [y]$. This gives us

$$(\mathrm{pr}_1 + [x] + [y])^*\mathscr{L} \otimes \mathrm{pr}_1^*\mathscr{L} \otimes [x]^*\mathscr{L} \otimes [y]^*\mathscr{L} \simeq ([x]+[y])^*\mathscr{L} \otimes (\mathrm{pr}_1 + [x])^*\mathscr{L} \otimes (\mathrm{pr}_1 + [y])^*\mathscr{L} \otimes (e^*\mathscr{L})_S^{-1}$$

over S. We have that $\mathrm{pr}_1 + [x]$ is the composition $A_T \xrightarrow{t_x} A_T \to A \in A(S)$ and $[x] + [y] = [x+y] \in A(S)$. Thus, this simplifies to

$$t_{x+y}^*\mathscr{L}_T \otimes \mathscr{L}_T \otimes x^*\mathscr{L}_T \otimes y^*\mathscr{L}_T \simeq (x+y)^*\mathscr{L}_T \otimes t_x^*\mathscr{L}_T \otimes t_y^*\mathscr{L}_T \otimes (e^*\mathscr{L})_T^{-1}$$

rearranging this gives us the corollary. □

We will now prove a series of nice properties of the maps $\phi_{\mathscr{L}}$.

PROPOSITION 1.1.3.8. $\phi_{\mathscr{L}} \circ [-1]_A = \phi_{\mathscr{L}^{-1}}$ as maps $A \to \mathrm{Pic}^0_{A/k}$.

Proof. Note that because this is just checking an equality of maps, it suffices to check on geometric points. So, for $x \in A(\overline{k})$, we want $t_{-x}^*\mathscr{L} \otimes \mathscr{L}^{-1} \simeq t_x^*(\mathscr{L}^{-1}) \otimes \mathscr{L}$, or equivalently $t_{-x}^*\mathscr{L} \otimes t_x^*\mathscr{L} \simeq \mathscr{L} \otimes \mathscr{L}$, but by the Theorem of the Square (1.1.3.7), the left side is isomorphic to $t_{x-x}^*\mathscr{L} \otimes \mathscr{L} \simeq \mathscr{L} \otimes \mathscr{L}$. □

PROPOSITION 1.1.3.9. We have $\phi_{\mathscr{L}_1 \otimes \mathscr{L}_2} = \phi_{\mathscr{L}_1} + \phi_{\mathscr{L}_2}$

Proof. Again, we may check on geometric points. For $x \in A(\overline{k})$, $\phi_{\mathscr{L}_1 \otimes \mathscr{L}_2}(x) = t_x^*(\mathscr{L}_1 \otimes \mathscr{L}_2) \otimes \mathscr{L}_1^{-1} \otimes \mathscr{L}_2^{-1}$. But this is just $t_x^*\mathscr{L}_1 \otimes \mathscr{L}_1^{-1} \otimes t_x^*\mathscr{L}_2 \otimes \mathscr{L}_2^{-1}$ which is the right side. □

PROPOSITION 1.1.3.10. For $x \in A(K)$ for any field K/k, $\phi_{t_x^*\mathscr{L}} = \phi_{\mathscr{L}}$ for $\mathscr{L}$ on A_K.

Proof. As with the previous propositions, it suffices to check on $\overline{K}$-points, so we may assume $K = \overline{K}$. So, for all $y \in \overline{K}$, we must check

$$\phi_{t_x^*\mathscr{L}}(y) = \phi_{\mathscr{L}}(y).$$

This is equivalent to

$$t_y^*(t_x^*\mathscr{L}) \otimes t_x^*\mathscr{L}^{-1} = t_y^*\mathscr{L} \otimes \mathscr{L}^{-1},$$

i.e.,

$$t_{x+y}^*\mathscr{L} \otimes \mathscr{L} = t_x^*\mathscr{L} \otimes t_y^*\mathscr{L}$$

which is the statement of the Theorem of the Square (1.1.3.7). □

REMARK 1.1.3.11. This shows that any line bundle of the form $\mathscr{M} = t_x^*\mathscr{L} \otimes \mathscr{L}^{-1}$ has trivial $\phi_{\mathscr{M}}$. We will later show that any line bundle in $\mathrm{Pic}^0_{A/k}$ is of this form, and further that the only line bundles $\mathscr{M}$ such that $\phi_{\mathscr{M}}$ is trivial are exactly those in $\mathrm{Pic}^0_{A/k}$. Thus for abelian varieties, $\mathrm{Pic}^0_{C/k}$ will be characterized by having trivial $\phi_{\mathscr{L}}$, or equivalently $K(\mathscr{L}) = A$ as defined below. This is analogous with $\mathrm{Pic}^0_{C/k}$ is characterized by being degree 0 for curves C.

DEFINITION 1.1.3.12. Let $\mathscr{L}$ be a line bundle on A. We define

$$K(\mathscr{L}) := \ker \phi_{\mathscr{L}} = \phi_{\mathscr{L}}^{-1}(0)$$

Projectivity of Abelian Varieties

As result of all our work on line bundles, we get another surprising result about how strict the structure is on abelian varieties!

THEOREM 1.1.3.13. Any abelian variety A over a field k is projective.

Note that projectivity is equivalent to having an ample line bundle. Thus, our strategy for the proof of projectivity will be to produce an ample line bundle. We first reduce to the case $k = \overline{k}$ via the following proposition:

PROPOSITION 1.1.3.14. For a proper k-scheme X, if $X_{\overline{k}}$ is projective over $\overline{k}$ then so is X over k.

Proof. To show X is projective, the goal is to build an ample line bundle $\mathscr{L}$ on X. Since $X_{\overline{k}}$ is projective over $\overline{k}$, we may choose an embedding $j : X_{\overline{k}} \hookrightarrow \mathbf{P}^n_{\overline{k}}$. Since X and $\mathbf{P}^n_k$ are finite type, j descends to a closed immersion $X_K \hookrightarrow \mathbf{P}^n_K$ for some finite extension K/k which we may assume is normal by passing to a larger finite extension if necessary. Then since K/k factors as $K/K_0/k$ where K/K_0 is purely inseparable and K_0/k is Galois, we may check that projectivity descends in these two cases (purely inseparable and Galois).

Firstly, we will take care of the purely inseparable case. Assume K/k is purely inseparable and let $\pi : X_K \to X$ be the projection. The ideas is to use the Frobenius morphism: there is some $e > 0$ such that

$f^{p^e} \in k$ for all $f \in K$. Thus, for an ample $\mathscr{L}$ on X_K, the gluing data for the line bundle lives over k, and we will show $\mathscr{L}^{p^e}$ is the pullback of a line bundle on X which we will show is ample.

To begin, we have a natural map $\mathscr{O}_X \hookrightarrow \pi_*\mathscr{O}_{X_K}$, and $t \mapsto t^{p^e}$ gives a map $\pi_*\mathscr{O}_{X_K}^\times \hookrightarrow \mathscr{O}_X^\times$ (as may be easily checked locally over affines in X). Therefore, $\iota : \mathscr{O}_X^\times \hookrightarrow \pi_*(\mathscr{O}_{X_K}^\times)$ has cokernel killed by p^e. Since K/k is inseparable, π is a homeomorphism, so a line bundle over X_K trivializes Zariski-locally over X. In other words, there is an open cover $\{U_\alpha\}$ of X such that $\mathscr{L}$ is trivial on each $\pi^{-1}(U_\alpha)$. This implies $\mathrm{Pic}(X_K) = H^1(X, \pi_*\mathscr{O}_{X_K}^\times)$ (which is not easy to see since $\pi_*\mathscr{O}_{X_K}^\times$ is not quasi-coherent). Thus, the map $H^1(X, \mathscr{O}_X^\times) \to H^1(X, \pi_*\mathscr{O}_{X_K}^\times)$ is identified with the natural map $\mathrm{Pic}(X) \to \mathrm{Pic}(X_K)$, so the cokernel of this latter map is killed by p^e since the cokernel of the former is. Thus, for any $\mathscr{L}$ on X_K, $\mathscr{L}^{p^e} \simeq \mathscr{M}_K$ for some line bundle $\mathscr{M}$ on X.

In general, for a line bundle $\mathscr{N}$ on X, we claim that if $\mathscr{N}_K$ is ample, then $\mathscr{N}$ must be ample as well. By Serre's criterion for ampleness, $\mathscr{N}$ is ample if and only if for each coherent $\mathscr{F}$ on X, any $i > 0$ and sufficiently large m,

$$H^i(X, \mathscr{F} \otimes \mathscr{N}^{\otimes m}) = 0.$$

But by the theorem on cohomology and flat base change,

$$H^i(X, \mathscr{F} \otimes \mathscr{N}^{\otimes m})_K = H^i(X_K, \mathscr{F}_K \otimes \mathscr{N}_K^{\otimes m})$$

and this vanishes for $m \geq 0$ since $\mathscr{N}_K$ is ample.

If we let $\mathscr{L}$ be ample, then clearly $\mathscr{L}^{p^e}$ is. Thus, our $\mathscr{M}$ from before is ample, so we have produced an ample line bundle on X. This takes care of the K/k purely inseparable case.

If instead K/k is Galois, given ample $\mathscr{L}$ on X_K, we define our $\mathscr{N} = \bigotimes_{\gamma \in \mathrm{Gal}(K/k)} \gamma^*(\mathscr{L})$. The line bundle $\mathscr{N}$ is ample because the tensor product of ample line bundles is again ample and each $\gamma^*\mathscr{L}$ is ample on X_K since they are the pullbacks of an ample line bundle by an automorphism of X_K. It is clear that this line bundle descends (as it is clearly invariant under $\mathrm{Gal}(K/k)$), so $\mathscr{N} = \mathscr{M}_K$ for a line bundle $\mathscr{M}$ on X, and ampleness of $\mathscr{N}$ on X_K is inherited by its descent $\mathscr{M}$ on X (as above, this is clear by Serre's criterion for ampleness). □

We will now construct the divisor of our candidate line bundle. Let U be an affine open neighborhood of 0 in A. Consider $(A - U)_{\mathrm{red}}$, which is of pure codimension 1, a general fact given here:

> LEMMA 1.1.3.15. Let Y be a normal, locally noetherian separated scheme, U a dense affine open in Y. Then $Y - U$ has pure codimension 1.

Proof. See [Stacks, Tag 0BCQ] □

We claim that $\mathscr{O}_A(D)$ is ample. To prove this, we first consider the stabilizer of D under translation.

> LEMMA 1.1.3.16. For $D = A - U$ as above, the group $H = \{x \in A(\overline{k}) \mid t_x^*(D_{\overline{k}}) = D_{\overline{k}} \text{ topologically}\}$ is finite.

Proof. Assume without loss of generality that $k = \overline{k}$. Firstly, we will show that H is closed and hence proper. Since D has finitely many irreducible components, all of the same dimension, the condition

$t_x^* D = D$ is equivalent to $t_x^*(D) \subset D$. Thus, $x \in H$ if $x + d \in D$ for all $d \in D(k)$ or equivalently $x \in \cap_{d \in D(k)}(-d + D)$, which is clearly closed.

Now note that if t_x preserves D, then it also preserves U. Hence since $e \in U$, $H \subset U$. But since H is proper and contained in the affine U, H is finite. □

It turns out that on an abelian variety, the finiteness of this stabilizer is equivalent to ampleness. Before we prove this fact, we need

LEMMA 1.1.3.17. If C is a curve on A and E is an irreducible divisor on A such that $C \cap E = \emptyset$ then $t^*_{x_1 - x_2} E = E$ for all $x_i \in C$.

Proof. Let $\mathscr{M} = \mathscr{O}(E)$, so that $\mathscr{M}|_C = \mathscr{O}_C$ since C and E are disjoint. We note that $\deg(t_a^* \mathscr{M}|_C)$ is independent of a and is therefore equal to 0 for all $a \in A(k)$.

This is because $m : A \times C \to A$ is a proper flat map, and generally for $X \to S$ proper, $\mathscr{F}$ coherent and S-flat, $\chi(\mathscr{F}_s)$ is independent of $s \in S$. So, in particular, for a line bundle $\mathscr{L}$ on a scheme X proper and flat over S, $\chi(\mathscr{L}_s^{\otimes n})$ is independent of s. But, for a curve, the degree of a line bundle is a coefficient of this polynomial in n. Hence, by considering the pullback $m^* \mathscr{L}$ on $A \times C$, the degree of $t_a^*(\mathscr{M})|_C$ is independent of $a \in A(k)$.

But this implies that $t_x(C)$ and E can never intersect in a non-empty finite set of points, since that would imply that $t_x^*(\mathscr{L})|_C$ has positive degree. So, for all $a \in A(k)$, either $t_a^*(C)$ and E are disjoint, or $t_x(C) \subset E$. Let $x_1, x_2 \in C$ and $y \in E$. Then $t_{y - x_2}(C)$ and E meet at y, so $t_{y-x_2}(C) \subset E$ hence $y - x_2 + x_1 \in E$. Therefore $t_{x_1 - x_2}(E) \subset E$ which is what we wanted. □

Now we are prepared to prove the equivalences in the following

PROPOSITION 1.1.3.18. Let A be an abelian variety, $D \subset A$ an effective reduced Weil divisor, and $\mathscr{L} = \mathscr{O}(D) = \mathscr{I}_D^{-1}$. The following are equivalent:

(i) $\mathscr{L}^{\otimes 2}$ has no basepoints (i.e., $\mathscr{O}_Z \otimes \Gamma(\mathscr{L}^{\otimes 2}) \to \mathscr{L}^{\otimes 2}$ is surjective) and the associated k-morphism $A \to \mathbf{P}(\Gamma(\mathscr{L}^{\otimes 2}))$ is finite.

(ii) $\mathscr{L}$ is ample.

(iii) $K(\mathscr{L}) = \{x \in A : t_x^*(\mathscr{L}) \simeq \mathscr{L}\}$ is finite;

(iv) $H = \{x \in A(\overline{k}) | t_x^*(D_{\overline{k}}) = D_{\overline{k}} \text{ topologically}\}$ is finite;

Proof. We will show (i) $\implies$ (ii) $\implies$ (iii) $\implies$ (iv) $\implies$ (i).

(i) $\implies$ (ii): For the finite map $\varphi : A \to \mathbf{P}(\Gamma(\mathscr{L}^{\otimes 2}))$, we have $\varphi^*(\mathscr{O}(1)) \simeq \mathscr{L}^{\otimes 2}$. But, under a finite morphism, the pullback of an ample line bundle is ample (e.g. by Serre's cohomological criterion, and the fact that $H^i(A, \mathscr{F}) = H^i(\mathbf{P}(\Gamma(\mathscr{L}^{\otimes 2})), \varphi_* \mathscr{F})$ since φ is finite).

(ii) $\implies$ (iii): If $K(\mathscr{L})$ is not finite, let $Y = (K(\mathscr{L}))^0_{\text{red}}$, so that Y is an abelian subvariety of positive dimension, and the restriction $\mathscr{L}|_Y$ is ample on Y. Futhermore, $t_y^*(\mathscr{L}|_Y) = \mathscr{L}|_Y$ for all $y \in Y$. Hence, by the Seesaw Theorem, $\Lambda(\mathscr{L}) = m^* \mathscr{L} \otimes \text{pr}_1^* \mathscr{L}^{-1} \otimes \pi_2^* \mathscr{L}^{-1}$ is trivial on $Y \times Y$. Pulling back by the morphism $(\text{id}, [-1]) : Y \to Y \times Y$ (given by $y \mapsto (y, -y)$), we get that $\mathscr{L}|_Y \otimes [-1]^*(\mathscr{L}|_Y)$ is trivial on Y. But, $\mathscr{L}|_Y$

is ample, and so is $[-1]^*(\mathscr{L}|_Y)$ since $[-1]$ is an automorphism of Y, so $\mathscr{L}|_Y \otimes [-1]^*(\mathscr{L}|_Y)$ must be ample. But this is a contradiction since $\dim Y > 0$.

(iii) $\implies$ (iv): This is trivial since $K(\mathscr{L}) \supset H$.

(iv) $\implies$ (i): To show that $\mathscr{L}^{\otimes 2}$ has no basepoints, we first note that this is equivalent to the condition that for any $a \in A(k)$, there is an effective divisor $D' \sim 2D$ such that $a \notin \mathrm{Supp}(D')$. By the Theorem of the Square (1.1.3.7), we have that for any $x \in A(k)$, $t_x^*(D) + t_{-x}^* D \sim 2D$. But, for any $a \in A(k)$, we can produce some $x \in A(k)$ such that $a \pm x \notin D$. This is because this is equivalent to producing $x \notin \pm t_{-a}(D)$ which is a union of two irreducible codimension 1 subspaces and thus an x clearly exists. This proves $\mathscr{L}^{\otimes 2}$ has no basepoints.

For the claim that the resulting map $f : A \to \mathbf{P}(\Gamma(\mathscr{L}^{\otimes 2}))$ is finite, assume not. Then there must exist some closed irreducible curve $C \subset A$ lying in a fiber of f, i.e. such that $f(C)$ is a point. But the fibers of the map f are exactly the effective divisors in the linear system $|2D|$. Thus for every $D' \simeq 2D$, either $C \subset |D'|$ or $C \cap |D'| = \emptyset$. In particular, for $x \in A(k)$, either C is disjoint from $|t_x^*(D)| \cup |t_{-x}^*(D)|$ or C is contained in $|t_x^*(D)|$ or $|t_{-x}^*(D)|$. Since D is effective, we may assume $D = \sum n_i D_i$ with D_i irreducible. Then by the lemma 1.1.3.17 above, each D_i is invariant under translation by all points $x_1 - x_2$, $x_i \in C$. This contradicts (i) as there can only be finitely many points fixing D. Thus, we are done. □

We now have all of our ingredients and can return to the projectivity of A:

Proof of Theorem 1.1.3.13. Proposition 1.1.3.18, together with Lemma 1.1.3.16 shows that for $D = A - U$, U an affine open neighborhood of 0, the corresponding line bundle $\mathscr{O}_A(D)$ is ample and thus A is projective. □

§1.1.4 Torsion

We will now study the torsion subgroups of our abelian varieties. Studying the torsion subgroups allows us to focus on the much more simple structure of finite commutative group schemes. This begs the question: how much about an abelian variety can we deduce from the structure of its torsion? We will later prove a surprising result which shows that in fact, up to isogeny, the structure of an abelian variety over a finite field is determined by the structure of its torsion subgroups. There are other similar results for different classes of fields, and in general, the study of torsion subgroups gives a surprising amount of information about an abelian variety.

For an abelian variety A, we define

$$\begin{aligned} [n]_A : A &\to A \\ a &\mapsto na := \underbrace{a + a + \cdots + a}_{n \text{ times}} \end{aligned}$$

functorially for all $a \in A(S)$, or equivalently it is the composition

$$[n]_A : A \xrightarrow{\Delta} \underbrace{A \times \cdots \times A}_{n \text{ times}} \xrightarrow{m} A$$

where m denotes n-fold multiplication on A. We define the scheme theoretic kernel $A[n] := \ker([n]) = [n]^{-1}(\{0\})$. When the abelian variety which we are working with is clear, we will often denote this map

simply by $[n]$. Following the theme of deducing results about abelian varieties from their line bundles, we first give a result about how this map interacts with line bundles:

THEOREM 1.1.4.1. For $\mathscr{L}$ a line bundle on abelian variety A, we have

$$\begin{aligned}[n]^*\mathscr{L} &\simeq \mathscr{L}^{\otimes \frac{n^2+n}{2}} \otimes ([-1]^*\mathscr{L})^{\otimes \frac{n^2-n}{2}} \\ &\simeq \mathscr{L}^{\otimes n} \otimes (\mathscr{L} \otimes [-1]^*\mathscr{L})^{\otimes \frac{n^2-n}{2}} \\ &\simeq \mathscr{L}^{\otimes n^2} \otimes (\mathscr{L} \otimes [-1]^*\mathscr{L}^{-1})^{\otimes \frac{n-n^2}{2}}.\end{aligned}$$

Proof. To prove the first equality, we will induct on n (upwards for positive n and downwards for negative n). For $n = -1, 0, 1$, the equality is clear. For general n, apply Corollary 1.1.3.4 with $a_1 = [n+1]$, $a_2 = [1]$, and $a_3 = [-1]$. It then follows that

$$[n+1]^*\mathscr{L} \otimes [n+2]^*\mathscr{L}^{-1} \otimes [n]^*\mathscr{L}^{-1} \otimes [n+1]^*\mathscr{L} \otimes (\mathscr{L} \otimes [-1]^*\mathscr{L}) \simeq \mathscr{O}_A.$$

Rearranging this, we have

$$[n+2]^*\mathscr{L} \otimes [n+1]^*\mathscr{L}^{-2} \otimes [n]^*\mathscr{L} \simeq \mathscr{L} \otimes [-1]^*\mathscr{L}.$$

Hence, by induction, for some line bundles $\mathscr{M}_1$ and $\mathscr{M}_2$, we must have

$[n]^*\mathscr{L} \simeq (\mathscr{L} \otimes [-1]^*\mathscr{L})^{\frac{n^2-n}{2}} \otimes \mathscr{M}_1^n \otimes \mathscr{M}_2$ and letting $n = 0$ shows $\mathscr{M}_2$ is trivial and letting $n = 1$ shows $\mathscr{M}_1 \simeq \mathscr{L}$. The second and third lines are obvious. □

An abelian variety is a "divisible group" in the following sense:

THEOREM 1.1.4.2. The map $[n] : A \to A$ is a finite flat surjection for $n \neq 0$.

Proof. We claim that it is actually enough to show that $\ker([n])$ is finite. If so, clearly $[n]$ is quasi-finite. So, by dimension considerations, the image must be dense. Additionally, since A is proper, $[n]$ is also proper (so universally closed), and thus, its image is closed and hence $[n]$ is surjective. Since in this case, $[n]$ is quasi-finite and proper, $[n]$ is finite. Therefore, we see that for all $a \in A(k)$,

$$\dim \mathscr{O}_{A,a} = \dim \mathscr{O}_{A,[n](a)} + \dim \mathscr{O}_{[n]^{-1}([n](a)),a} = \mathscr{O}_{A,[n](a)}$$

since $\dim \mathscr{O}_{[n]^{-1}([n](a)),a} = 0$ by the finiteness of $\ker([n])$. So, by the Miracle Flatness theorem (1.1.4.3 below), $[n]$ is flat.

Now we will show that the kernel is finite. Let $\mathscr{L}$ be an ample line bundle on A. Then by Theorem 1.1.4.1,

$$[n]^*\mathscr{L} \simeq \mathscr{L}^{\frac{n^2+n}{2}} \otimes ([-1]^*\mathscr{L})^{\frac{n^2-n}{2}}.$$

Since $[-1]$ is an automorphism, $[-1]^*\mathscr{L}$ is also ample (this takes care of the $[-1]^*$ case) and for $|n| > 1$, $\frac{n^2+n}{2}$ and $\frac{n^2-n}{2}$ are both positive. Thus, for all $n \neq 0$, $[n]^*\mathscr{L}$ is also ample.

If we let $Z = A[n]$, then $[n]^*\mathscr{L}|_Z = \mathscr{O}_Z$ and so $\mathscr{O}_Z$ is ample. But, if Z is a proper scheme such that $\mathscr{O}_Z$ is ample over a noetherian affine base, then Z is finite. This follows since Serre's vanishing theorem implies that repeatedly twisting by $\mathscr{O}_Z$ eventually kills the higher cohomology of coherent sheaves, but since twisting by $\mathscr{O}_Z$ does nothing, the higher cohomology $H^i(Z, \mathscr{O}_Z) = 0$ for $i > 0$, so by Serre's criterion for

affineness tells us that Z is affine. But it is also proper, so $H^0(Z, \mathscr{O}_Z)$ is finite and thus $Z = \operatorname{Spec} H^0(Z, \mathscr{O}_Z)$ is finite. □

THEOREM 1.1.4.3 (Miracle Flatness). Suppose $f : A \to B$ is a local map of local noetherian rings such that

$$\dim B = \dim A + \dim(B/f(\mathfrak{m}_A)B),$$

A is regular, and B is Cohen-Macaulay. Then f is flat.

Proof. See [Mat87, §23.1]. □

Next, we would like to compute $\deg([n])$. This makes sense since $[n] : A \to A$ is finite flat and A is connected and noetherian, so the geometric fibers each have the same degree. Also note that $[-n]$ and $[n]$ have the same degree since $[-1]$ is an isomorphism. So, we may assume $n > 0$. We first begin by giving the following definition:

DEFINITION 1.1.4.4. Let X be a proper k-scheme of dimension g and $\mathscr{L}$ a line bundle on X. Let $\mathscr{F}$ be a coherent sheaf on X. Then we consider $\chi(\mathscr{F} \otimes \mathscr{L}^{\otimes r})$. This is a numeric polynomial in r of degree $\leq g = \dim X$ of the form

$$\frac{\deg_{\mathscr{L}}(\mathscr{F})}{g!} r^g + \text{lower-order terms}$$

for some $\deg_{\mathscr{L}}(\mathscr{F}) \in \mathbf{Z}$ which we call **the degree of $\mathscr{F}$ with respect to $\mathscr{L}$**. When $\mathscr{F} = \mathscr{O}_X$, we call $\deg_{\mathscr{L}}(\mathscr{O}_X)$ the **degree of** $\mathscr{L}$ and denote it by $\deg \mathscr{L}$. In particular, $\chi(\mathscr{L}^{\otimes r}) = \frac{\deg \mathscr{L}}{g!} r^g +$ lower-order terms.

REMARK 1.1.4.5. Note that if we replace $\mathscr{L}$ in the above definition by a positive power, then we may absorb that power into r and hence, for any $i \in \mathbf{Z}_{>0}$,

$$\deg(\mathscr{L}^i) = i^g \deg(\mathscr{L}).$$

Additionally, if $\mathscr{L}$ is ample, then $\deg \mathscr{L} > 0$ by [EGA, IV_2, 5.3]. This will be important in describing $\deg[n]$ (Thm. 1.1.4.7). But before we begin, we prove

LEMMA 1.1.4.6. Let X be a proper integral k-scheme, $f : X' \to X$ a finite map, and $\deg f = d$. Then for all line bundles $\mathscr{L}$ on X,

$$\deg(f^*\mathscr{L}) = d \cdot \deg(\mathscr{L}).$$

Proof. Since the degree of a line bundle is determined by $\chi((f^*\mathscr{L})^{\otimes r}) = \chi(f^*(\mathscr{L}^{\otimes r}))$ for all r, we will study this. But since cohomology acts well under pushforward by a finite map, we have that $\chi(f_* f^* \mathscr{L}^{\otimes r}) = \chi(f_*(\mathscr{O}_{X'}) \otimes \mathscr{L}^{\otimes r})$ by the projection formula. Let $\mathscr{F} = f_*(\mathscr{O}_{X'})$ which is a coherent sheaf on X with generic stalk of rank d.

We want to show that $\deg_{\mathscr{L}}(\mathscr{F}) = \operatorname{rank}(\mathscr{F}) \deg(\mathscr{L})$ (see Def. 1.1.4.4). More generally, we can assert this for any $\mathscr{F} \in \operatorname{Coh}(X)$ with $\operatorname{rank}(\mathscr{F}) := \dim \mathscr{F}_\eta$ where η is the generic point of X. I.e., we want to show

$$\chi(\mathscr{F} \otimes \mathscr{L}^{\otimes r}) \stackrel{?}{=} \frac{\operatorname{rank}(\mathscr{F}) \cdot \deg(\mathscr{L})}{g!} r^g + \text{lower-order terms} \tag{1.1.4-I}$$

for all $r \in \mathbf{Z}$. To prove this, we will apply Grothendieck's unscrewing technique. Let $\mathscr{C} \subset \mathrm{Coh}(X)$ be the subcategory of all coherent $\mathscr{F}$ satisfying equation 1.1.4-I. We will show this is actually the entire category.

Now, χ is additive over short exact sequences,

$$0 \to \mathscr{F}' \to \mathscr{F} \to \mathscr{F}'' \to 0,$$

so if $\mathscr{F}', \mathscr{F}'' \in \mathscr{C}$, then $\mathscr{F} \in \mathscr{C}$. Grothendieck's unscrewing lemma [EGA, III, Thm 3.1.2] asserts that if $\mathscr{C}$ is an exact full subcategory of $\mathrm{Coh}(X)$, and if for all irreducible and reduced closed subschemes $Y \subset X$, there exists some $\mathscr{G}_Y \in \mathscr{C}$ with $\mathrm{Supp}(\mathscr{G}_Y) = Y$ and $\mathrm{rank}(\mathscr{G}_Y) = 1$, then $\mathscr{C} = \mathrm{Coh}(X)$.

Returning to our case, by what we noted above, $\mathscr{C}$ is exact and $\mathscr{G}_Y = \mathscr{O}_Y \in \mathscr{C}$ works for the second condition since for $Y = X$, this is clear by the definition of degree, and for $Y \subsetneq X$, both sides are 0. □

THEOREM 1.1.4.7. If A is an abelian variety of dimension g, then $\deg([n]) = n^{2g}$

REMARK 1.1.4.8. In [Mum74, Application 3, §6], there is an intersection-theoretic proof for this fact over an algebraically closed field which may be of interest. We give a different argument using line bundles here.

Proof. We will prove this by leveraging the degree of line bundles and using Lemma 1.1.4.6. Choose an ample line bundle $\mathscr{L}$ on A and replace $\mathscr{L}$ with $\mathscr{L} \otimes [-1]^*\mathscr{L}$ so that $\mathscr{L} \simeq [-1]^*\mathscr{L}$ (i.e. $\mathscr{L}$ is *symmetric*). Then by Theorem 1.1.4.1,

$$[n]^*\mathscr{L} \simeq \mathscr{L}^{\frac{n^2+n}{2}} \otimes ([-1]\mathscr{L})^{\frac{n^2-n}{2}} \simeq \mathscr{L}^{n^2}.$$

By Remark 1.1.4.5, we know that

$$\deg([n]^*\mathscr{L}) = \deg(\mathscr{L}^{n^2}) = (n^2)^g \deg \mathscr{L} = n^{2g} \deg \mathscr{L} \neq 0$$

So, by 1.1.4.6, we must have $\deg[n] = n^{2g}$.

□

We are now prepared to examine the structure of $A[n]$. Firstly, we want to reduce the problem to the prime power case via the following lemma:

LEMMA 1.1.4.9. Let $G \to S$ be a commutative group scheme killed by $n \neq 0$. Suppose $n = ab$ where $\gcd(a, b) = 1$. Then

$$G[a] \times G[b] \xrightarrow{m} G$$

is an isomorphism.

Proof. For any S-scheme T, $G(T) = G(T)[a] \times G(T)[b]$ since $G(T)$ is a $\mathbb{Z}/n\mathbb{Z}$-module. □

Then to examine $A[n]$ (which is the fiber over 0 of $[n]$), we use the following general

LEMMA 1.1.4.10. For $f : X \twoheadrightarrow Y$ a finite surjection between geometrically integral k-schemes of finite type, there exists a dense open $U \subset Y$ such that for all $u \in U(\overline{k})$, the size of the fiber over u is the

separable degree of the extension of function fields, i.e. $\#X_u(\overline{k}) = \deg_s(f)$.

Proof. Since we assumed X, Y are geometrically integral, and since $X_u(\overline{k}) = (X_u)_{\overline{k}}(\overline{k})$, we may assume without loss of generality that $k = \overline{k}$. Furthermore, we may assume $X = \operatorname{Spec} B$ and $Y = \operatorname{Spec} A$ are affine. Since f is surjective and A, B are domains, this corresponds to an injective morphism $A \to B$. Now, let $k(X)$ and $k(Y)$ be the respective function fields and let K be the separable closure of $k(Y)$ in $k(X)$. So we have the tower $k(X)/K/k(Y)$ where $k(X)/K$ is purely inseparable and $K/k(Y)$ is separable. Thus, by taking C to be the ring of integers of K in B and letting $Z = \operatorname{Spec} C$, we have reduced to proving the result for $X \to Z$ with function fields separable and $Z \to Y$ with function fields purely inseparable.

First suppose that we are in the separable case and as before, let $X = \operatorname{Spec} B$ and $Y = \operatorname{Spec} A$ and let $[k(X) : k(Y)] = n$. By the primitive element theorem, we may choose an element $\alpha \in \operatorname{Frac}(B)$ such that

$$\operatorname{Frac}(B) = \operatorname{Frac}(A)[\alpha]$$

with characteristic polynomial $P(T) \in k(Y)[t]$ a separable polynomial of degree n. Replacing A with a localization if necessary, we may assume the element t is integral over A. Since B is finitely generated, by localizing further, we may assume

$$A \subset B \subset A[\alpha].$$

Since B and $A[\alpha]$ have the same fraction field, by yet again localizing further, we may assume

$$B = A[\alpha] \simeq A[t]/(P(t)).$$

Potentially further localizing, we may assume the discriminant of $P(t)$, $\Delta \in A$ and $\Delta \neq 0$ since $P(t)$ is *separable*. Let $U = D(\Delta) \subset Y$ be the distinguished open set corresponding to Δ. Then for any $y \in Y$, the fiber $f^{-1}(y)$ is given by $k(y)[t]/(\overline{P}(t))$ where $k(y)$ is the residue field and $\overline{P} \in k(y)[T]$ is the reduction of $P(t)$. Since its discriminant is $\Delta(y) \neq 0$ for all $y \in U$, we have the it has n distinct roots (since $k(y) = k = \overline{k}$). Thus, we have our result for separable extensions.

If instead $k(X)/k(Y)$ is purely inseparable, we can break it down into a finite tower of field extensions

$$k(X) = K_0/K_1/\ldots/K_m = k(Y)$$

such that $[K_i : K_{i-1}] = p$ for all $i > 0$. So, we may treat the case $[k(X) : k(Y)] = p$. In this case, we know that the extension is primitive so that $\operatorname{Frac}(B) = \operatorname{Frac}(A)[\alpha] \simeq \operatorname{Frac}(A)[t]/(t^p - a)$ for some $a \in \operatorname{Frac}(A)$. Reducing as in the separable case to

$$B = A[\alpha] \simeq A[t]/(t^p - a),$$

$a \in A$, we conclude that each fiber $f^{-1}(y)$ is given by $k(y)[t]/(t^p - \overline{a})$, but since $k(y) = k = \overline{k}$, this splits as $k(y)[t]/(t - \overline{a})^p$ which has a single geometric point, so we are done. □

PROPOSITION 1.1.4.11. If $\operatorname{char} k \nmid n$, then $A[n](\overline{k}) \simeq (\mathbb{Z}/n\mathbb{Z})^{2g}$.

Proof. By Lemma 1.1.4.9, we may assume that $n = \ell^e$ where $\ell \neq \operatorname{char} k$ and $e \geq 1$. Since $[\ell]$ is finite, flat, and surjective, we have a short exact sequence

$$0 \to A[\ell](\overline{k}) \to A[\ell^e](\overline{k}) \xrightarrow{[\ell]} A[\ell^{e-1}](\overline{k}) \to 0.$$

So, by induction $\#A[\ell^e](\overline{k}) = \#A[\ell](\overline{k})^e$. By the structure of finite abelian ℓ-groups, we know that $A[\ell^e](\overline{k})$ is a quotient of $(\mathbb{Z}/\ell^e\mathbb{Z})^d$ with $\ell^d = \#A[\ell](\overline{k})$, but by cardinality reasons, this must be an isomorphism. Thus, we have reduced to showing that $A[\ell](\overline{k}) = (\mathbb{Z}/\ell\mathbb{Z})^{2g}$.

To show this, it suffices to show that $[\ell]$ induces a separable degree ℓ^{2g} extension on function fields. If this is the case, by Lemma 1.1.4.10, the geometric fibers would have size ℓ^{2g} over some dense open subset of A. But, homogeneity of group schemes, this implies that all geometric fibers of $[\ell]$ have the same size ℓ^{2g}, in particular $A[\ell] = [\ell]^{-1}(0)$ has size ℓ^{2g}. But, we know by Theorem 1.1.4.7 that $[\ell]$ is of degree ℓ^{2g}, so induces a field extension $k(A) \to k(A)$ of degree ℓ^{2g}. Since we assumed $\ell \neq \operatorname{char} k$, this is not divisible by $\operatorname{char} k$ and hence is separable. Thus, we are done. □

If $\ell = p = \operatorname{char} k$, we expect that the p-torsion would have fewer geometric points as can be seen in the case of elliptic curves – for example, "supersingular" elliptic curves have no nontrivial p-torsion $\overline{k}$-points!

Proposition 1.1.4.12. If k has characteristic $p > 0$ then $\dim_{\mathbb{F}_p} A[p](\overline{k}) \leq g$. Thus, $A[p^e](\overline{k}) \simeq (\mathbb{Z}/p^e\mathbb{Z})^i$ for all $e \geq 1$ and some $i \leq g$.

Proof sketch. Rather than give a full proof here, we will give a sketch and give a full proof in §1.3.2 in which we will call i the p-rank. We do so because the reason for this bound on i will become much clearer once we can sufficiently analyze $A[p^e]$ as a *group scheme*, not just its $\overline{k}$-points.

As before, by induction, we may reduce to the $e = 1$ case. Since $\deg[p] = p^{2g}$, by Lemma 1.1.4.10, we want to show that p^g divides the inseparable degree $\deg_i([p])$. To do so, we note that the map $[p]$ factors through the Frobenius (this will be clear from the discussion at the end of §1.3.2):

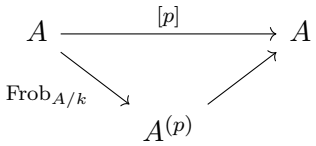

Then $\operatorname{Frob}_{A/k}$ has degree p^g because by the smoothness of A, it is Zariski-locally étale over $\mathbb{A}^g_k$. Thus, the separable degree $\deg_s([p]) \leq \deg([p])/\deg_i(\operatorname{Frob}_{A/k}) = p^g$, so we conclude by Lemma 1.1.4.10.

□

To summarize, we have proven the following:

Proposition 1.1.4.13. We have

1. $\deg[n] = n^{2g}$,
2. $[n]$ separable $\iff p \nmid n$
3. if $p \nmid n$, $A[n](\overline{k}) \simeq (\mathbb{Z}/n\mathbb{Z})^{2g}$, and
4. there is an integer i with $0 \leq i \leq g$ such that for all $e \geq 1$, $A[p^e](\overline{k} \simeq (\mathbb{Z}/p^e\mathbb{Z})^i$.

The study of torsion on abelian varieties will be continued in the second chapter in which it will play a crucial role in studying the structure of abelian varieties. We will now briefly introduce an object which will be at the center of Chapter 2.

DEFINITION 1.1.4.14. For prime ℓ, define the **ℓ-adic Tate module**

$$T_\ell(A) = \varprojlim A[\ell^n](\overline{k}) \simeq \begin{cases} \mathbf{Z}_\ell^{2g} & \ell \neq \operatorname{char} k \\ \mathbf{Z}_\ell^i, i \leq g & \ell = \operatorname{char} k. \end{cases}$$

We will eventually study the map

$$\mathbf{Z}_\ell \otimes \operatorname{Hom}_k(A, B) \to \operatorname{Hom}_{\mathbf{Z}_\ell}(T_\ell(A), T_\ell(B))$$

and will show it is actually injective! This is quite an incredible fact and will be an important step in proving Tate's Conjecture for finite fields.

Section 1.2

The Dual Abelian Variety

In this section we will show that $\widehat{A} := \mathrm{Pic}^0_{A/k}(A)$ is an abelian variety and we will show that this is a duality theory, i.e. that $\widehat{\widehat{A}} \simeq A$ canonically. In addition, we will study the maps $\phi_{\mathscr{L}} : A \to \widehat{A}$ (Def 1.1.3.5), which for ample $\mathscr{L}$ provides examples of "polarizations."

§1.2.1 The Poincaré Bundle

The dual abelian variety is akin to the dual of a vector space, and the first object we will study, the Poincaré bundle, will be in analogue to the evaluation pairing.

Definition 1.2.1.1. The **Poincaré bundle**, $\mathscr{P}_A$, on $A \times \widehat{A} = A \times \mathrm{Pic}^0(A)$ is the restriction of the universal line bundle on $A \times \mathrm{Pic}(A)$. I.e., the bundle such that for all $\alpha \in \widehat{A}$, the restriction $\mathscr{P}_A|_{A\times\{\alpha\}}$ represents the corresponding line bundle under the identification $\widehat{A} = \mathrm{Pic}^0(A)$. Moreover, we require $\mathscr{P}_A|_{\{0\}\times\widehat{A}} \simeq \mathscr{O}_{\widehat{A}}$. (Note that this uniquely characterizes $\mathscr{P}_A$ by the Seesaw Theorem 1.1.3.1).

The universal property of the Poincaré bundle $\mathscr{P}_A$ over $A \times \mathrm{Pic}^0_{A/k}$ is that any map $f : T \to \mathrm{Pic}^0_{A/k}$, is equivalent to a pr_2-trivialized line bundle $\mathscr{M} := (1 \times f)^*(\mathscr{P}_A)$ on $A \times T$, where this equivalence is given by $f(t) = \mathscr{M}_t \in \mathrm{Pic}^0_{A/k}(k(t))$ on A_t for all $t \in T$ (note that by remark 1.1.2.3, this is equivalent to giving an element of $\mathrm{Pic}_{A/k}(T)$). Applying this to the maps $\phi_{\mathscr{L}} : A \to \mathrm{Pic}^0_{A/k}$, this says that $\phi_{\mathscr{L}}$ is determined by the line bundle

$$\mathscr{M} := (1 \times \phi_{\mathscr{L}})^*(\mathscr{P}_A)$$

on $A \times A$. So, by the universal property, $\mathscr{M}|_{A\times\{x\}} = t_x^*(\mathscr{L}_K) \otimes \mathscr{L}_K^{-1}$ for all field extensions K/k. Here is a more familiar, useful description of $\mathscr{M}$:

Proposition 1.2.1.2. We have a canonical isomorphism

$$(1 \otimes \phi_{\mathscr{L}})^*\mathscr{P}_A \simeq \Lambda(\mathscr{L}) := m^*(\mathscr{L}) \otimes \mathrm{pr}_1^*(\mathscr{L})^{-1} \otimes \mathrm{pr}_2^*(\mathscr{L})^{-1}$$

as $e \times 1_A$-trivialized line bundles.

Proof. We will use the Seesaw Theorem (1.1.3.1), which gives us that it is enough to check an isomorphism on $\{e\} \times A$ and $A \times \{x\}$ for all $x \in A(\overline{k})$.

To see why this follows from the Seesaw Theorem, first note that by taking the "difference" of the two line bundles, it suffices to check that if S is a geometrically reduced k-scheme (such as A) then a line bundle $\mathscr{M}$ on $A \times S$ is trivial if and only if it is trivial on $\{e\} \times S$ and $A \times \{x\}$ for all $x \in S(\overline{k})$. By the Seesaw theorem, the second triviality condition implies that $\mathscr{M}$ is pulled back from a line bundle on S, and the first condition guarantees that this line bundle on S is trivial. Alternatively, the first triviality condition says that $\mathscr{M}$ is classified by a map $S \to \mathrm{Pic}^0_{A/k}$, hence the map $S \to \mathrm{Pic}^0_{A/k}$ is the trivial map (since S is geometrically reduced).

Returning to checking the equality of our specific line bundles, on $\{e\} \times A$, the left side is trivial because $\mathscr{P}_A$ is $e \times \mathrm{id}_A$-trivialized. On the right side, we get $\mathscr{L} \otimes (e^*\mathscr{L})_A^{-1} \otimes \mathscr{L}^{-1} \simeq \mathscr{O}_A$. On $A \times \{x\}$, the left side is $t_x^*\mathscr{L}_{\overline{k}} \otimes \mathscr{L}_{\overline{k}}^{-1}$. On the right, we get $t_x^*\mathscr{L}_{\overline{k}} \otimes \mathscr{L}_{\overline{k}}^{-1} \otimes (x^*\mathscr{L})_{A_{\overline{k}}}^{-1}$. Again, these are identified because $x^*\mathscr{L}$ is trivial. □

To justify some of the analogy between the Poincaré bundle with the evaluation pairing, we note that it gives us the ability to construct a morphism $\iota_A : A \to \widehat{\widehat{A}}$, which we will prove is an isomorphism in §1.2.6. Consider the flip isomorphism $\sigma : \widehat{A} \times A \simeq A \times \widehat{A}$. We see that $\sigma^*(\mathscr{P}_A)$ is a line bundle on $\widehat{A} \times A$ trivialized along $\widehat{e} \times 1$. Thus, by the universal property of the Picard scheme $\mathrm{Pic}_{\widehat{A}/k}$, this line bundle is classified by a map $A \to \mathrm{Pic}_{\widehat{A}/k}$ satisfying $e \mapsto 0$ (by the trivialization). This map clearly factors through the identity component, so is a pointed map $\iota_A : A \to \widehat{\widehat{A}}$. We will prove in §1.2.2 that $\widehat{A}$ and thus $\widehat{\widehat{A}}$ are abelian varieties. Thus, ι_A being a map between abelian varieties sending $e \mapsto e$, Corollary 1.1.1.3 tells us that ι_A is a homomorphism.

§1.2.2 Smoothness of $\mathrm{Pic}^0_{A/k}$

To finally prove that $\mathrm{Pic}^0_{A/k}$ is an abelian variety, we just need to show it is smooth. We have already showed that it is proper by Proposition 1.1.2.6 since A is projective by Theorem 1.1.3.13. To do so, we will use dimension considerations on tangent spaces.

Firstly, we know that for ample $\mathscr{L}$, $K(\mathscr{L}) := \ker \phi_{\mathscr{L}}$ is finite, so $\phi_{\mathscr{L}}$ has finite fibers. This follows from Proposition 1.1.3.18 for $\mathscr{L} = \mathscr{O}_A(D)$, D effective. For general ample $\mathscr{L}$, there is some n such that $\mathscr{L}^n = \mathscr{O}_A(D)$ for effective D and $\ker \phi_{\mathscr{L}} \circ [n] = \ker \phi_{\mathscr{L}^n}$ but by Theorem 1.1.4.2, $A[n] := \ker[n]$ is finite. Thus, by the properness of $\phi_{\mathscr{L}}$ (which follows from properness of the structure maps of A and $\mathrm{Pic}^0_{A/k}$), $\phi_{\mathscr{L}}$ is finite. It follows from quasi-finiteness that $\dim \mathrm{Pic}^0_{A/k} \geq g$. Thus, for smoothness, it suffices to show that the dimension of the tangent space $\dim T_0(\mathrm{Pic}^0_{A/k}) \leq g$ as this would imply that $\dim \mathrm{Pic}^0_{A/k} = g$, and that the finite $\phi_{\mathscr{L}}$ is surjective and flat (by generic flatness and homogeneity considerations over $\overline{k}$). We begin studying $T_0(\mathrm{Pic}^0_{A/k})$ with a general lemma:

> LEMMA 1.2.2.1. Let X be a projective, geometrically integral scheme over a field k. Then $\mathrm{Pic}_{X,k}$ has tangent space at the identity, $T_0(\mathrm{Pic}^0_{X/k})$ naturally identified with $H^1(X, \mathscr{O}_X)$ as a k-vector space.

Proof. Consider $S = k[\varepsilon]/(\varepsilon)^2$, so that the tangent space $T_0(\mathrm{Pic}^0_{X/k})$ is naturally identified with the group $\mathrm{Hom}_0(S, \mathrm{Pic}^0_{X/k})$ where Hom_0 denotes the set of morphisms of S to $\mathrm{Pic}^0_{X/k}$ mapping the unique point $s_0 \in S$ to 0. By the definition of $\mathrm{Pic}^0_{X/k}$, we have that

$$\begin{aligned}\mathrm{Hom}_0(S, \mathrm{Pic}^0_{X/k}) &= \{\text{Line bundles on } S \times X \text{ trivial on } \{s_0\} \times X\} \\ &= \ker\left(H^1(S \times X, \mathscr{O}_{S\times X}^\times) \to H^1(\{s_0\} \times X, \mathscr{O}_X^\times)\right)\end{aligned}$$

under the identification of $\mathrm{Pic}(X)$ with $H^1(X, \mathscr{O}_X^\times)$. But now we have an exact sequence of multiplicative sheaves

$$1 \to 1 + \varepsilon\mathscr{O}_X \to \mathscr{O}_{S\times X}^\times \to \mathscr{O}_X^\times \to 1$$

and as sheaves of abelian groups, $1+\varepsilon\mathscr{O}_X \simeq \mathscr{O}_X$. Therefore taking cohomology gives us the exact sequence

$$0 \to H^1(X, \mathscr{O}_X) \to H^1(S \times X, \mathscr{O}_{S\times X}^\times) \to H^1(X, \mathscr{O}_X^\times),$$

which, in light of our discussion above, gives us a natural isomorphism of $H^1(X, \mathscr{O}_X) \simeq T_0(\mathrm{Pic}^0_{X/k})$ as abelian groups. To see that this further gives us an isomorphism of k-vector spaces, we note that k-multiplication on the vector space $T_0(\mathrm{Pic}^0_{X/k})$ corresponds to multiplying ε by $a \in k$, which under the identification of $1+\varepsilon\mathscr{O}_X$ with $\mathscr{O}_X$ and extending to $H^1(X, \mathscr{O}_X)$ clearly gives us multiplication by a. □

Now to finish our proof that $\dim \mathrm{Pic}^0_{A/k} = g$ and $\phi_{\mathscr{L}}$ is surjective and flat, it suffices to show the following

PROPOSITION 1.2.2.2. $\dim H^1(A, \mathscr{O}_A) \leq g$.

Proof. We will reduce this to a general theorem of Borel given below (see Thm 1.2.2.3) applied to the cohomology ring of A.

Let $R = \bigoplus_{i\geq 0} H^i(A, \mathscr{O}_A)$ considered with its $\mathbf{Z}_{\geq 0}$-grading, and $R^j = 0$ for all $j > g$. This has an anti-commutative cup product $R \otimes_k R \to R$ denoted $a \otimes b \mapsto a \smile b$ (i.e., if $a \in R^i$, $b \in R^j$ then $a \smile b \in R^{i+j}$ and $b \smile a = (-1)^{ij}(a \smile b)$). Furthermore, $R^0 = k$ (since A is proper), and we may also say something about $R^g = H^g(A, \mathscr{O}_A)$: since $\Omega^1_{A/R} \simeq \mathscr{O}_A \otimes \mathrm{Lie}(A)^*$, so that $\Omega^g_{A/k} \simeq \mathscr{O}_A$ (since $\dim \mathrm{Lie}(A)^* = g$, but note that this is non-canonical), we have that $R^g = H^g(A, \Omega^g_{A/k})$. But this is 1-dimensional by Serre duality, so $R^g = k$. Likewise, the cup product $R^i \times R^{g-i} \to R^g$ is a perfect pairing by Serre duality (since the canonical bundle is trivial).

Consider the diagram

$$A \rightrightarrows A \times A \xrightarrow{m} A$$

where the inclusions are along identity elements ($A \to \{0\} \times A$ and $A \to A \times \{0\}$), we get a diagram

$$H^\bullet(A, \mathscr{O}_A) \leftleftarrows H^\bullet(\mathscr{O}_{A\times A}) \xleftarrow{m^*} H^\bullet(A, \mathscr{O}_A).$$

Combining this with the Künneth formula for $H^\bullet(\mathscr{O}_{A\times A})$, we have that

$$H^\bullet(A, \mathscr{O}_A) \otimes_k H^\bullet(A, \mathscr{O}_A) \xrightarrow{\sim} H^\bullet(A \times A, \mathscr{O}_{A\times A})$$

is an isomorphism. So, we get factorizations of the identity

$$R \leftleftarrows R \otimes R \xleftarrow{\mu} R$$

(tensor product as graded associative k-algebras) where μ is a k-algebra map giving R a structure of a Hopf algebra over k where the left two maps use that $R^0 = k$. We conclude that $\mu(x) = x \otimes 1 + 1 \otimes x +$ (higher degree terms) for all x. The result then follows from Theorem 1.2.2.3 below. □

THEOREM 1.2.2.3 (Borel). Let R be a finite-dimensional $\mathbf{Z}_{\geq 0}$-graded associative Hopf algebra over a field k such that $R^0 = k = R^g$, and $R^j = 0$ for all $j > g$. Assume multiplication $R^i \times R^{g-i} \to R^g$ is a perfect pairing for $0 \leq i \leq g$. Then

$$\dim R^1 \leq g,$$

and equality holds if and only if the natural map $\bigwedge(R^1) \to R$ is an isomorphism.

Proof. We may assume without loss of generality that $k = \overline{k}$ (since cohomology behaves well under flat base change). Then over perfect fields, Borel proved a general structure theorem for Hopf algebras which gives this result as a special case. See [Mil86, Lemma 15.2]. □

This concludes the proof that $\mathrm{Pic}^0_{A/k}$ is smooth, and thus an abelian variety. This allows us to define:

DEFINITION 1.2.2.4. If A is an abelian variety over k, we define the **dual abelian variety** $\widehat{A} := \mathrm{Pic}^0_{A/k}$.

§1.2.3 A Characterization of Line Bundles in $\mathrm{Pic}^0_{A/k}$

We have now seen that for ample $\mathscr{L}$, $\phi_{\mathscr{L}} : A \to \widehat{A}$ is an isogeny (i.e., finite flat surjective homomorphism). Continuing the remark after Proposition 1.1.3.10, by the surjectivity of these $\phi_{\mathscr{L}}$, we can begin characterizing the line bundles in $\mathrm{Pic}^0_{A/k}$:

LEMMA 1.2.3.1. For some ample $\mathscr{N}$ on A, a line bundle $\mathscr{M} \in \mathrm{Pic}(A) = \mathrm{Pic}_{A/k}(k)$ lies in $\mathrm{Pic}^0_{A/k}(k) = \widehat{A}(k)$ if and only if $\mathscr{M}_{\overline{k}} \simeq \phi_{\mathscr{N}}(x) = t_x^*(\mathscr{N}_{\overline{k}}) \otimes \mathscr{N}_{\overline{k}}^{-1}$ for some $x \in A(\overline{k})$.

Proof. Since $\phi_{\mathscr{N}} : A \to \mathrm{Pic}^0_{A/k}$ is surjective, it is surjective on $\overline{k}$-points. □

WARNING. This x may not be found in $A(k)$

We refine this with a partial characterization of line bundles in $\mathrm{Pic}^0_{A/k}$:

PROPOSITION 1.2.3.2. If $\mathscr{M} \in \widehat{A}(k)$, then $\phi_{\mathscr{M}} = 0$.

Proof. As we may check this equality over $\overline{k}$, without loss of generality, assume $k = \overline{k}$. Let $\mathscr{N}$ be ample so that by the preceding lemma, $\mathscr{M} \simeq t_x^*\mathscr{N} \otimes \mathscr{N}^{-1}$ for some $x \in A(k)$. By Proposition 1.1.3.9, 1.1.3.10, and 1.1.3.8 respectively,

$$\phi_{t_x^*\mathscr{N} \otimes \mathscr{N}^{-1}} = \phi_{t_x^*\mathscr{N}} + \phi_{\mathscr{N}^{-1}} = \phi_{\mathscr{N}} + [-1]_{\widehat{A}} \circ \phi_{\mathscr{N}} = 0.$$

□

We will now work towards proving the converse, which is much more difficult. Before proving it, we need the following lemma:

LEMMA 1.2.3.3. If $\mathscr{L} \not\simeq \mathscr{O}_A$ and $\phi_{\mathscr{L}} = 0$, the $H^j(A, \mathscr{L}) = 0$ for all j.

Proof. We will use induction on $j \geq 0$.

Before we start, we will make some general observations. Recall that the "Mumford construction" $\Lambda(\mathscr{M}) := m^*\mathscr{M} \otimes \mathrm{pr}_1^*\mathscr{M}^{-1} \otimes \mathrm{pr}_2^*\mathscr{M}^{-1}$ is isomorphic to the pullback $(1 \times \phi_{\mathscr{M}})^*\mathscr{P}_A$ under the map

$$A \times A \xrightarrow{1 \times \phi_{\mathscr{M}}} A \times \widehat{A}.$$

So if $\phi_{\mathscr{M}} = 0$ then $\Lambda(\mathscr{M})$ is trivial, so

$$m^*\mathscr{M} \simeq \mathrm{pr}_1^*\mathscr{M} \otimes \mathrm{pr}_2^*\mathscr{M}.$$

Then, for any k-scheme S and $(a,b) \in (A \times A)(S) = A(S) \times A(S)$, pulling back the above equation along $(a,b) : S \to A \times A$ gives

$$(a+b)^*\mathscr{M} \simeq a^*\mathscr{M} \otimes g^*\mathscr{M}.$$

In particular, we get that $[n]^*\mathscr{M} \simeq \mathscr{M}^{\otimes n}$ for all $n \in \mathbf{Z}$. So, we have

$$[-1]^*\mathscr{M} \simeq \mathscr{M}^{-1}. \tag{1.2.3-I}$$

Now, let's begin the induction, starting with the base case:

$\underline{j = 0}$: Assume there exists some non-zero global section $s \in \mathscr{M}(A)$, so $\mathscr{O}_A \to \mathscr{M}$ defined locally by $f \mapsto fs$ is a nonzero map of invertible sheaves. This dualizes to a nonzero map of invertible sheaves $\mathscr{O}_A \hookleftarrow \mathscr{M}^{-1}$, so it must be an injection and thus identifies $\mathscr{M}^{-1}$ with an ideal sheaf $\mathscr{I}_D$ for an effective Cartier divisor D. Thus, $\mathscr{M} \simeq \mathscr{I}_D^{-1} = \mathscr{O}_A(D)$.

Then the observed relation $[-1]^*\mathscr{M} \simeq \mathscr{M}^{-1}$ implies that $\mathscr{O}_A([-1]^*D) \simeq \mathscr{O}_A(-D)$, so $[-1]^*D + D$ is an effective Weil divisor linearly equivalent to 0. But a nonzero rational function whose divisor is effective must be constant, so $[-1]^*D + D = 0$ and hence $D = 0$. Thus, $\mathscr{M}$ is trivial, a contradiction.

$\underline{j > 0}$: Now, we will relate the higher cohomology groups to the lower ones. Assume $H^i(A, \mathscr{M}) = 0$ for all $i < j$. Now, factor the identity map as

$$\mathrm{id} : A \xrightarrow{(\mathrm{id},0)} A \times A \xrightarrow{m} A,$$

yeilding a factorization of the identity endomorphism

$$H^j(A, \mathscr{M}) \leftarrow H^j(A \times A, m^*\mathscr{M}) \xleftarrow{m^*} H^j(A, \mathscr{M}).$$

It now suffices to show that the middle term is 0. Since $m^*\mathscr{M} \simeq \mathrm{pr}_1^*\mathscr{M} \otimes \mathrm{pr}_2^*\mathscr{M}$, we have

$$H^j(A \times A, m^*\mathscr{M}) = H^j(A \times A, \mathrm{pr}_1^*\mathscr{M} \otimes \mathrm{pr}_2^*\mathscr{M})$$

and we may apply the Künneth formula to get that this is

$$\bigoplus_{i+i'=j} H^i(A, \mathscr{M}) \otimes_k H^{i'}(A, \mathscr{M}).$$

But either $i < j$ or $i' < j$ since $j > 0$, so by the induction hypothesis, this vanishes. □

Remark 1.2.3.4. If a line bundle $\mathscr{L}$ satisfies $[-1]^*\mathscr{L} \simeq \mathscr{L}$, we call $\mathscr{L}$ **symmetric** and similarly if $[-1]^*\mathscr{L} \simeq \mathscr{L}^{-1}$ we say $\mathscr{L}$ is **anti-symmetric**. Equation 1.2.3-I in the above proof shows that all line bundles in Pic^0 are anti-symmetric.

We can now prove the converse to Proposition 1.2.3.2, giving us a complete characterization of the line bundles in $\mathrm{Pic}^0_{A/k}$.

Theorem 1.2.3.5. If $\phi_{\mathscr{L}} = 0$, then $\mathscr{L}_{\overline{k}} \simeq \phi_{\mathscr{N}}(x)$ for some $x \in A(\overline{k})$. I.e. $\mathscr{L} \mapsto \phi_{\mathscr{L}}$ induces an

injective homomorphism

$$\mathrm{NS}(A) := \mathrm{Pic}_{A/k}(\overline{k}) \Big/ \mathrm{Pic}^0_{A/k}(\overline{k}) \hookrightarrow \mathrm{Hom}_{\overline{k}}(A_{\overline{k}}, \widehat{A}_{\overline{k}})$$

or equivalently ([Stacks, Tag 04XV]) a closed k-subgroup inclusion between étale k-groups

$$\mathrm{Pic}_{A/k} \Big/ \mathrm{Pic}^0_{A/k} \to \underline{\mathrm{Hom}}(A, \widehat{A}).$$

Remark 1.2.3.6. The group $\mathrm{NS}(A)$ is called the Néron-Severi group and is important in its own right. It will be explored in section 1.2.5.

This characterization essentially says that $\phi_{\mathscr{L}}$ detects whether or not one can "continuously deform" a line bundle to another in the moduli space of line bundles on A.

Proof. We may assume without loss of generality that $k = \overline{k}$. The idea is to make a family out of the collection of line bundles $\{t_a^*\mathscr{N} \otimes \mathscr{N}^{-1} \otimes \mathscr{L}^{-1}\}_{a \in A(k)}$. To make this formal, what we want is a line bundle $\mathscr{K}$ on $A \times A$ such that $\mathscr{K}|_{\{a\}\times A} \simeq t_a^*\mathscr{N} \otimes \mathscr{N}^{-1} \otimes \mathscr{L}^{-1}$ for all $a \in A(k)$, and we will get a contradiction if these are all nontrivial. Let $\mathscr{K} := \Lambda(\mathscr{N}) \otimes \mathrm{pr}_2^*\mathscr{L}^{-1}$ (where $\Lambda(\mathscr{N}) = m^*\mathscr{N} \otimes \mathrm{pr}_1^*\mathscr{N}^{-1} \otimes \mathrm{pr}_2^*\mathscr{N}^{-1}$, the so-called "Mumford construction"). It is clear that $\mathscr{K}|_{\{a\}\times A} = t_a^*\mathscr{N} \otimes \mathscr{N}^{-1} \otimes \mathscr{L}^{-1}$ since pulling back to $\{a\} \times A$ turns m^* into t_a^*, kills pr_1^*, and turns pr_2^* into id^*.

Now we study $H^j(A \times A, \mathscr{K})$ using the Leray spectral sequences associated to the two maps $\mathrm{pr}_1, \mathrm{pr}_2 : A \times A \rightrightarrows A$. Assume $\mathscr{K}|_{\{a\}\times A} \not\simeq \mathscr{O}_A$ for all $a \in A(k)$. We'll deduce from this assumption that we must have $H^j(A \times A, \mathscr{K}) = 0$ for all j (using $\phi_{\mathscr{L}} = 0$), and get a contradiction from this.

From our first spectral sequence, we have $H^i(A, R^j \mathrm{pr}_{1,*}\mathscr{K}) \Rightarrow H^{i+j}(A \times A, \mathscr{K})$. We will show that $R^j \mathrm{pr}_{1,*}\mathscr{K} = 0$ for all j. By the theorem on cohomology and base change (which is applicable since pr_1 is proper and flat), it suffices to show that $H^j(A, \mathscr{K}|_{\{a\}\times A}) = 0$ for all j and all $a \in A(k)$. By design, $\mathscr{K}|_{\{a\}\times A} \simeq t_a^*\mathscr{N} \otimes \mathscr{N}^{-1} \otimes \mathscr{L}^{-1}$ is nontrivial, and since $\phi_{\mathscr{M}}$ is a "discrete" invariant of $\mathscr{M}$, we have $\phi_{\mathscr{K}|_{\{a\}\times A}} = \phi_{\mathscr{L}^{-1}} = -\phi_{\mathscr{L}}$, which vanishes by hypothesis on $\mathscr{L}$. By Lemma 1.2.3.3, we have that $H^j(A, \mathscr{M}) = 0$ for all j. Thus, $R^j \mathrm{pr}_{1,*}\mathscr{K} = 0$ for all j, so the Leray spectral sequence implies that $H^j(A \times A, \mathscr{K}) = 0$ for all j.

Now, let's study the spectral sequence for the *other* projection:

$$H^i(A, R^j \mathrm{pr}_{2,*}\mathscr{K}) \Rightarrow H^{i+j}(A \times A, \mathscr{K}) = 0.$$

We have $\mathscr{K}|_{A\times\{a\}} = \phi_{\mathscr{N}}(a)$, and $\mathscr{K} = \Lambda(\mathscr{N}) \otimes \mathrm{pr}_2^*\mathscr{L}^{-1}$. This has vanishing ϕ by Proposition 1.2.3.2. Thus, again by Lemma 1.2.3.3, we have vanishing of fibral cohomology: $H^j(A \times \{a\}, \mathscr{K}|_{A\times\{a\}}) = 0$ for all j. Therefore, by cohomology and base change, the stalks of the higher direct images at such points vanish: $R^j \mathrm{pr}_{2,*}\mathscr{K}_a = 0$ for all j when $a \notin \ker \phi_{\mathscr{N}}$.

But, by the ampleness of $\mathscr{N}$, $\ker \phi_{\mathscr{N}}$ is finite (as noted at the beginning of Section 1.2.2). Thus, the coherent sheaves $R^j \mathrm{pr}_{2,*}\mathscr{K}$ on A have support contained in the finite k-scheme $\ker \phi_{\mathscr{N}}$. These coherent sheaves are then the pushforwards to A from k-finite closed subschemes, so $H^i(A, R^j \mathrm{pr}_{2,*}\mathscr{K}) = 0$ for all $i > 0$ and all j. Thus, the Leray spectral sequence vanishes away from the row consisting of terms

$$H^0(A, R^j \mathrm{pr}_{2,*}\mathscr{K})$$

whose members are the global sections of the coherent sheaves $R^j \operatorname{pr}_{2,*} \mathscr{K}$ with k-finite support, so $R^j \operatorname{pr}_{2,*} \mathscr{K}$ vanishes for all j. But, by the theorem on cohomology and base change applied to the proper flat pr_2, it follows via descending induction from larger cohomological degrees that the fibral cohomologies $H^j(A, \mathscr{K}|_{\operatorname{pr}_2^{-1}(a)})$ vanish for all $a \in A(k)$ and all $j \geq 0$. But, taking $j = 0$ and $a = 0$ gives that $H^0(A, \mathscr{O}_A) = 0$ which is false. We thus have a contradiction. □

§1.2.4 Dual Morphisms

We now want to see how morphisms behave under our duality theory.

Definition 1.2.4.1. Let $f : A \to B$ be a map of abelian varieties. We define the **dual morphism** $\widehat{f} : \widehat{B} \to \widehat{A}$ functorially by $\mathscr{L} \rightsquigarrow f_T^* \mathscr{L}$ on T-valued points (for T a k-scheme). Note that this is just the restriction of the pullback map $f^* : \operatorname{Pic}_{B/k} \to \operatorname{Pic}_{A/k}$.

Theorem 1.2.4.2. Let $f, g : A \rightrightarrows B$ be two morphisms of abelian varieties. Then $(f+g)^\wedge = \widehat{f} + \widehat{g}$ as maps $\widehat{B} \to \widehat{A}$.

Proof. Without loss of generality, we may assume $k = \overline{k}$ and compare on k-points, since we can check equality of maps via equality on $\overline{k}$-points. So, we want to show that for $\mathscr{M} \in \widehat{B}(k) = \operatorname{Pic}^0_{B/k}(k)$ (or equivalently $\mathscr{M}$ such that $\phi_{\mathscr{M}} = 0$ by Theorem 1.2.3.5),

$$(f+g)^* \mathscr{M} = f^* \mathscr{M} \otimes g^* \mathscr{M}$$

in $\operatorname{Pic}(B) = \operatorname{Pic}_{B/k}(k)$.

But, the left side is the pullback of $m_B^* \mathscr{M}$ under $(f, g) : A \to B \times B$ and the right side is the pullback of $\operatorname{pr}_1^* \mathscr{M} \otimes \operatorname{pr}_2^* \mathscr{M}$ under $(f, g) : A \to B \times B$. But by the Theorem of the Square (1.1.3.7), whenever $\phi_{\mathscr{M}} = 0$, $m_B^* \mathscr{M} \simeq \operatorname{pr}_1^* \mathscr{M} \otimes \operatorname{pr}_2^* \mathscr{M}$. □

We now have the important consequence:

Corollary 1.2.4.3. We have $\widehat{[n]_A} = [n]_{\widehat{A}}$.

Warning. Both Theorem 1.2.4.2 and Corollary 1.2.4.3 are false on the entire Picard scheme. As an example of the failure of Theorem 1.2.4.2 on the entire group, for $k = \mathbf{C}$, we see that

$$\operatorname{Pic} / \operatorname{Pic}^0 \hookrightarrow H^2(A(\mathbf{C}), \underline{\mathbf{Z}}(1)) = \wedge^2(H^1(A(\mathbf{C}), \mathbf{Z})$$

upon choosing a basis of $\mathbf{Z}(1)$, and this ambient group is not additive in the lattice. Essentially, the component group of the Picard scheme is a "quadratic" functor whereas its identity component behaves additively.

For an example of a failure of the analogue of Corollary 1.2.4.3 we know that the component group can be identified with a subgroup of $\operatorname{Hom}(A, \widehat{A})$ on which the the combined functorial effect of $[n]_A$ and $[n]_A^\wedge = [n]_{\widehat{A}}$ is multiplication by n^2, not n (see Lemma 1.2.5.7 below).

§1.2.5 The Néron-Severi Group

So far we have only focused on $\mathrm{Pic}^0_{A/k}$. However, the rest of the Picard scheme, packaged as the Néron-Severi group, also plays a very important role in the study of abelian varieties!

As stated in §1.1.2, if X is a proper, geometrically connected, geometrically reduced k-scheme then $\mathrm{Pic}_{X/k}$ is a k-group scheme locally of finite type. We call its geometric component:

Definition 1.2.5.1. The **Néron-Severi** group of X is

$$\mathrm{NS}(X) = \mathrm{Pic}_{X/k}(\overline{k}) \Big/ \mathrm{Pic}^0_{X/k}(\overline{k}).$$

This is the group of geometric points of the étale k-group $\mathrm{Pic}_{X/k} / \mathrm{Pic}^0_{X/k}$ since the formation of Picard schemes commutes with extension of the ground field. Geometrically, $\mathrm{NS}(X)$ is the group of line bundle on $X_{\overline{k}}$ up to algebraic equivalence, where we define algebraic equivalence as follows:

Definition 1.2.5.2. Let $\mathscr{L}_1, \mathscr{L}_2$ be line bundles on $X_{\overline{k}}$. We say that $\mathscr{L}_1$ and $\mathscr{L}_2$ are **algebraically equivalent** if there exists a connected $\overline{k}$-scheme S, and points $s_1, s_2 \in S(K)$ for some extension field $K/\overline{k}$, and a line bundle $\mathscr{N}$ on $X \times S$ such that $\mathscr{N}_{s_i} = (\mathscr{L}_i)_K$.

Remark 1.2.5.3. This is equivalent to $\mathscr{L}_1$ and $\mathscr{L}_2$ being contained in the same geometric connected component of $\mathrm{Pic}_{X/k}$ because any line bundle $\mathscr{N}$ on $X \times S$ is classified (up to twisting by an element of $f^* \mathrm{Pic}(S)$) by a $\overline{k}$-morphism $S \to \mathrm{Pic}_{X_{\overline{k}}} = (\mathrm{Pic}_{X/k})_{\overline{k}}$, and such a map must land in a connected component since S is connected.

It turns out that for abelian varieties, ampleness of a line bundle $\mathscr{L}$ can be detected on the level of the Néron-Severi group as follows:

Proposition 1.2.5.4. For any line bundles $\mathscr{L}, \mathscr{N}$ on A with $\mathscr{N} \in \mathrm{Pic}^0_{A/k}(j)$, $\mathscr{L}$ is ample if and only if $\mathscr{L} \otimes \mathscr{N}$ is ample; i.e., ampleness only depends on $[\mathscr{L}] \in \mathrm{NS}(A)$.

Remark 1.2.5.5. The Nakai-Moishezon criterion for geometrically integral, projective k-schemes X gives a characterization of when a line bundle $\mathscr{L}$ is ample on X in terms of intersection theoretic properties against closed subvarieties of $X_{\overline{k}}$. By local constancy of this data (due to the local constancy of the Hilbert polynomial), it is clear that ampleness only depends on its class in the Néron-Severi group. The proof we give does not use this, but instead uses the nice properties of abelian varieties we have already been building!

Proof. It suffices to prove the forward direction since if $\mathscr{L}$ ample $\implies \mathscr{L} \otimes \mathscr{N}$ ample, then the other direction can be obtained by tensoring with $\mathscr{N}^{-1}$. Note that ampleness is insensitive to ground field extension, so we may and do assume that $k = \overline{k}$. Then, we know that $\phi_{\mathscr{L}} : A(k) \to \overline{A}(k)$ is surjective since $\mathscr{L}$ is ample, so $\mathscr{N} = \phi_{\mathscr{L}}(x) = t_x^* \mathscr{L} \otimes \mathscr{L}^{-1}$ for some $x \in A(k)$. Hence, $\mathscr{L} \otimes \mathscr{N} \simeq t_x^* \mathscr{L}$, and this is ample since t_x is an automorphism. □

In fact, we can go further and show that $\phi_{\mathscr{L}}$ can "detect" whether $\mathscr{L}$ is ample or not. Our strategy is to first attempt to recover $\mathscr{L}$ from $\phi_{\mathscr{L}}$ by pulling back the Poincaré bundle along

$$(1, \phi_{\mathscr{L}}) = (1 \times \phi_{\mathscr{L}}) \circ \Delta : A \to A \times \widehat{A}$$

to get $\Delta^*(1 \times \phi_{\mathscr{L}})^* \mathscr{P}_A$. Now the bundle $(1 \times \phi_{\mathscr{L}})^* \mathscr{P}_A$ on $A \times A$ is just the Mumford construction $\Lambda(\mathscr{L}) = m^*\mathscr{L} \otimes \mathrm{pr}_1^* \mathscr{L}^{-1} \otimes \mathrm{pr}_2^* \mathscr{L}^{-1}$. So, when we restrict to the diagonal, we get $[2]^*\mathscr{L} \otimes \mathscr{L}^{-2}$. The following Lemma 1.2.5.7 then gives us that $[2]^*\mathscr{L} \otimes \mathscr{L}^{-2} \simeq \mathscr{L}^2 \otimes \mathscr{N}$, $\mathscr{N} \in \mathrm{Pic}^0_{A/k}(k)$, so we can recover the class of $\mathscr{L}^2$ in $\mathrm{NS}(A)$ from $\phi_{\mathscr{L}}$. Thus by Proposition 1.2.5.4, we may detect ampleness with $\phi_{\mathscr{L}}$ (as $\mathscr{L}^2$ is ample if and only if $\mathscr{L}$ is). In particular, $\mathscr{L}$ is ample if and only if $(1, \phi_{\mathscr{L}})^* \mathscr{P}_A$ is ample.

Remark 1.2.5.6. In accordance with our running analogy to linear algebra, we note that this extra factor of 2 coming from the conversion between $\mathscr{L}$ and $\phi_{\mathscr{L}}$ is similar to that coming from a quadratic from q and its bilinear form B_q.

Lemma 1.2.5.7. For all $\mathscr{L} \in \mathrm{Pic}(X)$,

$$[n]^*\mathscr{L} \simeq \mathscr{L}^{n^2} \otimes (\text{something in } \mathrm{Pic}^0(A))$$

i.e., the effect of $[n]$ on $\mathrm{NS}(A)$ is multiplication by n^2.

Proof. We showed in Theorem 1.1.4.1 that

$$[n]^*\mathscr{L} \simeq \mathscr{L}^{n^2} \otimes (\mathscr{L} \otimes [-1]^*\mathscr{L}^{-1})^{\frac{n^2-n}{2}}.$$

So we want to show that $\mathscr{L} \otimes [-1]^*\mathscr{L}^{-1} \in \mathrm{Pic}^0_{A/k}(k)$, or equivalently (by 1.2.3.5) that

$$\phi_{\mathscr{L}} = \phi_{[-1]^*\mathscr{L}}$$

as maps $A \to \widehat{A}$. As we may check equality of maps on $\overline{k}$-points, we may assume without loss of generality that $k = \overline{k}$ and compute on $x \in A(k)$:

$$\begin{aligned}
\phi_{[-1]^*\mathscr{L}}(x) &= t_x^*([1]^*\mathscr{L}) \otimes [-1]^*\mathscr{L}^{-1} \\
&\simeq [-1]^*(t_{-x}^*\mathscr{L}) \otimes [-1]^*\mathscr{L}^{-1} \\
&\simeq [-1]^*(t_{-x}^*\mathscr{L} \otimes \mathscr{L}^{-1}) \\
&\simeq \widehat{[-1]}\phi_{\mathscr{L}}(-x).
\end{aligned}$$

So we have reduced to showing $\widehat{[-1]}\phi_{\mathscr{L}}(-x) = \phi_{\mathscr{L}}(x)$. But $\widehat{[n]_A} = [n]_{\widehat{A}}$ for any integer n by Corollary 1.2.4.3. Thus, $\widehat{[-1]}\phi_{\mathscr{L}}(-x) = -\phi_{\mathscr{L}}(-x) = \phi_{\mathscr{L}}(x)$. □

We have the following corollary, which will be useful later.

Corollary 1.2.5.8. For any $\mathscr{L} \in \mathrm{Pic}_{A/k}(\overline{k})$, we may write $\mathscr{L} = \mathscr{L}_s \otimes \mathscr{L}_0$ where $\mathscr{L}_s$ is symmetric (i.e., $[-1]^*\mathscr{L}_s = \mathscr{L}_s$) and $\mathscr{L}_0 \in \mathrm{Pic}^0_{A/k}(\overline{k})$.

Proof. Firstly, since $[2]_{\widehat{A}}$ is surjective and thus surjective on $\overline{k}$-points, for any $\mathscr{N}' \in \mathrm{Pic}^0_{A/k}(\overline{k})$, there exists some $\mathscr{N} \in \mathrm{Pic}^0_{A/k}(\overline{k})$ such that $\mathscr{N}^2 = \mathscr{N}'$. Then by Lemma 1.2.5.7, $[-1]^*\mathscr{L} \simeq \mathscr{L} \otimes \mathscr{L}_0'$ for $\mathscr{L}_0' \in \mathrm{Pic}^0_{A/k}(k)$.

Let $\mathscr{L}_0 \in \mathrm{Pic}^0_{A/k}(\overline{k})$ be such that $\mathscr{L}_0^2 = \mathscr{L}_0'$ so that $[-1]^*\mathscr{L} \simeq \mathscr{L} \otimes \mathscr{L}_0^2$. Then for $\mathscr{L}_s = \mathscr{L} \otimes \mathscr{L}_0$, we have that

$$[-1]^*\mathscr{L}_s = [-1]^*\mathscr{L} \otimes [-1]^*\mathscr{L}_0^{-1} = \mathscr{L} \otimes \mathscr{L}_0^2 \otimes \mathscr{L}_0^{-1} = \mathscr{L} \otimes \mathscr{L}_0 = \mathscr{L}_s$$

as $[-1]^*$ acts by inversion by Corollary 1.2.4.3> □

Now we will discuss some deep results about $\mathrm{NS}(X)$ for general X. Note that this will not be used in the rest of the paper. The first of which is the Néron-Lang "Theorem of the Base," which says that $\mathrm{NS}(X)$ is *always finitely generated* over $\mathbf{Z}$ (in particular, its torsion subgroup if finite). We may now ask whether we can characterize the bundles corresponding to torsion in the Néron-Severi group. It turns out that these are precisely the line bundles with are undetectable by intersection theory:

DEFINITION 1.2.5.9. We say that a line bundle $\mathscr{L}$ on $X_{\overline{k}}$ is **numerically equivalent to** 0 if for all integral curves $C \subset X_{\overline{k}}$, the intersection number $C \cdot \mathscr{L} := \deg_C(\mathscr{L}|_C)$ vanishes (see [BLR90b, §9.1] for several definitions of degree of a line bundle on general integral proper curves over fields).

We have:

THEOREM 1.2.5.10. We have the following results:

(1) (Theorem of the Base) $\mathrm{NS}(X)$ is finitely generated over Z.

(2) If $\mathrm{Pic}^\tau_{X/k}$ denotes the open and closed subgroup scheme of $\mathrm{Pic}_{X/k}$ corresponding to $\mathrm{NS}(X)_{\mathrm{tor}}$ then $\mathrm{Pic}^\tau_{X/k}(\overline{k})$ consists exactly of those $\mathscr{L}$ on $X_{\overline{k}}$ that are numerically equivalent to 0.

Proof. Part (1) in [SGA6, Exp. XIII, Thm. 5.1] and part (2) in [SGA6, Exp XIII, Thm. 4.6]. □

Additionally, in SGA6, a relative version is proved: for any proper flat finitely presented morphism of schemes $X \to S$, the rank and torsion of the Néron-Severi groups are bounded locally on S as one varies through the geometric fibers, whereas the locus where the rank or torsion subgroup are equal to some specified value are typically non-constructible.

Returning to our case of abelian varieties, for any abelian variety A, we have an embedding

$$\mathrm{NS}(A) \hookrightarrow \mathrm{Hom}_{\overline{k}}(A_{\overline{k}}, \widehat{A}_{\overline{k}})$$

given by $[\mathscr{L}] \mapsto \phi_{\mathscr{L}}$. Since $\mathrm{Hom}_{\overline{k}}(A_{\overline{k}}, \widehat{A}_{\overline{k}})$ is torsion-free, it follows that $\mathrm{NS}(A)$ is torsion-free. In §2.1.2 (Theorem 2.1.2.9), we will show that this Hom group is finitely generated, so we get finite-generation of $\mathrm{NS}(A)$ for free. To do so, we will need a discussion of duality of morphisms between abelian varieties (tackled in §1.2.4 and §1.2.6) and a robust theory of isogenies (to be tackled in §2.1.2).

We now return to our discussion of duality on Abelian varieties.

§1.2.6 More on $\mathscr{P}_A$ and the Duality Hypothesis

The goal of this section is to prove the duality hypothesis, i.e. that there is a canonical isomorphism $\iota_A : A \xrightarrow{\sim} \widehat{\widehat{A}}$. Along the way, we will prove some other important results, most notably that for an isogeny $f : A \to B$ of abelian varieties, $\widehat{f}$ is an isogeny of the same degree as f.

We will first introduce some lemmas which we will use in the following theorem. The proofs of these lemmas are long homological calculations, but are also fairly straightforward, so the proofs are omitted. Maybe add proofs in final draft

LEMMA 1.2.6.1. Let $f : X \to Y$ be a proper morphism of noetherian schemes with $Y = \operatorname{Spec} A$ affine, and $\mathscr{F}$ a coherent sheaf on X, flat over Y. There is a finite complex

$$K^\bullet : 0 \to K^0 \to K^1 \to \cdots \to K^n \to 0$$

of finitely generated projective A-modules and an isomorphism of functors

$$H^p(X \times_Y \operatorname{Spec} B, \mathscr{F} \otimes_A B) \simeq H^p(K^\bullet \otimes_A B), \quad (p \geq 0)$$

on the category of A-algebras B.

Proof. See [Mum74, §5, Theorem 2] □

LEMMA 1.2.6.2. Let R be Cohen-Macaulay (e.g. regular) local ring of dimension $g \geq 0$. Let $K^\bullet$ be a bounded complex of R-flat (e.g. finite free) R-modules

$$0 \to K^0 \to \cdots \to K^n \to 0$$

such that all the $H^j(K^\bullet)$ have finite R-length. Then $H^j(K^\bullet) = 0$ for all $0 \leq j < g$.

Proof. See [Mum74, §13, Lemma 1]. □

Firstly, we will examine the cohomology of $\mathscr{P}_A$. The proceeding calculation will be quite long.

THEOREM 1.2.6.3. $H^n(A \times \widehat{A}, \mathscr{P}_A) = \begin{cases} 0 & n \neq g, \\ 1 & n = g. \end{cases}$

Proof. We will use the Leray spectral sequence to relate the cohomology of the fibers of the projection $A \times \widehat{A} \xrightarrow{\pi} \widehat{A}$ to the total cohomology. Tautologically, for $\mathscr{L}_{\widehat{a}}$ associated to $\widehat{a} \in \widehat{A}(k)$, since $\mathscr{P}_A$ is defined as the universal bundle,

$$\mathscr{P}_A|_{\pi^{-1}(\widehat{a})} \simeq \mathscr{L}_{\widehat{a}}$$

(where $\pi^{-1}(\widehat{a}) \simeq A$ since π is a projection, and so we treat $\mathscr{L}_{\widehat{a}}$ as a bundle on this). The Leray spectral sequence may be described as follows:

$$E_2^{i,j} = H^i(\widehat{A}, R^j\pi_*\mathscr{P}_A) \Rightarrow H^{i+j}(A \times \widehat{A}, \mathscr{P}_A).$$

We will show that $R^j\pi_*\mathscr{P}_A$ is supported on $\widehat{0}$ for all j (where $\widehat{0}$ is the point of $\widehat{A}$ corresponding to $\mathscr{O}_A$). Since π is proper and flat, by cohomology and base changed, it's enough to show that $H^j(A, \mathscr{P}_A|_{A\times\{\widehat{a}\}} = 0$ for all j and all $\widehat{a} \in \widehat{A}(\overline{k}) - \{\widehat{0}\}$. But since $\mathscr{P}_A|_{A\times\{\widehat{a}\}} = \mathscr{L}_{\widehat{a}}$, it suffices to show $H^j(A, \mathscr{L}) = 0$ for all $\mathscr{L} \in \operatorname{Pic}^0_{A/k}(k)$ such that $\mathscr{L} \not\simeq \mathscr{O}_A$. But we already proved this in Lemma 1.2.3.3.

Thus, $R^j\pi_*\mathscr{P}_A$ is a coherent sheaf on $\widehat{A}$ vanishing on $\widehat{A} - \{\widehat{0}\}$, so it arises from a coherent sheaf on an infinitesimal closed subscheme $Z_j \subset \widehat{A}$ supported at $\widehat{0}$. Thus, for dimension reasons, $H^i(\widehat{A}, R^j\pi_*\mathscr{P}_A)$

vanishes for $i > 0$. Hence, the only possible non-zero terms on the E_2 page are contained in the column $i = 0$. Thus the spectral sequence degenerates to give

$$H^n(A \times \widehat{A}, \mathscr{P}_A) = (R^n\pi_*\mathscr{P}_A)_{\widehat{0}}.$$

For $n > g$, this is 0 by the Theorem on Formal Functions because the fibers have dimension g. Additionally, for all n, this is k-finite. What remains to show is that $(R^n\pi_*\mathscr{P}_A)_{\widehat{0}}$ vanishes for $n < g$ and is k for $n = g$.

Let $R = \mathscr{O}_{\widehat{A},\widehat{0}}$, a regular local ring. Consider the base change

$$\begin{array}{ccc} A_R & \longrightarrow & A \times \widehat{A} \\ \downarrow & & \downarrow \\ \operatorname{Spec} R & \xrightarrow{\text{flat}} & \widehat{A}. \end{array}$$

By compatibility of cohomology with flat base change, we have

$$R^j\pi_*(\mathscr{P}_A)_{\widehat{0}} = H^j(A_R, (\mathscr{P}_A)_{A_R}).$$

Since $(\mathscr{P}_A)_{A_R}$ is coherent and A_R is proper over R, $H^j(A_R, (\mathscr{P}_A)_{A_R})$ is finitely generated over R.

Now, by Lemma 1.2.6.1 below and since a projective module over a local ring is free (see [Stacks, Tag 058Z]), there is a a bounded complex of finite free R-modules

$$K^\bullet : 0 \to K^0 \to \cdots \to K^N \to 0$$

and an isomorphism of functors

$$H^p(A_R \times_R \operatorname{Spec} S, (\mathscr{P}_A)_{A_R}) \simeq H^p(K^\bullet \otimes_R S)$$

for all R-algebras S.

Now we have a bounded complex $K^\bullet$ of finite free R-modules with the property that $H^j(K^\bullet)$ are all k-finite and vanish for $j > g = \dim R$. By Lemma 1.2.6.2, the homology automatically vanishes in dimension below g as well. To conclude with the 1-dimensionality in degree g, we're going to use the universality of the Poincaré bundle relative to $\widehat{A}$ as a moduli scheme (by using infinitesimal considerations).

To recap, we now have a complex

$$\underbrace{0 \to K^0 \to \cdots \to}_{\text{exact by Lemma}} K^g \to K^{g+1} \underbrace{\to \cdots \to K^N \to 0}_{\text{exact from setup}}.$$

Now, if $N \geq g$, $\ker(K^N \twoheadrightarrow K^{N+1})$ is finite and R-flat (as it is the kernel of a surjective map between flat modules) and thus free, and its formation commutes with base change on R, so we can replace K^{N+1} by 0 and K^N with that kernel to produce a chain complex with the same cohomology.

Thus, without loss of generality, we may assume $N = g$, so we have reduced to another sequence

$$0 \to K^0 \to \cdots \to K^{g-1} \to (K')^g \to 0$$

with the same cohomology. Adding the cokernel to the end of our sequence (and renaming back to $K^\bullet$ for ease of notation), we have built a finite free resolution of the R-module $H^g(K^\bullet) = R^g\pi_*(\mathscr{P}_A)_{\widehat{0}}$, which we name K^{g+1}:

$$0 \to K^0 \to \cdots \to K^g \to \underbrace{R^g\pi_*(\mathscr{P}_A)_{\widehat{0}}}_{H^g(K^\bullet)} \to 0$$

and it remains to prove that $R^g\pi_*(\mathscr{P}_A)_{\widehat{0}}$ has length 1. Dualizing (applying $\mathrm{Hom}_R(-,R)$ to) the non-extended sequence, we obtain

$$0 \to \widehat{K^g} \to \cdots \to \widehat{K^0} \to 0$$

whose final cohomology computes $\mathrm{Ext}^\bullet_R(R^g\pi_*(\mathscr{P}_A)_{\widehat{0}}, R)$. But $R^g\pi_*(\mathscr{P}_A)_{\widehat{0}}$ has finite R-length, so the Ext-modules also have finite R-length. Thus, we may apply Lemma 1.2.6.2 again to get that $\widehat{K^\bullet}$ is exact before the final term. Letting $M = \mathrm{Ext}^g(R^g\pi_*(\mathscr{P}_A)_{\widehat{0}}, R)$, the cohomology at the final term, so we have a finite free resolution

$$0 \to \widehat{K^g} \to \cdots \to \widehat{K^1} \to M \to 0.$$

Thus, as with last times, we may compute $\mathrm{Ext}^\bullet_R(M,R)$ by applying $\mathrm{Hom}_R(-,R)$ again to the non-extended resolution. But, this undoes the original dualization, as each K^j is a finite free R-module, so $R^g\pi_*(\mathscr{P}_A)_{\widehat{0}} = \mathrm{Ext}^g_R(M,R)$. To show that it has length 1, we will analyze M.

It is a general fact (see [Stacks, Tag 0DW6]) that for Gorenstein (e.g. regular) local rings of dimension g that $\mathrm{Ext}^g_R(R/\mathfrak{m}_R, R)$ is 1-dimensional. So, it is enough to show that M is 1-dimensional over k. We know that M is finitely generated over R, and we claim that it is non-zero and generated by a single element. By definition, $M = \mathrm{coker}(\widehat{K^1} \to \widehat{K^0})$ of a map of finite free modules, so

$$M/\mathfrak{m}M = \mathrm{coker}((K^1 \bmod \mathfrak{m})^\wedge \to (K^0 \bmod \mathfrak{m})^\wedge)$$

since tensoring is right exact, and dualizing commutes with scalar extension for finite free modules. But, this is just $\ker(K^0 \otimes R/\mathfrak{m} \to K^1 \otimes R/\mathfrak{m})$, which is

$$H^0(K^\bullet \otimes R/\mathfrak{m}) = H^0(A, \underbrace{\mathscr{P}_A|_{A\times\widehat{0}}}_{=\mathscr{O}_A}) = k;$$

by the universal property of the Poincaré bundle. By Nakayama's lemma, M must be non-zero and principally generated, so $M = R/J$ for J a proper ideal $J \subsetneq R$. This is the annihilator ideal for the finite-length M, so J is $\mathfrak{m}$-primary. It remains to check $J \supset \mathfrak{m}$.

For any $\mathfrak{m}$-primary $J' \subset R$,

$$\begin{aligned}\mathrm{Hom}(M, R/J') &= \ker((K^0 \bmod J') \to K^1 \bmod J')) \\ &= H^0(K^\bullet \bmod J') \\ &= H^0(A \times \mathrm{Spec}(R/J'), \mathscr{P}_A|_{A\times\mathrm{Spec}(R/J')}).\end{aligned}$$

Let's apply this to both $J' = J$ and to $J' = \mathfrak{m}$.

So we get a commutative diagram of natural maps

$$\begin{array}{ccc}\mathrm{Hom}(M, R/J) & \xrightarrow{\sim} & \Gamma(A_R \times \mathrm{Spec}(R/J), \mathscr{P}_A) \\ \downarrow & & \downarrow \\ \mathrm{Hom}(M, R/\mathfrak{m}) & \xrightarrow{\sim} & \Gamma(A, \mathscr{O}_A)\end{array}$$

The left side is trivially surjective as $J \subset \mathfrak{m}$, so the right side is surjective. This shows that the generating section of $\Gamma(A, \mathscr{O}_A)$ lifts, so the line bundle $\mathscr{P}_A|_{A_R\times\mathrm{Spec}(R/J)}$ on the infinitesimal thickening $A_R \otimes_R (R/J) = A \otimes_k (R/J)$ of A has a global section which is non-zero in all fibers, i.e., is a generating section. Thus, this line bundle is trivial.

By the universal property of $\mathscr{P}_A$, it follows that the infinitesimal closed immersion $\mathrm{Spec}(R/J) \to \widehat{A}$

factors through the inclusion $\mathrm{Spec}(R/\mathfrak{m}) \to \widehat{A}$, so $J \supset \mathfrak{m}$. Thus, $M = k$, so $\mathrm{Ext}^q_R(M, R)$ is and thus $R^g\pi_*(\mathscr{P}_A)_{\widehat{0}} = k$ as required. □

COROLLARY 1.2.6.4. $\chi(\mathscr{P}_A) = (-1)^g$.

We can now study the dual morphism more carefully. We will begin with a short lemma which will help us study the degree of the dual morphism to an isogeny.

LEMMA 1.2.6.5. Let $f : A \to B$ be a homomorphism between abelian varieties over k, and $\widehat{f} : \widehat{B} \to \widehat{A}$ the dual morphism. Then the formation of $\widehat{f}$ commutes with extension of the base field, and it is uniquely characterized by the property that

$$(1_A \times \widehat{f})^* \mathscr{P}_A \simeq (f \times 1_{\widehat{B}})^* \mathscr{P}_B.$$

Proof. For the first part, let K/k be a field extension. Note that for any line bundle $\mathscr{L}$ on B_T, $\widehat{f}^*_{T_K}\mathscr{L}_K = (\widehat{f}^*_T\mathscr{L})_K$, so we are done. We will now prove the second claim. First, we check that both line bundles are trivial on $\{0\} \times \widehat{B}$:

$$\begin{aligned}\left((1_A \times \widehat{f})^* \mathscr{P}_A\right)|_{\{0\}\times\widehat{B}} &\simeq (1_A \times \widehat{f})^*(\mathscr{P}_A|_{\{0\}\times\widehat{f}(\widehat{B})}) \\ &\simeq (1_A \times \widehat{f})^*(\mathscr{O}_{\widehat{f}(\widehat{B})}) \\ &\simeq \mathscr{O}_{\widehat{B}}\end{aligned}$$

and

$$\begin{aligned}\left((f \times 1_{\widehat{B}})^* \mathscr{P}_B\right)|_{\{0\}\times\widehat{B}} &\simeq (f \times 1_{\widehat{B}})^*(\mathscr{P}_B|_{\{0\}\times\widehat{B}}) \\ &\simeq (f \times 1_{\widehat{B}})^*(\mathscr{O}_{\widehat{B}}) \\ &\simeq \mathscr{O}_{\widehat{B}}\end{aligned}$$

(note that by abuse of notation, we are identifying $\mathscr{O}_{\widehat{B}}$ with the trivial bundle on $\{0\} \times \widehat{B}$). Thus, by the Seesaw theorem (1.1.3.1), both line bundles are parameterized by maps $\widehat{B} \to \mathrm{Pic}^0_{A/k} = \widehat{A}$ (note that it factors through the identity component because $\widehat{B}$ is connected and it clearly maps $\mathscr{O}_B$ to $\mathscr{O}_A$). We claim that in both cases this map is $\widehat{f} : \widehat{B} \to \widehat{A}$. As this is simply checking equality of maps, we may check over geometric points. So, without loss of generality (using that we proved its formation commutes with base change), assume $k = \overline{k}$. Let $\widehat{b} \in \widehat{B}(k)$ and $\widehat{a} = \widehat{f}(\widehat{b})$ and let $\mathscr{L}_{\widehat{a}} \in \mathrm{Pic}(A)$ be a line bundle corresponding to $\widehat{a}$. We have

$$\begin{aligned}\left((1_A \times \widehat{f})^* \mathscr{P}_A\right)|_{A\times\{\widehat{b}\}} &\simeq (1_A \times \widehat{f})^*(\mathscr{P}_A|_{A\times\{\widehat{a}\}}) \\ &\simeq (1_A \times \widehat{f})^*(\mathscr{L}_{\widehat{a}}) \\ &\simeq \mathscr{L}_{\widehat{a}}\end{aligned}$$

and

$$\begin{aligned}\left((f \times 1_{\widehat{B}})^* \mathscr{P}_B\right)|_{A\times\{\widehat{b}\}} &\simeq (f \times 1_{\widehat{B}})^*(\mathscr{P}_B|_{f(A)\times\{\widehat{b}\}}) \\ &\simeq (f \times 1_{\widehat{B}})^*(\mathscr{L}_{\widehat{b}}|_{f(A)})\end{aligned}$$

$$\simeq f^*\mathscr{L}_{\widehat{b}}$$
$$\simeq \mathscr{L}_{\widehat{a}}.$$

□

We can now prove

PROPOSITION 1.2.6.6. If $f : A \to B$ is an isogeny of degree d, then $\widehat{f} : \widehat{B} \to \widehat{A}$ is an isogeny of degree d.

Proof. Let $n > 0$ be an integer killing the finite k-group scheme $\ker f$ (e.g., we may take $n = \#\ker f$ by [TO70, §1 Deligne's Theorem]), we have that $A[n] \supset \ker f$. Thus, f factorizes:

$$A \xrightarrow[f]{} B \xrightarrow{[n]_A} A$$

and dualizing gives

$$\widehat{A} \xleftarrow[\widehat{f}]{} \widehat{B} \xleftarrow{[n]_{\widehat{A}}} \widehat{A}$$

(where of course $\widehat{[n]_A} = [n]_{\widehat{A}}$ by Corollary 1.2.4.3). This implies $\widehat{f} : \widehat{B} \to \widehat{A}$ is surjective, and since A and B and thus $\widehat{A}$ and $\widehat{B}$, are the same, $\widehat{f}$ has 0-dimensional fibers over a dense open of $\widehat{A}$ and thus by translation arguments over $\overline{k}$, all fibers are finite. Thus, $\widehat{f}$ is a finite surjective map. It is also generically flat and thus everywhere flat again by homogeneity considerations over $\overline{k}$. Thus, $\widehat{f}$ is an isogengy.

Now, we must show the degree of $\widehat{f}$ is d. Consider the diagram

$$\begin{array}{ccc} A \times \widehat{B} & \xrightarrow{1\times\widehat{f}} & A \times \widehat{A} \\ {\scriptstyle f\times 1}\downarrow & & \\ B \times \widehat{B}. & & \end{array}$$

By Lemma 1.2.6.5 $(1 \times \widehat{f})^*\mathscr{P}_A \simeq (f \times 1)^*\mathscr{P}_B$. But $1 \times \widehat{f}$ and $f \times 1$ are isogenies of degrees $\widehat{d}$ and d respectively. To do so, we will compare their Euler characteristics via

LEMMA 1.2.6.7. If $\varphi : A \to B$ is a degree-m isogeny between abelian varieties over k, then

$$\chi(\varphi^*\mathscr{L}) = m\chi(\mathscr{L})$$

for any line bundle $\mathscr{L}$ on B.

Granting this lemma, we get that

$$\widehat{d}\chi(\mathscr{P}_B) = d\chi(\mathscr{P}_A).$$

But since A and B have the same dimension $\chi(\mathscr{P}_B) = (-1)^g = \chi(\mathscr{P}_A)$ (by Corollary 1.2.6.4) is a common non-zero integer, so $\widehat{d} = d$.

□

Proof of Lemma 1.2.6.7. For the finite k-group scheme $G = \ker\varphi$, the map φ expresses X as a G-torsor

over Y for the fppf topology. We may therefore drop the abelian variety and isogeny hypotheses and instead prove the result for general G-torsors X over Y for a finite k-group scheme G with $m := \dim_k k[G]$. Furthermore, in this generality, we claim that $\chi(\varphi^*\mathscr{F}) = m\chi(\mathscr{F})$ for every coherent sheaf $\mathscr{F}$ on Y.

As in the proof of Proposition 1.1.4.6, we may use the Grothendieck Unscrewing Lemma (but with a mild variant [EGA, III, Cor. 3.1.3]) which gives us that it suffices to for every closed integral subscheme $V \subset Y$, the assertion holds for *one* coherent sheaf $\mathscr{F}$ on Y with coherent annihilator ideal $\mathscr{I}_V \subset \mathscr{O}_Y$. By passing to the G-torsor $\varphi^{-1}(V) \to V$, we may rename V as Y and just need to check for one $\mathscr{F}$ with annihilator ideal (0). For this purpose, set $\mathscr{F}$ to be $\varphi_*(\mathscr{O}_X)$. The G-torsor condition implies that the map

$$\theta : G \times X \to X \times_Y X$$
$$(g, x) \mapsto (gx, x)$$

is an isomorphism. By compatibility of (finite) pushforward of coherent sheaves with flat base change, we have $\varphi^*(\varphi_*(\mathscr{O}_X)) = (q_2)_*(\mathscr{O}_{X\times_Y X})$ for the second projection $q_2 : X \times_Y X \to X$. Via the X-isomorphism θ, the latter pushforward is identified with $k[G] \otimes_k \mathscr{O}_X = \mathscr{O}_X^{\oplus m}$ as an $\mathscr{O}_X$-module, so

$$\chi(\varphi^*(\mathscr{F})) = \chi(\mathscr{O}_X^{\oplus m}) = m\chi(\mathscr{O}_X) = m\chi(\varphi_*(\mathscr{O}_X)) = m\chi(\mathscr{F}),$$

where the final equality holds since φ is finite.

□

The Duality Hypothesis

Before we prove the duality hypothesis (Cor 1.2.6.12), we will make a very important digression which will come up multiple times throughout this paper. Our goal is to show that given a closed k-subgroup scheme, G^0_{red} is an abelian subvariety. To do so, we will use the schematic density of torsion points:

> LEMMA 1.2.6.8. Let $\ell \neq \operatorname{char} k$ be prime. The collection
>
> $$A[\ell^\infty] := \bigcup_{n=1}^{\infty} A[\ell^n]$$
>
> is relatively schematically dense in A, i.e., for any k-scheme S, a closed subscheme of A_S containing $A[\ell^n]$ for all $n \geq 1$ is equal to A_S.

Proof. Consider the Zariski closure, Z of $A[\ell^\infty]_S$ in A_S. Note that because in the map $S \to \operatorname{Spec} k$, the generic points of the irreducible components map to generic points (since $\operatorname{Spec} k$ is a point, this is trivial), we have that Zariski closure commutes with base change to S, i.e. $Z = Z'_S$ where Z' is the Zariski closure of $A[\ell^\infty]$ in A. Thus, if we can prove that $Z' = A$, we will have our result. So, we may assume $S = \operatorname{Spec} k$. We claim that Z is an abelian variety. It is a subgroup scheme because the Zariski closure of any subgroup scheme is a subgroup scheme. It is clearly proper. Finally, it is smooth since smoothness may be checked over $\overline{k}$, Zariski closure commutes with flat base change, and closure is given the reduced structure and is thus smooth (reducedness over an algebraically closed field implies the existence of a smooth point and homogeneity considerations show that every point is smooth). Thus, Z is an abelian variety with torsion points $Z[\ell] = A[\ell] \simeq (\mathbb{Z}/\ell\mathbb{Z})^{2g}$. But since abelian varieties of dimension g in characteristic not equal to ℓ automatically have torsion points isomorphic to $(\mathbb{Z}/\ell\mathbb{Z})^{2g}$, the dimension of Z must be g. But this implies

that $Z = A$. □

THEOREM 1.2.6.9. If $G \subset A$ be a closed k-subgroup scheme, G^0_{red} is smooth and thus an abelian subvariety.

Proof. Let Z be the Zariski-closure of the collection of étale k-subgroups $G[n] = G \cap A[n]$ for $\operatorname{char} k \nmid n$. We claim that Z is a smooth k-subgroup scheme of G and $Z^0 = G^0_{\text{red}}$. Firstly, we claim that the formation of Z commutes with extension of base field. This is because $[n]$ does, i.e., for field extension K/k, $G[n]_K = G_K[n]$, and Zariski closure commutes with flat base change (because generic points map to generic points), so in particular for field extension K/k.

Thus, without loss of generality, we may let $k = \overline{k}$. In this case, as the Zariski closure is given the reduced structure, it is clear that Z is a reduced group scheme and therefore smooth (as reduced groups schemes over an algebraically closed field are automatically smooth). Now, we want to show that $Z^0 = G^0_{\text{red}}$. But, since $k = \overline{k}$ and G^0_{red} is reduced, thus smooth, and therefore is an abelian variety. So, by Lemma 1.2.6.8, the collection of ℓ^n-torsion points is dense and thus its closure, $Z^0 = G^0_{\text{red}}$. □

Now that we have gotten this result out of the way, we may now move on to the duality hypothesis. To prove it, we will use the general theory of divisoral correspondences, which, under the analogy with vector spaces, behave like non-degenerate bilinear forms. Then, we will use $\mathscr{P}_A$ to give the double duality pairing.

DEFINITION 1.2.6.10. A **divisoral correspondence** between two abelian varieties A and B of the same dimension is a line bundle $\mathscr{Q}$ on $A \times B$ such that $\mathscr{Q}|_{\{0\}\times B}$ and $\mathscr{Q}|_{A\times\{0\}}$ are trivial.

The fact that a divisoral correspondence $\mathscr{Q}$ trivializes on $A \times \{0\}$ gives a homomorphism $f_{\mathscr{Q}} : A \to \widehat{B}$ and the fact that it trivializes on $\{0\} \times B$ gives a homomorphism $f'_{\mathscr{Q}} : B \to \widehat{A}$.

THEOREM 1.2.6.11. Let A, B be abelian varieties of dimension $g > 0$ and $\mathscr{Q}$ a divisoral correspondence between A and B, and let $f_{\mathscr{Q}} : A \to \widehat{B}$ and $f'_{\mathscr{Q}} : B \to \widehat{A}$ be the associated homomorphisms. The following are equivalent.

1. $|\chi(\mathscr{Q})| = 1$,
2. $f : A \to \widehat{B}$ or $f' : B \to \widehat{A}$ is an isomorphism,
3. both f and f' are isomorphism.

Proof. The flip $\sigma : B \times A \simeq A \times B$ clearly satisfies $\chi(\sigma^*\mathscr{Q}) \simeq \chi(\mathscr{Q})$, and $f'_{\mathscr{Q}} = f_{\sigma^*(\mathscr{L})}$, so it suffices to show that $|\chi(\mathscr{Q})| = 1 \iff f_{\mathscr{Q}}$ is an isomorphism.

Firstly, if $f_{\mathscr{Q}}$ is an isomorphism then $(1 \times f_{\mathscr{Q}})^*(\mathscr{P}_B) = \mathscr{Q}$ (by the definition of $f_{\mathscr{Q}}$) so since $1 \times f_{\mathscr{Q}} : A \times B \to B \times \widehat{B}$ is an isomorphism, we have $\chi(\mathscr{Q}) = \chi(\mathscr{P}_B) = \pm 1$.

Conversely, suppose $|\chi(\mathscr{Q})| = 1$, so $\chi(\mathscr{Q}) \in \mathbf{Z}$ is not divisible by any prime. We want $f_{\mathscr{Q}} : A \to \widehat{B}$ to be an isomorphism. Since $\dim A = \dim B$, $f_{\mathscr{Q}}$ is an isomorphism if and only if $\ker(f_{\mathscr{Q}}) = 0$ as a scheme. This is because finiteness of the kernel implies $f_{\mathscr{Q}}$ is an isogeny (by dimension and flatness considerations), its degree matches the order of its kernel scheme, and an isogeny of degree 1 is an isomorphism.

We may assume $k = \overline{k}$. Suppose $K := \ker(f_{\mathscr{Q}}) \neq 0$. We know by Theorem 1.2.6.9 that K^0_{red} is an abelian subvariety of A. If K is not finite, then $K^0_{\mathrm{red}}[\ell] \neq 0$ for some prime $\ell \neq \operatorname{char} k$ (in fact for all $\ell \neq \operatorname{char} k$ as we have proven in Proposition 1.1.4.11). Therefore, if $K \neq 0$, it always contains a finite closed k-subgroup scheme $H \neq 1$. Thus, we have a factorization

$$\begin{array}{ccc} A & \xrightarrow{\quad f_{\mathscr{Q}} \quad} & \widehat{B} \\ & {\scriptstyle h}\searrow \quad \nearrow{\scriptstyle g} & \\ & A/H & \end{array}$$

with h an isogeny of degree $\#H > 1$. Then by Lemma 1.2.6.7,

$$\chi(\mathscr{L}) = \chi((1 \times f_{\mathscr{Q}})^* \mathscr{P}_B) = \chi(h^*(\mathscr{M})) = (\deg h)\chi(\mathscr{M})$$

for $\mathscr{M} = (1 \times g^*)^* \mathscr{P}_B$ on $B \times A/H$. But this implies $\chi(\mathscr{Q})$ is divisible by $\deg h = \#H > 1$, a contradiction. □

Corollary 1.2.6.12 (The duality hypothesis). For any abelian variety A, the canonical morphism $\iota_A : A \to \widehat{\widehat{A}}$ defined by the Poincaré bundle $\mathscr{P}_A$ on $A \times \widehat{A}$ is an isomorphism.

Proof. Apply Theorem 1.2.6.11 to $\mathscr{P}_A$, which is a divisoral correspondence by definition and satisfies (1) by 1.2.6.4. □

Now that we know that A is canonically isomorphic to $\widehat{\widehat{A}}$, we introduce the notion

Definition 1.2.6.13. A homomorphism $f : A \to \widehat{A}$ is **symmetric** if it splits as $f = \widehat{f} \circ \iota_A$:

$$f : A \xrightarrow{\iota_A} \widehat{\widehat{A}} \xrightarrow{\widehat{f}} A.$$

Remark 1.2.6.14. In our running analogy with linear algebra, symmetric morphisms correspond to symmetric bilinear forms $B : V \times V \to k$ in the sense that the corresponding linear map $V \xrightarrow{T_B} V^\wedge$ factors as $V \xrightarrow{\sim} (V^\wedge)^\wedge \xrightarrow{T_B^\wedge} V^\wedge$ through the canonical double dual isomorphism. Furthermore, for any bilinear form $B : V \times V \to k$ with induced map T_B, the composition

$$V \xrightarrow{(1,T_B)} V \times V^\wedge \xrightarrow{\text{eval}} k$$

recovers B. Thus the evaluation map has the same universal property as the Poincaré bundle.

Furthermore, in remark 1.2.5.6, we noted that the relationship between $\mathscr{L}$ and $\phi_{\mathscr{L}}$ is similar to that between a quadratic form q and its associated *symmetric* bilinear form B_q. This motivates Proposition 1.2.6.15.

Proposition 1.2.6.15. For any line bundle $\mathscr{L}$ on A, $\phi_{\mathscr{L}} : A \to \widehat{A}$ is symmetric.

Proof. We want to show

$$A \xrightarrow{\iota_A} \widehat{\widehat{A}} \xrightarrow{\widehat{\phi_{\mathscr{L}}}} \widehat{A}$$

is equal to $\phi_{\mathscr{L}}$. Equivalently, we want to show that $\phi_{\mathscr{L}} \circ \iota_A^{-1} : \widehat{\widehat{A}} \to \widehat{A}$ is equal to $\widehat{\phi_{\mathscr{L}}}$. Since maps to the dual abelian variety are classified by their pullback of the Poincaré bundle, we conclude that it is equivalent to prove

$$(1 \times (\phi_{\mathscr{L}} \circ \iota_A^{-1}))^* \mathscr{P}_A \simeq (\phi_{\mathscr{L}} \times 1)^* \mathscr{P}_{\widehat{A}}.$$

By Lemma 1.2.6.5, for general $f : A \to B$ with associated diagram

$$\begin{array}{ccc} A \times \widehat{B} & \xrightarrow{1 \times \widehat{f}} & A \times \widehat{A} \\ {\scriptstyle f \times 1}\downarrow & & \\ B \times \widehat{B} & & \end{array}$$

we know that

$$(1 \times \widehat{f})^* \mathscr{P}_A \simeq (f \times 1)^* \mathscr{P}_B$$

as line bundles on $A \times \widehat{B}$. Applying this to the map $\phi_{\mathscr{L}}$, we get

$$\begin{array}{ccc} A \times \widehat{\widehat{A}} & \xrightarrow{\phi_{\mathscr{L}} \times 1} & \widehat{A} \times \widehat{\widehat{A}} \\ {\scriptstyle 1 \times \iota_A^{-1}}\downarrow \wr & \nearrow_{\phi_{\mathscr{L}} \times \iota_A} & \\ A \times A & & \\ {\scriptstyle 1 \times \phi_{\mathscr{L}}}\downarrow & & \\ A \times \widehat{A} & & \end{array}$$

where the upper triangle is just an expression of the isomorphism $1 \times \iota_A^{-1}$. We have

$$\begin{aligned} (1 \times (\phi_{\mathscr{L}} \circ \iota_A^{-1}))^* \mathscr{P}_A &\simeq (1 \times \iota_A^{-1})^* (1 \times \phi_{\mathscr{L}})^* \mathscr{P}_{\widehat{A}} \\ &\simeq (\phi_{\mathscr{L}} \times 1)^* (1 \times \iota_A)^* \mathscr{P}_{\widehat{A}} \\ &\simeq (\phi_{\mathscr{L}} \times 1)^* \sigma^* (\mathscr{P}_A) \end{aligned}$$

where $\sigma : \widehat{A} \times A \xrightarrow{\sim} A \times \widehat{A}$ is the flip isomorphism and the last equality follows by the construction of double duality (see the end of §1.2.1). But, $(\phi_{\mathscr{L}} \times 1)^* \circ \sigma^* = (\sigma \circ (\phi_{\mathscr{L}} \times 1))^*$ and $\sigma \circ (\phi_{\mathscr{L}} \times 1) = (1 \times \phi_{\mathscr{L}}) \circ \sigma'$ where $\sigma' : A \times A \simeq A \times A$ is the flip isomorphism, and the line bundle $(1 \times \phi_{\mathscr{L}})^* \mathscr{P}_A = \Lambda(\mathscr{L})$ is canonically invariant under σ'-pullback (see Proposition 1.2.1.2). Thus, we see that it is isomorphic to $(1 \times \phi_{\mathscr{L}})^* \mathscr{P}_A$ as required.

□

Section 1.3

Applications of Duality

In this section, we will see some of the fruits of our labor in the previous two sections. We will develop the theory of the Weil pairing, which will be invaluable in studying the endomorphisms of abelian varieties in Chapter 2. We will then gain some greater granularity on the structure of the p-power torsion subgroups of abelian varieties, which allows us to get a greater grasp on the structure of positive characteristic abelian varieties. Finally, we will examine the cohomology of line bundles, giving the analogue for the Riemann-Roch theorem on abelian varieties, and proving the striking result that all non-degenerate line bundles have vanishing cohomology outside a single degree (called the index). As was the theme of many of the previous sections, getting insight into the structures of line bundles is the way to understand an abelian variety, and in line with this, these results will be important once we dive deeper into the structure of endomorphisms in Chapter 2.

§1.3.1 Cartier Duality and The Weil Pairing

The first goal for this section is to show that our newly constructed duality theory in fact coincides with Cartier duality on torsion points! We will then use the general theory of Cartier duality to construct a bi-additive pairing $A[n] \times \widehat{A}[n] \to \mu_n$. This will become very important in classifying endomorphisms of abelian varieties in Chapter 2.

Cartier Duality

We begin this section with a brief introduction to Cartier duality, which one can read more about in [Mum74, §14] for the case over $\overline{k}$, and [SGA3, §3] for the general theory over an arbitrary base.

Suppose G is a finite locally free commutative group scheme over some arbitrary scheme S (i.e., $\pi_*\mathscr{O}_G$ is a locally free $\mathscr{O}_S$-module of finite rank).

> Definition 1.3.1.1. The **Cartier dual** of G is the group scheme $\mathbf{D}(G)$ whose functor of points on S-schemes is
>
> $$T \rightsquigarrow \mathrm{Hom}_{T-\mathrm{gp}}(G_T, \mathbf{G}_{m,T}).$$

It turns out that this is represented by the group scheme whose underlying $\mathscr{O}_S$-module is the vector bundle

$$\pi_*(\mathscr{O}_{\mathbf{D}(G)}) = \underline{\mathrm{Hom}}_{\mathscr{O}_S}(\pi_*\mathscr{O}_G, \mathscr{O}_S).$$

Additionally, the formation of $\mathbf{D}(G)$ naturally commutes with base change, and there is an obvious bi-additive pairing $G \times_S \mathbf{D}(G) \to \mathbf{G}_m$ (given functorially by the pairing $G_T \times \mathrm{Hom}_{T-\mathrm{gp}}(G_T, \mathbf{G}_{m,T})$). This pairing induces an S-homomorphism

$$G \xrightarrow{\sim} \underline{\mathrm{Hom}}(\mathbf{D}G, \mathbf{G}_m) \simeq \mathbf{D}\mathbf{D}G$$

which is an isomorphism.

Duality of $A[n]$ and $\widehat{A}[n]$

We now want to show that $A[n]$ and $\widehat{A}[n]$ are Cartier dual to one another. In fact we will show more generally that for any map of abelian varieties $f : A \to B$:

> THEOREM 1.3.1.2. There is a natural Cartier duality pairing
>
> $$\langle \cdot, \cdot \rangle_f : \ker f \times \ker \widehat{f} \to \mathbf{G}_m.$$
>
> I.e., the associated map $\ker \widehat{f} \to \mathbf{D}(\ker f)$ is an isomorphism.

The particular cases in which we will be most interested are when $B = \widehat{A}$ and $B = A$.

WARNING. The relationship between

$$\langle \cdot, \cdot \rangle_{\widehat{f}} : \ker \widehat{f} \times \ker \widehat{\widehat{f}} \to \mathbf{G}_m$$

and $\langle \cdot, \cdot \rangle_f$ via $f \simeq \widehat{\widehat{f}}$ is not just via a flip of the components, but there is an additional sign error in the sense that

$$\langle \widehat{b}, \iota_A(a) \rangle_{\widehat{f}} = \langle a, \widehat{b} \rangle_f^{-1}.$$

We will prove this when it becomes relevant in Lemma 2.1.3.4.

Proof. We want to construct an isomorphism

$$(\ker \widehat{f})(T) \overset{?}{\simeq} \mathrm{Hom}_{T-\mathrm{gp}}(\ker(f_T), \mathbf{G}_{m,T})$$

naturally in k-schemes T. By definition, the left side is

$$\ker(\mathrm{Pic}^0_{B/k}(T) \xrightarrow{f^*} \mathrm{Pic}^0_{A/k}(T)).$$

Firstly, we would like to show that this is the same as

$$\ker(\mathrm{Pic}_{B/k}(T) \xrightarrow{f^*} \mathrm{Pic}_{A/k}(T)),$$

dropping the restriction to the identity components. Clearly, we may replace the target by $\mathrm{Pic}^0_{A/k}$ without changing the kernel. To show that we may do the same with the source, we first note that the statement of equality of these two kernels is Zariski-local, so we may assume T is local. Additionally, this problem is topological: it is equivalent to saying that if a T-valued point of $\mathrm{Pic}_{B/k}$ mapping to 0, then it comes from the identity component $\mathrm{Pic}_{B/k}$. Thus, without loss of generality, we may assume $T = \mathrm{Spec}\, K$ for a field K (we may even assume K is algebraically closed). Over K, f^* is an isogeny, hence invertible up to composing with a non-zero integer by the fact that its kernel is killed by multiplication by the degree (see Proposition 2.1.1.4 for full details), and $\mathrm{NS}(A_K)$ is torsion-free on which the effect of $[n]_A$ is multiplication by n^2 (Lemma 1.2.5.7). Thus, the induced map $\mathrm{NS}(B_K)_{\mathbb{Q}} \to \mathrm{NS}(A_K)_{\mathbb{Q}}$ is an isomorphism, so there is no kernel between Néron-Severi groups. This gives us the result on replacing the source.

We conclude that

$$(\ker \widehat{f})(T) = \left\{ \mathscr{L} \in \mathrm{Pic}(B_T) | f_T^* \mathscr{L} \simeq \mathscr{O}_{A_T} \right\} \Big/ \phi_T^*(\mathrm{Pic}(T)) \subset \widehat{B}(T)$$

(where $\phi_T : A_T \to T$ is the base change of the structure morphism). Furthermore, since the diagram

$$\begin{array}{ccc} A_T & \xrightarrow{f_T} & B_T \\ & {\scriptstyle\psi_T}\searrow \quad \swarrow{\scriptstyle\phi_T} & \\ & T & \end{array}$$

commutes, i.e., since f_T is a T-morphism, $f_T^*\phi_T^*\mathscr{M} = \psi_T^*\mathscr{M}$ for any $\mathscr{M} \in \mathrm{Pic}(T)$. Thus, $f_T^*(\mathscr{L} \otimes \phi_T^*\mathscr{M}) = f_T^*\mathscr{L} \otimes \psi_T^*\mathscr{M}$, but if this maps to the trivial bundle, i.e., if

$$f_T^*(\mathscr{L} \otimes \phi_T^*\mathscr{M}) \simeq \mathscr{O}_{A_T}$$

for some $\mathscr{M} \in \mathrm{Pic}(T)$, then we must have

$$f_T^*\mathscr{L} \simeq \psi_T^*(\mathscr{M}^{-1})$$

but if $\psi_T^*(\mathscr{M}^{-1}) \simeq \mathscr{O}_{A_T}$, then $\mathscr{M} = \mathscr{O}_T$ because $e_T^*\psi_T^*(\mathscr{M}^{-1}) = \mathscr{M}^{-1}$ for example (note that this relies on having a T-point, but this exists for all group schemes). Therefore each coset $[\mathscr{L}]$ in $\mathrm{Pic}(B_T)/\phi_T^*(\mathrm{Pic}(T))$ with some $f_T^*\mathscr{L} \simeq \mathscr{O}_{A_T}$ has a unique representative mapping to $\mathscr{O}_{A_T}$. I.e.,

$$\left\{\mathscr{L} \in \mathrm{Pic}(B_T) | f_T^*\mathscr{L} \simeq \mathscr{O}_{A_T}\right\} \Big/ \phi_T^*(\mathrm{Pic}(T)) = \ker(\mathrm{Pic}(B_T) \to \mathrm{Pic}(A_T)).$$

For $G := \ker f$ and any k-scheme T, we have $G_T = \ker f_T$, and likewise $B_T = A_T/G_T$. We need to identify $\ker \widehat{f}(T) = \ker(\mathrm{Pic}(B_T) \to \mathrm{Pic}(A_T))$ with $\mathrm{Hom}_{T-\mathrm{gp}}(G_T, \mathbf{G}_{m,T})$ naturally in T.

Now we will consider this problem in a more general setting.

> PROPOSITION 1.3.1.3. Let $f : X \to S$ be proper, flat, and finitely presented, with reduced and connected geometric fibers (so that $\mathscr{O}_S \simeq f_*\mathscr{O}_X$). Suppose that we are given $G \to S$ a finite locally free commutative S-group with a free action of G on X over S (i.e. $G(T)$ acts freely on $X(T)$ for all S-schemes T, or equivalently, the action map $G \times_S X \to X \times_S X$ is a monomorphism).
>
> Suppose further that there exists an S-scheme Y representing X/G. In particular such that the action map $G \times_S X \to X \times_Y X$ is an isomorphism. Then for all $T \to S$, we have
>
> $$\ker(\mathrm{Pic}(Y_T) \to \mathrm{Pic}(X_T)) \simeq \mathrm{Hom}_{T-\mathrm{gp}}(G_T, \mathbf{G}_{m,T})$$
>
> naturally in T.

Granting this, it is clear that given our $f : A \to B$ as above with $G = \ker f$ over $S = \mathrm{Spec}(k)$, all of these conditions are fulfilled, so we are done. □

Proof of Proposition 1.3.1.3. This will crucially use descent theory (see Chapter 6 of the book Néron Models [BLR90b] for good exposition on descent). Note also that $f_*\mathbf{G}_{m,X} = \mathbf{G}_{m,S}$ as functors on S-schemes. This is because for every S-scheme T, the natural map $\mathscr{O}_T \to (f_T)_*(\mathscr{O}_{X_T})$ of sheaves of rings is an isomorphism and so passing to the unit group yeilds the isomorphism $\mathscr{O}_T^\times \simeq ((f_T)_*(\mathscr{O}_{X_T}))^\times = (f_T)_*(\mathscr{O}_{X_T}^\times)$. Passing to global sections then gives $\mathbf{G}_m(T) \simeq \mathbf{G}_m(X_T)$.

Since our setup is compatible with base change, without loss of generality, we will just treat the case $T = S$. It will be clear that what we do for the rest of the proof is also compatible with base change. The kernel $\ker(\mathrm{Pic}(Y) \to \mathrm{Pic}(X))$ consists of descents of $\mathscr{O}_X$ to a line bundle on Y, up to isomorphism on X.

Consider the diagram

$$\begin{array}{ccc} G \times_S X & \xrightarrow{(g,x)\mapsto(gx,x)} & X \times_Y X \\ & \searrow^{(m,\mathrm{pr}_2)} & \downarrow{\scriptstyle \mathrm{pr}_1},\ {\scriptstyle \mathrm{pr}_2} \\ & & X. \end{array}$$

By descent theory $\ker(\mathrm{Pic}(Y) \to \mathrm{Pic}(X))$ is equal to the data

$$\{\text{isomorphisms } \mathrm{pr}_1^* \mathscr{O}_X \simeq \mathrm{pr}_2^* \mathscr{O}_X \text{ satisfying the cocycle condition as } \mathscr{O}_{X\times X}\text{-modules}\}.$$

Notably, we may exclude mention of coboundaries since $f_* \mathbf{G}_{m,X} = \mathbf{G}_{m,S}$, and so they are all trivial. The data of an isomorphism between the pr_1-pullback and the pr_2-pullback amounts to a unit on $G \times X$ satisfying a cocycle condition expressed in terms of the action. For any unit u on $G \times X$, S-scheme T, and point $g \in G(T)$, $u(g, \cdot)$ is a unit on X_T by pullback:

$$X_T \xrightarrow{(g,\mathrm{id})} G_T \times X_T.$$

But a unit on X_T is the same as a unit on T. Let us write $u(g)$ to denote this unit on T. Then, the cocycle condition on u as T varies expresses exactly that $G \to \mathbf{G}_m$ of group-valued functors on S-schemes defined by $g \mapsto u(g)$ is a homomorphism.

□

Before we begin studying the Weil pairing (which is just the pairing given for $f = [n]_A$), we need to give a more explicit description of the pairing for general f so that we may work with it easier. From the general theory of Cartier duality, the results of Theorem 1.3.1.2 implicitly gets us a pairing

$$\langle \cdot, \cdot \rangle_f : \ker f \times \ker \widehat{f} \to \mathbf{G}_m$$

for any isogeny $f : A \to B$. But, we would like an explicit description of these pairings. Specifically for the Weil pairing (corresponding to $[n]$) as it will allow us to extend it to Tate modules, which were briefly introduced at the end of Section 1.1.4 (Def. 1.1.4.14).

Let $a \in (\ker f)(T)$ and $\widehat{b} \in (\ker \widehat{f})(T)$ for some k-scheme T. We would like to describe $\langle a, \widehat{b} \rangle \in \mathbf{G}_m(T)$. The T-point $\widehat{b} \in (\ker \widehat{f})(T)$ corresponds to a class of line bundles $[\mathscr{L}] = \{\mathscr{L} \otimes \phi_T^* \mathscr{M} | \mathscr{M} \in \mathrm{Pic}(T)\}$ for some $\mathscr{L}$ on B_T. Note that by remark 1.1.2.3, there is a unique representative $\mathscr{L}$ of $[\mathscr{L}]$ which comes with a trivializing section σ along $e'_T \in B_T(T)$. Thus, we may think of $\widehat{b}$ as the pair $(\mathscr{L}, \sigma)$ and additionally, since $\widehat{b} \in \ker \widehat{f}$ there is a trivialization $\varphi : f_T^*(\mathscr{L}) \simeq \mathscr{O}_{A_T}$ respecting the trivializations on both sides along e_T (noting that $f \circ e_T = e'_T$).

Since $a \in (\ker f)(T)$, the following diagram commutes

$$\begin{array}{ccc} A_T & \xrightarrow{t_a} & A_T \\ & {\scriptstyle f_T} \searrow \quad \swarrow {\scriptstyle f_T} & \\ & B_T. & \end{array}$$

This provides a canonical isomorphism

$$t_a^* f_T^* \mathscr{L} \overset{\mathrm{can}}{\simeq} f_T^* \mathscr{L}.$$

But the trivializations $t_a^*(\sigma)$ and σ on the respective sides are possibly not compatible with this isomor-

phism. Thus there is some unit $u \in \mathbf{G}_m(A_T) = \mathbf{G}_m(T)$ such that $t_a^*(\sigma) = u\sigma$. We claim that

$$\langle a, \widehat{b} \rangle_f = \frac{t_a^*(\sigma)}{\sigma}. \tag{1.3.1-I}$$

To see this, we need to study the isomorphism we constructed $\mathbf{D}(\ker f) \simeq \ker \widehat{f}$. In the proof of Proposition 1.3.1.3, we saw that for a line bundle $\mathscr{L} \in \ker(\mathrm{Pic}(B_T) \to \mathrm{Pic}(A_T))$, the unit from a base k-scheme T is the unit giving the isomorphism between $\mathrm{pr}_1^* \mathscr{O}_A \simeq \mathrm{pr}_2^* \mathscr{O}_A$ corresponding to $\mathscr{L}$ under descent theory. Explicitly, we have that under the the equivalence of categories between line bundles $\mathscr{L}$ on B_T with pairs $(\mathscr{M}, \alpha)$ where $\mathscr{M}$ is a line bundle on A_T and $\alpha : \mathrm{pr}_1^* \mathscr{M} \to \mathrm{pr}_2^* \mathscr{M}$ is an isomorphism (for $\mathrm{pr}_1, \mathrm{pr}_2 : A_T \times_{B_T} A_T \rightrightarrows A_T$). This functor maps $\mathscr{L} \mapsto (f_T^* \mathscr{L}, \phi)$ where ϕ is the canonical isomorphism $\phi : \mathrm{pr}_1^* f_T^* \mathscr{L} \to \mathrm{pr}_2^* f_T^* \mathscr{L}$ coming from the associated coequalizer diagram. Under the identification of $A_T \times_{B_T} A_T$ with $\ker f_T \times_T A_T$, the image of a point (a, x), is (ax, x) and so the isomorphism ϕ restricted to $\{a\} \times A_T$ is exactly the canonical identification of $t_a^* f_T^* \mathscr{L}$ with $f_T^* \mathscr{L}$. Thus, the unit u on $\ker f \times_T A_T$ must restrict to $\{a\} \times A_T$ to give the isomorphism $t_a^* f_T^* \mathscr{L} \overset{\mathrm{can}}{\simeq} f_T^* \mathscr{L}$. As we may detect the identification of a unit on A_T with a the unit on T on a fiber of $A_T \to T$, picking the section e and tracing the trivialization σ gives us that the unit $u \in \mathbf{G}_m(A_T) = \mathbf{G}_m(T)$ is exactly the unit such that $t_a^*(\sigma) = u\sigma$.

Now that we have this description, we are able to discuss the functoriality of the Weil pairing

$$\langle \cdot, \cdot \rangle_{A,n} : A[n] \times \widehat{A}[n] \to \mu_n$$

with respect to *both* A and n. Firstly, let $h : A \to A$ be a homomorphism. We claim that the following diagram "commutes"

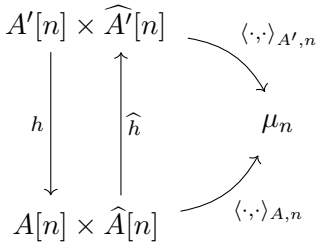

in the sense that for any k-scheme T, $a' \in A'[n](T)$ and $\widehat{a} \in \widehat{A}[n](T)$,

$$\langle h(a'), \widehat{a} \rangle_{A,n} = \langle a', \widehat{h}(\widehat{a}) \rangle_{A',n} \tag{1.3.1-II}$$

in $\mu_n(T)$. Let $\mathscr{L}$ on A_T correspond to $\widehat{a}$, which is equipped with a trivialization $\sigma : [n]_A^* \mathscr{N} \simeq \mathscr{O}_A$. The pullback $h^* \mathscr{L}$ corresponds to $\widehat{h}(\widehat{a})$ which admits the associated trivialization

$$[n]_{A'}^* h^* \mathscr{L} = h^* [n]_A^* \mathscr{L} \overset{h^*(\sigma)}{\simeq} \mathscr{O}_{A'}$$

Thus, by what we showed above, $\langle a', \widehat{h}(\widehat{a}) \rangle_{A',n}$ is the unit which under the canonical isomorphism

$$t_{a'}^*([n]_{A'}^* h^* \mathscr{L}) \simeq [n]_{A'}^* h^* \mathscr{L}$$

sends $t_{a'}^*(h^*\sigma) \mapsto \langle a', \widehat{h}(\widehat{a}) \rangle_{A',n} \cdot h^*\sigma$. To show equation 1.3.1-II, we will trace the trivialization σ across a list of equalities of line bundles. Firstly, we have

$$\begin{aligned} [n]_{A'}^* h^* \mathscr{L} &= t_{a'}^*([n]_{A'}^* h^* \mathscr{L}) \\ &= t_{a'}^*(h^* [n]_A^* \mathscr{L}) \end{aligned}$$

where the first equality is by a' being in $A'[n](T)$ and the second is by h being a homomorphism, so multiplication by n commutes with it. Next, we note that translations interact with homomorphisms as expected, i.e., the following diagram commutes:

$$\begin{array}{ccc} A' & \xrightarrow{t_{a'}} & A' \\ {\scriptstyle h}\downarrow & & \downarrow{\scriptstyle h} \\ A & \xrightarrow{t_{h(a')}} & A. \end{array}$$

This gives us

$$\begin{aligned} [n]_{A'}^* h^* \mathscr{L} &= t_{a'}^*([n]_{A'}^* h^* \mathscr{L}) \\ &= t_{a'}^*(h^* [n]_A^* \mathscr{L}) \\ &= h^*(t_{h(a')}^* [n]_A^* \mathscr{L}) \\ &= h^*([n]_A^* \mathscr{L}) \end{aligned}$$

and following our trivialization gives us

$$\begin{aligned} \langle a', \widehat{h}(\widehat{a})\rangle_{A',n} h^*\sigma &= t_{a'}^*(h^*(\sigma)) \\ &= h^*(t_{h(a')}^*(\sigma)) \\ &= \langle h(a'), \widehat{a}\rangle_{A,n} \cdot h^*\sigma. \end{aligned}$$

This gives us equation 1.3.1-II as desired.

Now, we want to also show "functoriality" in change in n, i.e.,

PROPOSITION 1.3.1.4. For $m, n \geq 1$, the following diagram commutes:

$$\begin{array}{ccc} A[nm] \times \widehat{A}[nm] & \xrightarrow{\langle\cdot,\cdot\rangle_{A,nm}} & \mu_{nm} \\ {\scriptstyle ([m]_A, [m]_{\widehat{A}})}\downarrow & & \downarrow{\scriptstyle t^m} \\ A[n] \times \widehat{A}[n] & \xrightarrow[\langle\cdot,\cdot\rangle_{A,n}]{} & \mu_n \end{array}$$

Proof. This is simple for n, m not divisible by the char k as we may check on geometric points since it only involves étale k-group schemes (this is covered by [Mum74, §20, Proposition 1]). We want to allow general $n, m \geq 1$ for completeness (as including the p-part is important for more advanced studies of abelian varieties through p-divisible groups).

Choose $a \in A(T)$ that is mn-torsion and choose a line bundle $\mathscr{L}$ on A_T for which $[mn]^*\mathscr{L}$ admits a trivializing section σ over A_T along the identity e. The choice of σ (which pins down the equivalence class $[\mathscr{L}] \in \mathrm{Pic}_{A/k}(T)$ and any two choices are related through a unique automorphism of the line bundle) specifies a trivialization $\sigma(e)$ of λ. We will pay great detail to our trivializing sections in order to use the explicit formula for the Weil pairing (1.3.1-I).

Let $\widehat{a} \in \widehat{A}[mn](T)$ be the T-point corresponding to $[\mathscr{L}]$. Since $[m]_{\widehat{A}} = \widehat{[m]_A}$ (Cor 1.2.4.3), so $\mathscr{L}^{\otimes m} \simeq [m]^*\mathscr{L}$ is the line bundle corresponding to the n-torsion point $m\widehat{a} \in \widehat{A}[n](T)$. In particular, σ may be identified with a trivializing section of the $[n]$-pullback of the line bundle $[m]^*\mathscr{L}$ corresponding to $m\widehat{a}$ via the canonical isomorphism $[mn]^*\mathscr{L} \simeq [n]^*[m]^*\mathscr{L}$. Hence, our task is to prove that the canonical

isomorphism

$$t_{ma}^*([n]^*[m]^*\mathscr{L}) \simeq [n]^*[m]^*$$

carries $t_{ma}^*(\sigma)$ to $u^m\sigma$ for the unit $u := \langle a, \widehat{a}\rangle_{A,mn}$ over T that satisfies $t_a^*(\sigma) \mapsto u\sigma$ via the canonical isomorphism $t_a^*([mn]^*\mathscr{L}) \simeq [mn]^*\mathscr{L}$ of line bundles on A_T.

The idea behind this is that t_{ma} is the m-fold composition of copies of t_a, so iterating m times over the "equality" $t_a^*(\sigma) = u\sigma$ should give the result, but the "equality" is actually over an isomorphism $t_{ma}^*([mn]^*\mathscr{L}) \simeq [mn]^*\mathscr{L}$, so we need to keep track of sections under a composition of isomorphisms rather than just use iterate this identity.

For any $x \in A[mn](T)$, let $\theta_x : t_x^*([mn]^*\mathscr{L}) \simeq [mn]^*\mathscr{L}$ be the canonical isomorphism expressing the equality $[mn] \circ t_x = [mn]$ of endomorphisms of A. The key point is that for any $y \in A[mn](T)$, θ_{x+y} splits as the composition

$$\theta_{x+y} = \theta_y \circ t_y^*(\theta_x) : t_y^*(t_x^*([mn]^*\mathscr{L})) \xrightarrow{\sim} t_y^*([mn]^*\mathscr{L}) \xrightarrow{\sim} [mn]^*\mathscr{L} \tag{1.3.1-III}$$

via the natural isomorphism of functors $t_y^* \circ t_x^* \simeq t_{x+y}^*$ on line bundles. This follows by tracking the equalities

$$[mn] \circ t_{x+y} = ([mn] \circ t_x) \circ t_y = [mn] \circ t_y = [mn].$$

Our aim is to show that for every $i \geq 1$, $\theta_{ia}(t_{ia}^*(\sigma)) = u^i \cdot \sigma$ for the unit u on T satisfying $\theta_a(t_a^*(\sigma)) = u \cdot \sigma$. Note that the choice of σ is irrelevant, since any two choices are related by scaling by a unit from T. It suffices to show more generally that for $a, b \in A[mn](T)$, if $\theta_a(t_a^*(\sigma)) = u \cdot \sigma$ and $\theta_b(t_b^*(\tau)) = v \cdot \tau$ for some (equivalently any) trivializing sections σ, τ of $[mn]^*\mathscr{L}$, then $\theta_{a+b}(t_{a+b}^*(\sigma)) = uv \cdot \sigma$ for any σ.

In view of the identity 1.3.1-III, we consider

$$(\theta_b \circ t_b^*(\theta_a))(t_b^*(t_a^*(\sigma))).$$

This is equal to

$$\theta_b(t_b^*(\theta_a(t_a^*(\sigma)))) = \theta_b(t_b^*(u\sigma)) = u\theta_b(t_b^*(\sigma)) = uv \cdot \sigma$$

where in the second equality, we may pass u through the maps because u is a unit from the base scheme T. Combining this with the isomorphism of functors $t_{a+b}^* \simeq t_b^* \circ t_a^*$ translates this final computation into the desired equality $\theta_{a+b}(t_{a+b}^*(\sigma)) = uv \cdot \sigma$. Applying this to $b = (i-1)a$ for $i \geq 2$, by induction, we have $\theta_{ia}(t_{ia}^*(\sigma)) = u^i \cdot \sigma$ and letting $i = m$ gives us the desired result. □

This allows us to extend the Weil pairings to our Tate modules. More specifically, as we have pairings $\langle \cdot, \cdot \rangle_{A,\ell^n}$ for all n and they are compatible under the maps in the above proposition, we may pass to the projective limit to get a $\mathbf{Z}_\ell$-bilinear pairing

$$T_\ell A \times T_\ell \widehat{A} \to \mathbb{Z}_\ell(1)$$

where we define $\mathbb{Z}_\ell(1) := \varprojlim \mu_{\ell^n}$ under ℓ-power maps. This will become important when we study Riemann forms in Chapter 2.

§1.3.2 The p-rank of an Abelian Variety

In this part, we will summarize the results of §14 and §15 of [Mum74] to give a more refined scheme structure on $A[p^n]$ which will tie up the loose end left by Proposition 1.1.4.11 at the end of §1.1.4. Many

of these results are not very relevant the rest of this paper so will largely be stated without proof. The curious reader can read §14 and §15 of [Mum74].

Finite commutative group schemes in characteristic p

We first need to build a general theory of finite commutative group schemes in characteristic p. Firstly note the important examples:

EXAMPLE 1.3.2.1. (1) Suppose G is a reduced finite group of order n prime to p. Then $\mathbf{D}(G)$ is a group of order n which is again reduced, and isomorphic to the constant reduced group scheme $\mathrm{Hom}(G, k^*)$.

(2) Next suppose G is reduced and isomorphic to the constant group scheme $\mathbb{Z}/p^n\mathbb{Z}$. Then, we have that

$$\mathbf{D}(\mathbb{Z}/p^n\mathbb{Z}) \simeq \mu_{p^n},$$

the group of p^nth roots of unity. And of course, by symmetry

$$\mathbf{D}(\mu_{p^n}) = \mathbb{Z}/p^n\mathbb{Z}.$$

Now, we are prepared to tackle classifying finite commutative group schemes in characteristic $p > 0$.

DEFINITION 1.3.2.2. Let G be any finite commutative group scheme. We say that G is of type l (local) or r (reduced) if the underlying space G is a single point of G is reduced respectively. We say that G is of type (l, l) (resp. (l, r), (r, l), (r, r)) if G is of type l and $\mathbf{D}(G)$ is of type l (resp. G of type l, $\mathbf{D}(G)$ of type r, etc.)

Then we have

PROPOSITION 1.3.2.3. Any finite commutative group scheme admits a unique decomposition as a product

$$G = G_{r,r} \times G_{r,l} \times G_{l,r} \times G_{l,l}$$

of groups of the indicated types.

Proof. In fact, if G^0 is the connected component of the identity, considered as an open and closed subscheme, and if G_{red} is the reduced group, the closed immersions $G^0 \hookrightarrow G$ and $G_{\mathrm{red}} \hookrightarrow G$ induce a homomorphism $G^0 \times G_{\mathrm{red}} \to G$ which is clearly an isomorphism. Further, this decomposition into a product of a reduced and local group is clearly unique. Thus, it suffices to show that each local (resp. reduced) group can be uniquely expressed as a product $G_{l,r} \times G_{l,l}$ (resp. $G_{r,r} \times G_{r,l}$). Now, if G is reduced, G is uniquely expressible as a product $G_1 \times G_2$ where G_1 is of order prime to p and G_2 is a p-group. By example 1.3.2.1 above, $\mathbf{D}(G_1)$ is again reduced and $\mathbf{D}(G_2)$ is local, which proves the assertion for reduced groups. When G is local, split $\mathbf{D}(G)$ into its local and reduced parts. Dualizing back, this implies a unique decomposition of a local G into types (l, r) and (l, l) respectively. This proves the assertion for local groups. (Note that we could have done the same thing for the reduced part, but we chose to use the classification given earlier to emphasize the structure of reduced commutative group schemes). □

It turns out that the only groups of type (r, r) are those reduced groups of order prime to p, hence direct products of cyclic prime power groups, the primes being distinct from p; that the only groups of type (r, l) are p-groups, hence direct products of $\mathbb{Z}/p^{n_i}\mathbb{Z}$'s; and that the only group of type (l, r) are duals of p-groups, hence direct products of $\mu_{p^{n_i}}$'s.

EXAMPLE 1.3.2.4. An example of a local-local group is α_{p^n} (which is self-dual), the Frobenius kernel of $\mathbf{G}_a$.

We continue to study p-groups specifically, starting with the following definition.

DEFINITION 1.3.2.5. A group scheme G is of **height one** if it consists of a single point e, and if $x^p = 0$ for all $x \in \mathfrak{m}_e$. (G need not be commutative).

And we state the following:

THEOREM 1.3.2.6. The functor $G \mapsto \mathrm{Lie}(G)$ sets up an equivalence of the categories of finite group schemes of height one and finite-dimensional p-Lie algebras over k (i.e., finite dimensional vector spaces with Lie bracket and pth power maps).

Proof. See [Mum74, §14, Theorem 1] for the commutative G case, which is all we need. □

COROLLARY 1.3.2.7. If G is a commutative group scheme of height one, the homomorphism $[p]_G$ is zero.

Proof. Multiplication by p kills $\mathrm{Lie}(G)$, so the result follows from Theorem 1.3.2.6. □

We also have the following

PROPOSITION 1.3.2.8. If V is any vector space with a p-linear map $x \mapsto x^{(p)}$, there is a unique decomposition, invariant under $x \mapsto x^{(p)}$:

$$V = V_s \oplus V_n,$$

such that V_s has a basis $x_1, \ldots, x_k$ for which $x_i^{(p)} = x_i$, and such that $x \to x^{(p)}$ is nilpotent when restricted to V_n.

V_s and V_n are called the **semisimple** and **nilpotent** subspaces of V respectively, and in the case of $\mathrm{Lie}\, G$, for G a finite commutative group scheme, we have:

$$(\mathrm{Lie}\, G)_s = \mathrm{Lie}\, G_{l,r}$$
$$(\mathrm{Lie}\, G)_n = \mathrm{Lie}\, G_{l,l}$$

Proof. See [Mum74, §14, Corollary 1] □

We now return to the case of kernels of isogenies $f : A \to B$ of abelian varieties. As we saw, $\ker f$ and $\ker \widehat{f}$ are naturally Cartier duals of one another. In the specific case of characteristic p, we have the following

DEFINITION 1.3.2.9. An isogeny $f : A \to B$ of abelian varieties is said to be of height one if, denoting by $k(A)$ and $k(B)$ the respective function fields, we have $k(A)^p \subset k(B)$.

We shall show that f is of height one if and only if $\ker f$ is of height one, hence the name.

THEOREM 1.3.2.10. For all abelian varieties A, there is a 1-to-1 correspondence between isogenies $f : A \to B$ of height one (up to isomorphism) and sub p-Lie algebras of $\mathrm{Lie}(A)$.

Proof. Isogenies of height one are uniquely determined up to isomorphism by their kernels, which are subgroup-schemes of the maximal height of one subgroup-scheme $A[p] = \mathrm{Spec}(\mathscr{O}_{A,0}/\mathfrak{m}_{A,0}^{(p)})$. But by Theorem 1.3.2.6, subgroup-schemes $A[p]$ are in natural 1-to-1 correspondence with p-Lie subalgebras of $\mathrm{Lie}\, A_p = \mathrm{Lie}\, A$. □

EXAMPLE 1.3.2.11. This allows us to give an example of an abelian variety A admitting infinitely many distinct isogenies $A \to B$ of height one.

In fact, for every prime $p > 0$, over an algebraically closed field, there is an elliptic curve E unique up to isogeny, such that the pth power map in $\mathrm{Lie}\, E$ is 0 (see Theorem 2.1.9.1). But if E is such a curve and $A = E \times E$, any 1-dimensional subspace of $\mathrm{Lie}(A)$ is stable for the pth power map, and hence defines an isogeny of height one.

The p-rank

We may now introduce the idea of the p-rank of an abelian variety. Let A be an abelian variety of dimension g in characteristic p and $n = p^r m \in \mathbb{Z}$, $r \geq 0$, $m \geq 1$, and $(p, m) = 1$. We want to analyze the structure of the finite group *scheme* $A[n] = \ker[n]_A$. Firstly, the multiplication map $A[m] \times A[p^r] \xrightarrow{m} A[n]$ is an isomorphism since $A[m]$ is the (r, r) part of $A[n]$ and $A[p^r]$ is a product of the $(r, l), (l, r)$, and (l, l) parts of $A[n]$. As we saw in Proposition 1.1.4.11, since the structure of a reduced commutative group scheme is determined by its $\overline{k}$-points, $A[m] \simeq (\mathbb{Z}/m\mathbb{Z})^{2g}$ as a *scheme*.

Thus, it suffices to study the structure of $A[p^n]$, which we rename G_n. Suppose that $(G_1)_{\mathrm{red}} \simeq (\mathbb{Z}/p\mathbb{Z})^r$ for some r (which we may do so since any reduced p-torsion groups are of this form). Since A is "divisible" (see Theorem 1.1.4.2), for any $n \geq 1$, we have an exact sequence

$$0 \to (G_1)_{\mathrm{red}} \to (G_{n+1})_{\mathrm{red}} \xrightarrow{p} (G_n)_{\mathrm{red}} \to 0$$

and one deduces by induction that for any $n \geq 1$, $(G_n)_{\mathrm{red}} = (\mathbb{Z}/p^n\mathbb{Z})^r$. By Corollary 1.2.4.3, $\mathbf{D}(G_n)$ is the kernel of $[p^n]_{\widehat{A}}$, so it follows that there is an integer s such that for any $n \geq 0$, $(\mathbf{D}(G_n))_{\mathrm{red}} = (\mathbb{Z}/p^n\mathbb{Z})^s$. Thus, the decomposition of G_n into its pieces is as follows:

$$\begin{aligned} G_n &= (\mathbb{Z}/p^n\mathbb{Z})^r \times \mathbf{D}(\mathbb{Z}/p^n\mathbb{Z})^s \times G_n^0 \\ &= (\mathbb{Z}/p^n\mathbb{Z})^r \times \mu_{p^n}^s \times G_n^0 \end{aligned}$$

where G_n^0 is local-local. Since G_n is of order p^{2ng}, we deduce that there is an integer $t \geq 0$ such that $r + s + t = 2g$, and G_n^0 is of order p^{nt}.

We will show that $r = r_A$, $s = s_A$, and $t = t_A$ are invariant under isogeny. It suffices to prove this for r and s since $r + s + t = 2g$. Further, since $s_A = r_{\widehat{A}}$, it even suffices to verify this for r. Let $f : A \to B$ be

an isogeny, with kernel of order k. Since $f(A[p^n]) \subset B[p^n]$, we have that the order of $(A[p^n])_{\text{red}}$ is at most k times that of $(B[p^n])_{\text{red}}$, that is $p^{nr_A} \leq kp^{nr_B}$ for all n, hence $r_A \leq r_B$. But by Proposition 2.1.1.4, there is some N such that $[N]_A$ factors as $A \xrightarrow{f} B \xrightarrow{g} A$ with $B \xrightarrow{g} A$ an isogeny, so $r_B \subset r_A$. Thus, r, s, and t are isogeny invariants. In particular, since A and $\widehat{A}$ are isogenous, we have $r_A = r_{\widehat{A}} = s_A$, so that

$$G_n = (\mathbb{Z}/p^n\mathbb{Z})^r \times (\mu_{p^n})^r \times G_n^0$$

with G_n^0 local-local of order p^{nt} and $2r + t = 2g$. In particular, we may deduce that $r \leq g$ (which proves Prop. 1.1.4.12).

DEFINITION 1.3.2.12. The integer r above is called the p-**rank** of A, and is an isogeny invariant.

Since $[p]_A$ induces the 0-map on Lie algebras, we see that $\operatorname{Lie} A = \operatorname{Lie}(\ker[p]_A) = \operatorname{Lie} G_1 = \operatorname{Lie}(\mu_p)^r \oplus \operatorname{Lie} G_1^0$. Since G_1^0 is local-local, the pth power map on $\operatorname{Lie} G_1^0$ is nilpotent, whereas $\operatorname{Lie}(\mu_p)^r$ admits a basis $e_1, \ldots, e_r$ such that $e_i^p = e_i$. We thus see that the p-rank of A is the dimension of the semisimple part of $\operatorname{Lie} A$ with respect to the pth power map. Clearly this also holds for $\widehat{A}$ since the p-rank of A and $\widehat{A}$ is the same. We have already established a canonical isomorphism $\operatorname{Lie} \widehat{A} \simeq H^1(A, \mathscr{O}_A)$ (see Lemma 1.2.2.1). Let $\pi : \mathscr{O}_A \to \mathscr{O}_A$ be the Frobenius homomorphism $\pi(\alpha) = \alpha^p$, which induces a p-linear map also denoted $\pi : H^1(A, \mathscr{O}_A) \to H^1(A, \mathscr{O}_A)$. We claim that the pth power map is sent to π under the isomorphism $\operatorname{Lie} \widehat{A} \simeq H^1(A, \mathscr{O}_A)$. Then we have the following method for computing the p-rank: the p-rank is the dimension of the semi-simple part of $H^1(A, \mathscr{O}_A)$ with respect to the Frobenius map π.

THEOREM 1.3.2.13. Under the natural isomorphism $\operatorname{Lie} \widehat{A} \xrightarrow{\sim} H^1(A, \mathscr{O}_A)$, the pth power operation on $\operatorname{Lie} \widehat{A}$ corresponds to the Frobenius map in $H^1(A, \mathscr{O}_A)$

Proof. See [Mum74, §15, Theorem 3]. □

§1.3.3 Cohomology of Line Bundles

We will discuss some final important results regarding the cohomology of line bundles on abelian varieties. The first of which is:

THEOREM 1.3.3.1 (Riemann-Roch for Abelian Varieties). For all line bundles $\mathscr{L}$ on A, if $\mathscr{L} \simeq \mathscr{O}_A(D)$, we have

$$\chi(\mathscr{L}) = \frac{(D^g)}{g!}$$
$$\chi(\mathscr{L})^2 = \deg \phi_{\mathscr{L}}$$

where (D^g) is the g-fold self-intersection number of D.

Proof. Intersection numbers behave well under change of base field, i.e. $(D \cdot E) = (D_K \cdot E_K)$ for any field extension K/k. Additionally, by compatibility of cohomology with flat base change, $\chi(\mathscr{L}) = \chi(\mathscr{L}_K)$ for any field extension K/k. Thus, we may and do assume $k = \overline{k}$.

If $\mathscr{L}_1$ and $\mathscr{L}_2$ are two line bundles on A such that $\mathscr{L}_1 \otimes \mathscr{L}_2^{-1} \in \operatorname{Pic}^0 A(k)$, then we claim that $\chi(\mathscr{L}_1) = \chi(\mathscr{L}_2)$. Consider the line bundle $\mathscr{M} := \operatorname{pr}_1^*(\mathscr{L}_1) \otimes \mathscr{P}_A$ on $A \times \widehat{A}$. Then the Euler characteristic of its

restriction $\mathscr{M}_{\widehat{a}}$ to $A \times \{\widehat{a}\}$ is locally constant, so independent of $\widehat{a}$ since $\widehat{A}$ is connected. But, $\mathscr{L}_1$ and $\mathscr{L}_2$ are both isomorphic to one of the $\mathscr{L}_{\widehat{a}}$'s, so $\chi(\mathscr{L}_1) = \chi(\mathscr{L}_2)$. Next, if $\mathscr{L}$ is symmetric, $[n]_A^*(\mathscr{L}^m) \simeq \mathscr{L}^{mn^2}$, so

$$\chi(\mathscr{L}^{mn^2}) = \chi([n]_A^*\mathscr{L}^m) = (\deg[n]_A)\chi(\mathscr{L}^m) = n^{2g}\chi(\mathscr{L}^m). \tag{1.3.3-I}$$

Since since $k = \overline{k}$, any $\mathscr{L}$ can be written $\mathscr{L}_1 \otimes \mathscr{L}_2$ where $\mathscr{L}_1$ is symmetric and $\mathscr{L}_2 \in \mathrm{Pic}^0 A$ (Cor. 1.2.5.8), equation 1.3.3-I for any $\mathscr{L}$. Since $\chi(\mathscr{L}^k)$ is a polynomial in k, equation 1.3.3-I shows that $\chi(\mathscr{L}^k)$ is a homogeneous polynomial of degree g. Let ψ be such that

$$\chi(\mathscr{L}^k) = \psi(\mathscr{L})\frac{k^g}{g!}$$

so that $\chi(\mathscr{L}) = \frac{\psi(\mathscr{L})}{g!}$, and we have only to establish that $\psi(\mathscr{L}) = (D^g)$.

Say this holds for $\mathscr{L}$ very ample and let $\mathscr{L}_1, \mathscr{L}_2$ be very ample. Then

$$P(n_1, n_2) = g!\chi(\mathscr{L}_1^{n_1} \otimes \mathscr{L}_2^{n_2})$$

is a polynomial in n_1 and n_2. Since $\chi(\mathscr{L}_1^{kn_1} \otimes \mathscr{L}_2^{kn_2}) = k^g\chi(\mathscr{L}_1^{n_1} \otimes \mathscr{L}_2^{n_2})$, P is homogeneous of degree g. If D_i is the divisor corresponding to $\mathscr{L}_i$, since $\mathscr{L}_1^{n_1} \otimes \mathscr{L}_2^{n_2}$ is very ample for $n_1, n_2 \geq 1$, we see that

$$P(n_1, n_2) = ((n_1D_1 + n_2D_2)^g) = \sum_{i=0}^{g} \binom{g}{i} n_1^i n_2^{g-i}(D_1^i D_2^{g-i})$$

if $n_1, n_2 \geq 1$, and it follows that the same equality holds for all $n_1, n_2 \in \mathbb{Z}$, in particular for $n_1 = 1$ and $n_2 = -1$. But any line bundle $\mathscr{L}$ on A can be written as $\mathscr{L}_1 \otimes \mathscr{L}_2^{-1}$ with $\mathscr{L}_i$ very ample (by divisoral considerations).

It thus only remains to show that $\psi(\mathscr{L}) = (D^g)$ for $\mathscr{L}$ very ample. We can then choose sections $\sigma_0, \ldots, \sigma_g$ of $\mathscr{L}$ on A such that (i) σ_i have no common zero, and (ii) the divisor of zeroes of $\sigma_1, \ldots, \sigma_g$ intersect transversely at (D^g) distinct points. Because of condition (i), we get a morphism $A \xrightarrow{\phi} \mathbb{P}^g$ defined by $a \mapsto (\sigma_0(a), \ldots, \sigma_g(a))$, and by (ii), the 0-cycle $\phi^{-1}((1, 0, \ldots, 0))$ is of degree (D^g), hence ϕ is of degree (D^g). Hence by Lemma 1.1.4.6, $\psi(\mathscr{L})$ is (D^g) times the leading coefficient of the polynomial defined by $g!\chi(\mathscr{O}_{\mathbb{P}^g}(n))$, i.e., $\psi(\mathscr{L}) = (D^g)$.

Next, we have to show that $\chi(\mathscr{L})^2 = \deg \phi_{\mathscr{L}}$. Suppose first that $K(\mathscr{L})$ is finite. Then, by definition of $\phi_{\mathscr{L}}$, we have

$$(1 \times \phi_{\mathscr{L}})^*\mathscr{P}_A \simeq \Lambda(\mathscr{L}) := m^*\mathscr{L} \otimes \mathrm{pr}_1^*\mathscr{L}^{-1} \otimes \mathrm{pr}_2^*\mathscr{L}^{-1}$$

on $A \times A$ by Proposition 1.2.1.2. By the proof of Theorem 1.2.6.3, the higher direct images of $\Lambda(\mathscr{L})$ are supported on $K(\mathscr{L})$. Therefore,

$$\begin{aligned} R^i \mathrm{pr}_{1,*}(\Lambda(\mathscr{L})) &\simeq R^i \mathrm{pr}_{1,*}(m^*\mathscr{L} \otimes \mathrm{pr}_2^*\mathscr{L}^{-1}) \otimes \mathscr{L}^{-1} \\ &\simeq R^i \mathrm{pr}_{1,*}(m^*\mathscr{L} \otimes \mathrm{pr}_2^*\mathscr{L}^{-1}). \end{aligned}$$

It follows that

$$\begin{aligned} \chi(\Lambda(\mathscr{L})) &= \sum_{i=0}^{g}(-1)^i\chi(R^i \mathrm{pr}_{1,*}(\Lambda(\mathscr{L}))) \\ &= \sum_{i=0}^{g}(-1)^i\chi(R^i \mathrm{pr}_{1,*}(m^*\mathscr{L} \otimes \mathrm{pr}_2^*\mathscr{L}^{-1})) \end{aligned}$$

$$= \chi(m^*\mathscr{L} \otimes \mathrm{pr}_2^*\mathscr{L}^{-1}).$$

Since $(m, \mathrm{pr}_2) : A \times A \to A \times A$ is an isomorphism, and $(m, \mathrm{pr}_2)^*[\mathrm{pr}_1^*\mathscr{L} \otimes \mathrm{pr}_2^*\mathscr{L}^{-1}] \simeq m^* \otimes \mathrm{pr}_2^*\mathscr{L}^{-1}$, we find

$$\begin{aligned}\chi(m^*\mathscr{L} \otimes \mathrm{pr}_2^*\mathscr{L}^{-1}) &= \chi(\mathrm{pr}_1^*\mathscr{L} \otimes \mathrm{pr}_2^*\mathscr{L}^{-1}) \\ &= \chi(\mathscr{L})\chi(\mathscr{L}^{-1}) \\ &= (-1)^g\chi(\mathscr{L})^2.\end{aligned}$$

Since $A \times \widehat{A} \simeq A \times (A/K(\mathscr{L}))$, we deduce that

$$\begin{aligned}(-1)^g \deg \phi_{\mathscr{L}} &= \deg \phi_{\mathscr{L}} \chi(\mathscr{P}_A) \\ &= \chi((1 \times \phi_{\mathscr{L}})^* \mathscr{P}_A) \\ &= \chi(\Lambda(\mathscr{L})) \\ &= (-1)^g\chi(\mathscr{L})^2.\end{aligned}$$

Finally, suppose $K(\mathscr{L})$ is not finite. We can choose a finite subgroup $S \subset K(\mathscr{L})$ of arbitrarily large order s. Thus $1 \times \phi_{\mathscr{L}}$ factors as

$$A \times A \to A \times A/S \to A \times \widehat{A},$$

so that since $m^*\mathscr{L} \otimes \mathrm{pr}_2^*\mathscr{L}^{-1} = (1 \times \phi_{\mathscr{L}})^*(\mathscr{P}_A \otimes \mathrm{pr}_1^*\mathscr{L})$, the Euler characteristic is divisible by s (by 1.2.6.7). Since this holds for arbitrarily large s, $\chi(m^*\mathscr{L} \otimes \mathrm{pr}_2^*\mathscr{L}^{-1}) = 0$. But as before, $\chi(m^*\mathscr{L} \otimes \mathrm{pr}_2^*\mathscr{L}^{-1}) = (-1)^g\chi(\mathscr{L})^2$, so $\chi(\mathscr{L}) = 0$ and $\deg \phi_{\mathscr{L}} = 0$. □

We will now prove the second result of this section. This result complements Lemma 1.2.3.3, and together they give heavy restrictions on cohomology at both ends of the spectrum (i.e., $K(\mathscr{L}) = A$ and $K(\mathscr{L})$ finite).

THEOREM 1.3.3.2 (The Vanishing Theorem). Let $\mathscr{L}$ be a non-degenerate line bundle A (i.e., $K(\mathscr{L})$ is finite), there is a unique integer $i(\mathscr{L})$, $0 \leq i(\mathscr{L}) \leq g$ such that $H^p(A, \mathscr{L}) = 0$ for $p \neq i(\mathscr{L})$ and $H^{i(\mathscr{L})}(A, \mathscr{L}) \neq 0$. Further $i(\mathscr{L}^{-1}) = g - i(\mathscr{L})$.

Proof. We may again assume that $k = \overline{k}$ since cohomology behaves well under flat base change. Consider the cartesian diagram

$$\begin{array}{ccc} A \times A & \xrightarrow{\mathrm{pr}_2} & A \\ {\scriptstyle 1\times\phi_{\mathscr{L}}}\downarrow & & \downarrow{\scriptstyle \phi_{\mathscr{L}}} \\ A \times \widehat{A} & \xrightarrow{\mathrm{pr}_2'} & \widehat{A}, \end{array}$$

and note that $\Lambda(\mathscr{L}) \simeq (1 \times \phi_{\mathscr{L}})^*\mathscr{P}_A$ on $A \times A$. Since $\phi_{\mathscr{L}}$ is flat, we have that

$$\begin{aligned}\phi_{\mathscr{L}}^*(R^q \mathrm{pr}_{2,*}'(\mathscr{P}_A)) &\simeq R^q \mathrm{pr}_{2,*}((1 \times \phi_{\mathscr{L}})^*\mathscr{P}_A) \\ &\simeq R^q \mathrm{pr}_{2,*}(\Lambda(\mathscr{L})).\end{aligned}$$

But we have seen in Theorem 1.2.6.3 that $R^q \mathrm{pr}_{2,*}'(\mathscr{P}_A) = 0$ if $q \neq g$ and $R^g \mathrm{pr}_{2,*}'(\mathscr{P}_A)$ is the residue field

$k(0)$ at $0 \in \widehat{A}$. Hence, we deduce that

$$R^q \operatorname{pr}_{2,*}(\Lambda(\mathscr{L})) = \begin{cases} 0 & q \neq g \\ \deg \phi_{\mathscr{L}} & q = g. \end{cases} \tag{1.3.3-II}$$

In this formula (like in the proof of Thm. 1.3.3.1), we may replace $\Lambda(\mathscr{L})$ by $m^*\mathscr{L} \otimes \operatorname{pr}_1^* \mathscr{L}^{-1}$ in equation 1.3.3-II. Taking cohomologies, we get that

$$\dim H^q(A \times A, m^*\mathscr{L} \otimes \operatorname{pr}_1^* \mathscr{L}^{-1}) = \begin{cases} 0 & q \neq g \\ \deg \phi_{\mathscr{L}} & q = g. \end{cases}$$

In this formula, we may (again, like in the proof of Thm 1.3.3.1) replace $m^*\mathscr{L} \otimes \operatorname{pr}_1^* \mathscr{L}^{-1}$ by $\operatorname{pr}_2^* \mathscr{L} \otimes \operatorname{pr}_1^* \mathscr{L}^{-1}$, so using the Künneth formula, we see that if $h^i(\mathscr{L}) = \dim H^i(A, \mathscr{L})$,

$$\sum_{i=0}^{q} h^i(\mathscr{L}) h^{q-i}(\mathscr{L}^{-1}) = \begin{cases} 0 & q \neq g \\ \deg \phi_{\mathscr{L}} & q = g. \end{cases}$$

Since all these h^i's are non-negative, it is easy to see that this can only hold if *only one* of the $h^i(\mathscr{L})$'s is positive and *only one* of the $h^i(\mathscr{L}^{-1})$'s is positive, and that the sum of these two i's is g. □

We will now prove that $i(\mathscr{L})$ can be computed as follows:

THEOREM 1.3.3.3. Let $\mathscr{L}$ be an ample line bundle on an abelian variety A, and $\mathscr{M}$ a non-degenerate line bundle. Let $P(t)$ be the polynomial defined by $P(n) = \chi(\mathscr{L}^n \otimes \mathscr{M})$. Then P has all its g roots real and

$$i(\mathscr{M}) = \#\left(\begin{smallmatrix}\text{positive roots} \\ \text{(with mult.) of } P(t)\end{smallmatrix}\right).$$

This proof is long and will be given in a series of propositions. First note that by the compatibility of cohomology with flat base change, we can and will (until the end of the proof of this theorem) assume $k = \overline{k}$. Before starting, we will introduce some notations. For $a = (a_1, \ldots, a_k) \in \mathbb{Z}^k$, we write $|a| = \sum_{i=1}^k |a_i|$, and if $\mathscr{L}_1, \ldots, \mathscr{L}_k$ are line bundles, we denote the line bundle $\mathscr{L}_1^{a_1} \otimes \cdots \otimes \mathscr{L}_k^{a_k}$ by $\mathscr{L}^a$. We will make heavy use of the following

LEMMA 1.3.3.4. Let X be a projective variety of dimension r and $\mathscr{L}_1, \ldots, \mathscr{L}_k$ line bundles on X. Then there is a constant c depending only on the $\mathscr{L}_i$ such that

$$\dim H^i(\mathscr{L}^a) \leq c(1 + |a|^r)$$

for $i \geq 0$ and all $a \in \mathbb{Z}^k$.

Proof. We can easily reduce ourselves to the case where all the $\mathscr{L}_i$ are very ample. To do so, choose a very ample $\mathscr{L}_{k+1}$ such that $\mathscr{L}_i \otimes \mathscr{L}_{k+1}$ are very ample for $1 \leq i \leq k$, and let $\mathscr{L}_i' = \mathscr{L}_i \otimes \mathscr{L}_{k+1}$, $\mathscr{L}_{k+1}' = \mathscr{L}_{k+1}$. Then there is a linear automorphism T of $\mathbb{Z}^{k+1}$ such that for $a = (a_1, \ldots, a_{k+1}) \in \mathbb{Z}^{k+1}$, $\mathscr{L}_1^{a_1} \otimes \cdots \otimes \mathscr{L}_{k+1}^{a_{k+1}} = (\mathscr{L}')^{T(a)}$ and $\alpha^{-1}|a| \leq |Ta| \leq \alpha|a|$ for suitable $\alpha > 0$.

Thus, we may assume the $\mathscr{L}_i$ are all very ample. We proceed by induction on the integer $\nu = \dim X + k$. When $\nu = 0$, the assertion is trivial. Thus, we may assume $\nu > 0$ and that the assertion holds for all smaller values of ν. If $k = 0$, the assertion is again clear. Thus, suppose $k > 0$ and let $\mathscr{L}'$ be the system

of line bundles $\{\mathscr{L}_1, \dots \mathscr{L}_{k-1}\}$ and for $a = (a_1, \dots, a_k) \in \mathbb{Z}^k$, set $a' = (a_1, \dots, a_{k-1}) \in Z^{k-1}$. Then the existence of a constant c for all $a \in \mathbb{Z}^k$ with $a_k = 0$ follows by induction hypothesis. For any $\mu \in \mathbb{Z}$, we have an exact sequence

$$0 \to (\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu} \to (\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu+1} \to ((\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu+1})|_{\mathcal{H}} \to 0$$

where $\mathcal{H}$ is a hyperplane section for the projective embedding of X given by $\mathscr{L}_k$. Suppose now that $a_k > 0$. We have the exact sequence

$$H^i(X, (\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu}) \to H^i(X, (\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu+1}) \to H^i(\mathcal{H}, ((\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu+1})|_{\mathcal{H}})$$

from which we have

$$\begin{aligned} \dim H^i((\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu+1}) - \dim H^i(X, (\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu}) &\leq \dim H^i(\mathcal{H}, ((\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu+1})|_{\mathcal{H}}) \\ &\leq c(|a|^{r-1} + 1) \end{aligned}$$

for $0 \leq \mu < a_k$ and a suitable constant c, by induction hypothesis. Summing over all μ with $0 \leq \mu < a_k$, we obtain

$$\begin{aligned} \dim H^i((\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{a_k}) &\leq C|a_k|(|a|^{r-1} + 1) + \dim H^i((\mathscr{L}')^{a'}) \\ &\leq C'(|a|^r + 1). \end{aligned}$$

Similarly, if $a_k < 0$, the exactness of

$$H^{i-1}(\mathcal{H}, (\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu+1})|_{\mathcal{H}}) \to H^i(X, (\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu}) \to H^i(X, (\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu+1})$$

gives, for $0 > \mu \leq a_k$, the inequalities

$$\begin{aligned} \dim H^i(X, (\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu}) - \dim H^i(X, (\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu+1}) &\leq \dim H^{i-1}(\mathcal{H}, (\mathscr{L}')^{a'} \otimes \mathscr{L}_k^{\mu+1})|_{\mathcal{H}}) \\ &\leq C(1 + |a|^{r-1}) \end{aligned}$$

and summing over μ with $0 < \mu \leq a_k$, we get as before the required inequality. □

PROPOSITION 1.3.3.5. Let $\mathscr{L}$ be any non-degenerate line bundle on an abelian variety A and $\mathscr{M}$ a very ample line bundle on A, and let $\mathscr{P}_A$ be the polynomial in two variables defined by $P(m, n) = \chi(\mathscr{L}^m \otimes \mathscr{M}^n)$. If $P(1, t) \neq 0$ for $0 \leq t \leq 1$, then $i(\mathscr{L}) = i(\mathscr{L} \otimes \mathscr{M})$.

Proof. Note that if $f : A \to B$ is an isogeny of abelian varieties of degree prime to p and $\mathscr{L}$ a line bundle on B with $\chi(\mathscr{L}) \neq 0$, then $i(f^*(\mathscr{L})) = i(\mathscr{L})$. In fact, in this case, we know that $\mathscr{L}$ is a direct summand of $f_* f^*(\mathscr{L})$, PROVE THIS FACT IN final draft ! and $H^i(A, f^*\mathscr{L}) \simeq H^i(B, f_* f^*\mathscr{L})$, fix in final draft so that if $H^i(A, \mathscr{L}) \neq 0$, then $H^i(B, f_* f^*\mathscr{L}) \neq 0$ and thus $H^i(A, f^*\mathscr{L}) \neq 0$.

Also if $\mathscr{L}_0 \in \mathrm{Pic}^0 A$, then $i(\mathscr{L}) = i(\mathscr{L} \otimes \mathscr{L}_0)$. This is because (by the surjectivity of $\phi_{\mathscr{L}}$ and since we assumed k is algebraically closed) for some $a \in A(k)$, $\mathscr{L} \otimes \mathscr{L}_0 \simeq t_a^*\mathscr{L}$, and since t_a is an automorphism, $H^i(A, \mathscr{L}) \neq 0$ implies $H^i(A, t_a^*\mathscr{L}) \neq 0$ and thus $H^i(A, \mathscr{L} \otimes \mathscr{L}_0) \neq 0$. In particular, since $[n]^*\mathscr{L} = \mathscr{L}^{n^2} \otimes \mathscr{L}_0$ for some $\mathscr{L}_0 \in \mathrm{Pic}^0 A$ (by Lemma 1.2.5.7), we see that $i(\mathscr{L}^{n^2}) = i(\mathscr{L})$ for any $n > 0$, with $p \nmid n$.

Suppose then that N is a large square, prime to p. If $i(\mathscr{L}) \neq i(\mathscr{L} \otimes \mathscr{M})$, then $i(\mathscr{L}^N) \neq i(\mathscr{L}^N \otimes \mathscr{M}^N)$,

so there is a least integer $0 < a \leq N$ with

$$i_1 = i(\mathscr{L}^N \otimes \mathscr{M}^a) \neq i(\mathscr{L}^N \otimes \mathscr{M}^{a-1}) = i(\mathscr{L}^N).$$

(Note that since $P(1,t)$ has no rational zeroes for $0 \leq t \leq 1$, all the bundles $\mathscr{L}^N \otimes \mathscr{M}^a$ have finite $K(\mathscr{L}^N \otimes \mathscr{M}^a)$ so $i(\mathscr{L}^N \otimes \mathscr{M}^a)$ is well-defined). The exact sequence

$$0 \to \mathscr{L}^N \otimes \mathscr{M}^{a-1} \to \mathscr{L}^N \otimes \mathscr{M}^a \to (\mathscr{L}^N \otimes \mathscr{M}^a)|_{\mathcal{H}} \to 0,$$

where $\mathcal{H}$ is a hyperplane section of A for the embedding given by $\mathscr{M}$, gives us that

$$0 \to H^{i_1}(A, \mathscr{L}^N \otimes \mathscr{M}^a) \to H^{i_1}(\mathcal{H}, (\mathscr{L}^N \otimes \mathscr{M}^a)|_{\mathcal{H}})$$

is exact. Thus,

$$\begin{aligned} \dim H^{i_1}(\mathcal{H}, (\mathscr{L}^N \otimes \mathscr{M}^a)|_{\mathcal{H}}) &\geq H^{i_1}(A, \mathscr{L}^N \otimes \mathscr{M}^a) \\ &= |\chi(\mathscr{L}^N \otimes \mathscr{M}^a)| \\ &= N^g P(1, \frac{a}{N}), \end{aligned}$$

but there is a lower bound $P(1,t) \geq c > 0$ if $0 \leq t \leq 1$, and since this holds for arbitrarily large N, while $\dim H = g - 1$, we get a contradiction to Lemma 1.3.3.4. □

PROPOSITION 1.3.3.6. If $\mathscr{L}_1$ and $\mathscr{L}_2$ are two line bundles on an abelian variety and $F(s,t)$ is the homogeneous polynomial defined by $F(m,n) = \chi(\mathscr{L}_1^m \otimes \mathscr{L}_2^n)$, and if $F(t, 1-t) \neq 0$ for $0 \leq t \leq 1$, thn $i(\mathscr{L}_1) = i(\mathscr{L}_2)$.

Proof. Choose a very ample $\mathscr{L}_3$ such that $\mathscr{L}_1 \otimes \mathscr{L}_2^{-1} \otimes \mathscr{L}_3$ is also very ample. Set $f(a,b,c) = \chi(\mathscr{L}_1^a \otimes \mathscr{L}_2^b \otimes \mathscr{L}_3^c)$, so that $f(a,b,0) = F(a,b)$. Since $F(t,1-t) \neq 0$ for $0 \leq t \leq 1$, by continuity, we can choose a square N, prime to p, so large that for $0 \leq r \leq N-1$, $0 \leq t \leq 1$,

$$f(N-r, r, t) = N^g f(1 - \frac{r}{N}, \frac{r}{N}, \frac{t}{N}) \neq 0,$$

$$f(N-r-1+t, r+1-t, t) = N^g f(1 - \frac{r+1-t}{N}, \frac{r+1-t}{N}, \frac{t}{N}) \neq 0.$$

The first equation together with Proposition 1.3.3.5 gives us that for $0 \leq r \leq N-1$, r an integer. Taken together, we get that for r in the same range,

$$i(\mathscr{L}_1^{N-r} \otimes \mathscr{L}_2^r \otimes \mathscr{L}_3) = i(\mathscr{L}_1^{N-r-1} \otimes \mathscr{L}_2^{r+1})$$

for $0 \leq r \leq N-1$, r an integer. Taken together, we get that for r in the same range,

$$i(\mathscr{L}_1^{N-r} \otimes \mathscr{L}_2^r) = i(\mathscr{L}_1^{N-r-1} \otimes \mathscr{L}_2^{r+1}),$$

so that we obtain

$$i(\mathscr{L}_1) = i(\mathscr{L}_1^N) = i(\mathscr{L}_2^N) = i(\mathscr{L}_2).$$

□

Which shows:

COROLLARY 1.3.3.7. If $\mathscr{L}$ is non-degenerate and $n \in \mathbb{Z}_{>0}$,

$$i(\mathscr{L}) = i(\mathscr{L}^n).$$

PROPOSITION 1.3.3.8. If $\mathscr{L}_1$, $\mathscr{L}_2$, and $\mathscr{L}_1 \otimes \mathscr{L}_2$ are non-degenerate,

$$i(\mathscr{L}_1 \otimes \mathscr{L}_2) \leq i(\mathscr{L}_1) + i(\mathscr{L}_2).$$

Proof. Let $\nu : A \times A \to A$ be the morphism $(a, b) \mapsto a - b$, and set $\mathscr{L} = \mathrm{pr}_1^* \mathscr{L}_1 \otimes \mathrm{pr}_2^* \mathscr{L}_2$. Then $\nu^{-1}(0) = \Delta_A \subset A \times A$ and if we identify it with A, $\mathscr{L}|_{\nu^{-1}(0)}$ becomes isomorphic to $\mathscr{L}_1 \otimes \mathscr{L}_2$, which shows that $\mathscr{L}|_{\nu^{-1}(0)}$ is non-degenerate for any $a \in A$ and has index $i = i(\mathscr{L}_1 \otimes \mathscr{L}_2)$. Hence the direct images $R^j\nu(\mathscr{L})$ vanish for $j < i$, and by the Leray spectral sequence $H^p(A \times A, \mathscr{L}) = 0$ for $p < i$. But now,

$$H^{i(\mathscr{L}_1)+i(\mathscr{L}_2)}(A \times A, \mathrm{pr}_1^*(\mathscr{L}_1) \otimes \mathrm{pr}_2^*(\mathscr{L}_2)) = H^{i(\mathscr{L}_1)}(A, \mathscr{L}_1) \otimes H^{i(\mathscr{L}_2)}(A, \mathscr{L}_2),$$

so that $i(\mathscr{L}_1) + i(\mathscr{L}_2) \geq i = i(\mathscr{L}_1 \otimes \mathscr{L}_2)$. □

PROPOSITION 1.3.3.9. Let $\mathscr{L}_1$ and $\mathscr{L}_2$ be non-degenerate line bundles on an abelian variety A such that $\mathscr{L}_1 \otimes \mathscr{L}_2^{-1}$ is ample. If $f(m, n) = \chi(\mathscr{L}_1^m \otimes \mathscr{L}_2^n)$, suppose $f(t, 1-t)$ has a unique zero $\tau \in [0, 1]$ of multiplicity λ. Then

$$0 \leq i(\mathscr{L}_2) - i(\mathscr{L}_1) \leq \lambda.$$

Proof. By Proposition 1.3.3.8, we have

$$i(\mathscr{L}_1) = i(\mathscr{L}_1 \otimes \mathscr{L}_2^{-1} \otimes \mathscr{L}_2) \leq i(\mathscr{L}_1 \otimes \mathscr{L}_2^{-1}) + i(\mathscr{L}_2) = i(\mathscr{L}_2)$$

which proves the left side of the inequality.

Since f is homogeneous and $i(\mathscr{L}^n) = i(\mathscr{L})$ for $n > 0$, we may replace $\mathscr{L}_1$ and $\mathscr{L}_2$ by suitable powers to suppose $\mathscr{L}_1 \otimes \mathscr{L}_2^{-1}$ is very ample. Let us denote by $\mathcal{H}, \mathcal{H}^2, \ldots,$ respectively a hyper plane section of A, a hyperplane section of $\mathcal{H}$, etc. for the projective embedding given by $\mathscr{L}_1 \otimes \mathscr{L}_2^{-1}$.

Let N be any large integer, which we let be coprime to the denominator of τ if $\tau \in \mathbb{Q}$. Then there is a unique integer r with $0 < r \leq N$ such that $\frac{r-1}{N} < \tau < \frac{r}{N}$. Let $s = N - r$, $i_1 = i(\mathscr{L}_1)$, and $i_2 = i(\mathscr{L}_2)$.

We first claim that for $k \leq r - 1$, $l = N - k$, we have

$$H^p(\mathcal{H}^\alpha, \mathscr{L}_1^k \otimes \mathscr{L}_2^l) \stackrel{?}{=} 0$$

if $p \leq i_2 - \alpha - 1$. We will prove this by induction on α. The assertion is obvious for $\alpha = 0$, since by Prop 1.3.3.6,

$$i(\mathscr{L}_1^{r-1} \otimes \mathscr{L}_2^{s+1}) = i(\mathscr{L}_1^{r-2} \otimes \mathscr{L}_2^{s+2}) = \cdots = i(\mathscr{L}_2^N) = i_2.$$

Suppose then that $\alpha > 0$ and that the above holds for $\alpha - 1$. From the exact sequence

$$0 \to (\mathscr{L}_1^{k-1} \otimes \mathscr{L}_2^{l+1})|_{\mathcal{H}^{\alpha-1}} \to (\mathscr{L}_1^k \otimes \mathscr{L}_2^l)|_{\mathcal{H}^{\alpha-1}} \to (\mathscr{L}_1^k \otimes \mathscr{L}_2^l)|_{\mathcal{H}^\alpha} \to 0$$

we get an exact sequence

$$H^p(\mathcal{H}^{\alpha-1}, (\mathscr{L}_1^k \otimes \mathscr{L}_2^l)|_{\mathcal{H}^{\alpha-1}}) \to H^p(\mathcal{H}^\alpha, (\mathscr{L}_1^k \otimes \mathscr{L}_2^l)|_{\mathcal{H}^\alpha}) \to H^p(\mathcal{H}^{\alpha-1}, (\mathscr{L}_1^{k-1} \otimes \mathscr{L}_2^{l+1})|_{\mathcal{H}^{\alpha-1}})$$

and the assertion follows from the inductive hypothesis. But, the exactness of

$$H^p(\mathcal{H}^{\alpha-1}, (\mathscr{L}_1^{r-1} \otimes \mathscr{L}_2^{s+1})|_{\mathcal{H}^{\alpha-1}}) \to H^p(\mathcal{H}^{\alpha-1}, (\mathscr{L}_1^{r} \otimes \mathscr{L}_2^{s})|_{\mathcal{H}^{\alpha-1}}) \to H^p(\mathcal{H}^{\alpha}, (\mathscr{L}_1^{r} \otimes \mathscr{L}_2^{s})|_{\mathcal{H}^{\alpha}})$$

together with the claim gives us that the map

$$H^p(\mathcal{H}^{\alpha-1}, (\mathscr{L}_1^{r} \otimes \mathscr{L}_2^{s})|_{\mathcal{H}^{\alpha-1}}) \to H^p(\mathcal{H}^{\alpha}, (\mathscr{L}_1^{r} \otimes \mathscr{L}_2^{s})|_{\mathcal{H}^{\alpha}})$$

is injective for $p \leq i_2 - \alpha$. Taking $p = i_1$, we deduce that

$$H^{i_1}(A, \mathscr{L}_1^r \otimes \mathscr{L}_2^s) \to H^i(\mathcal{H}^{i_2-i_1}, (\mathscr{L}_1^{r} \otimes \mathscr{L}_2^{s})|_{\mathcal{H}^{i_2-i_2}})$$

is injective. Thus, we deduce that

$$\begin{aligned} \dim H^{i_1}(A, \mathscr{L}_1^r \otimes \mathscr{L}_2^s) &\leq \dim H^i(\mathcal{H}^{i_2-i_1}, (\mathscr{L}_1^{r} \otimes \mathscr{L}_2^{s})|_{\mathcal{H}^{i_2-i_2}}) \\ &= |\chi(\mathscr{L}_1^r \otimes \mathscr{L}_2^s)| \\ &= N^g \left| f(\frac{r}{N}, \frac{s}{N}) \right|. \end{aligned}$$

By the first lemma, we therefore deduce that there is a constant $c > 0$ such that for all large N (prime to the denominator of τ if τ is rational),

$$N^g \left| f(\frac{r}{N}, \frac{s}{N}) \right| \leq cN^{g-(i_2-i_1)},$$

or

$$\left| f(\frac{r}{N}, \frac{s}{N}) \right| \leq \frac{c}{N^{i_2-i_1}}.$$

Now, $f(t, 1-t) = (t-\tau)^\lambda g(t)$ where g is non-vanishing at τ. Thus for all large N as above, there is some $c' > 0$ such that

$$\left| \frac{r}{N} - \tau \right|^\lambda \leq \frac{c'}{N^{i_2-i_1}},$$

or

$$\left| \frac{r}{N} - \tau \right| \leq c'' N^{-(i_2-i_1)/\lambda}.$$

If τ were rational and N coprime to the denominator q of τ, we must have

$$\left| \frac{r}{N} - \tau \right| \geq \frac{1}{Nq}$$

for all r, so that we must have

$$\frac{i_2 - i_1}{\lambda} \leq 1,$$

hence $i_2 \leq i_1 + \lambda$. If τ were irrational, by Kronecker's theorem on the density of fractional parts of $N\tau$ for τ irrational, for a sequence N_i if integers going to infinity, the distance of $N_i\tau$ from the nearest integer tends to $\frac{1}{2}$, so that we again deduce that $\frac{i_2-i_1}{\lambda} \leq 1$, hence $i_2 \leq i_1 + \lambda$. □

Proof of Theorem 1.3.3.3. We finally return to the proof of the theorem. As a reminder, we are given an ample $\mathscr{L}$ and a non-degenerate $\mathscr{M}$. Let $P(n, m) = \chi(\mathscr{L}^n \otimes \mathscr{M}^m)$. We must prove that all the g zeroes of $P(t, 1) = 0$ are real, and that $i(\mathscr{M})$ is the number of positive ones ($t = 0$ is not a real zero, since $\mathscr{M}$ is non-degenerate). Choose a positive integer q so that the real zeroes $t_1, \ldots, t_k$ of $P(t, 1)$ are divided up as

follows:

$$r_1 < \cdots < r_k, r_i \in \mathbb{Z}$$

$$q^{-1}(r_i - 1) < t_i < q^{-1} r_i.$$

Let λ_i be the multiplicity of the root t_i. Then by Propositions 1.3.3.6 and 1.3.3.9, we conclude that if $i(r) = i(\mathscr{L}^r \otimes \mathscr{M}^q)$, then

$$i(R) = i(r_k) \leq i(r_k - 1) = i(r_{k-1}) \leq i(r_{k-1} - 1) = \dots$$

for $R \geq r_k$ and

$$\cdots = i(r_2) \leq i(r_2 - 1) = i(r_1) \leq i(r_1 - 1) = i(S),$$

for $S \leq r_1 - 1$, and

$$i(r_1 - 1) - i(r_l) \leq \lambda_l.$$

But for $r \gg$, $\mathscr{L}^r \otimes \mathscr{M}^q$ is ample, hence $i(r) = 0$. And $r \ll 0$, $(\mathscr{L}^r \otimes \mathscr{M}^q)^{-1}$ is ample, hence $i(r) = g$ by the last statement of the vanishing theorem. Therefore

$$g = i\left(\begin{smallmatrix}\text{very}\\ \text{negative}\\ r\end{smallmatrix}\right) - i\left(\begin{smallmatrix}\text{very}\\ \text{positive}\\ r\end{smallmatrix}\right) = \sum_{l=1}^{k} (i(r_l - 1) - i(r_l)) \leq \sum_{i=1}^{k} \lambda_l.$$

But, $\sum \lambda_l$, the number of real zeroes of $P(t, 1) = 0$ is at most g, so equality holds everywhere, and $i(0) = i(\mathscr{M})$ is just the sum of the λ_l's for which the corresponding root t_l is positive. □

The final result of this section concerns when when an ample line bundle is very ample.

THEOREM 1.3.3.10. For any ample line bundle $\mathscr{L}$ on an abelian variety A, $\mathscr{L}^n$ is very ample if $n \geq 3$.

Proof. Since very ampleness is invariant under change of base field, we may assume $k = \overline{k}$, in particular that k is infinite. For simplicity, we shall prove that $\mathscr{L}^3$ is very ample (since for $n > 3$, the same proof applies). Since $\dim H^0(A, \mathscr{L}) = \chi(\mathscr{L}) > 0$ (by the Vanishing Theorem (1.3.3.2) and the fact that ample line bundles have nonzero global sections), we can choose an effective divisor D such that $\mathscr{O}_A(D)$ is isomorphic to the sheaf of sections of $\mathscr{L}$. Further, D can be chosen so as not to have any multiple components, for if mE occurs in D with E an irreducible divisor and $m > 1$, mE is linearly equivalent to $\sum_{i=1}^{m} t_{a_i}^* E$ for $\sum_{i=1}^{m} a_i = 0$ and for suitable choice of the a_i, the $t_{a_i}^* E$ are all distinct and distinct from the other components of D (this is possible since k is infinite).

Thus, we assume D has no components with multiplicity. Note that for any $a, b \in A(k)$, $t_a^* D + t_b^* D + t_{-a-b}^* D \in |3D|$, the linear system of divisors. We now have to establish the following results:

(1) Given $a_0, a_1 \in A$ with $a_0 \neq a_1$, there is a $D' \in |3D|$ such that $a_0 \in \operatorname{Supp} D'$, $a_1 \notin \operatorname{Supp} D'$.

(2) Given any tangent vector t to A at a_0, there is a $D' \in |3D|$ such that $a_0 \in \operatorname{Supp} D'$ and t is not tangential to D' (i.e., if $\phi = 0$ is a local equation of D' at a_0, $\langle t, d\phi \rangle \neq 0$).

Since we are making these assertions for any ample $\mathscr{L}$, it suffices to prove (1) and (2) for any ample $\mathscr{L}$ with $a_0 = 0$, since the general case follows by applying the result to a translate of $\mathscr{L}$.

Thus, if (1) were not true (with $a_0 = 0$), we would have that for any D as above and any $a, b \in A$, $0 \in \operatorname{Supp} t_a^* D$, so $x_1 \in \operatorname{Supp} t_a^* D \cup \operatorname{Supp} t_b^* D \cup \operatorname{Supp} t_{-a-b}^* D$. Since we may clearly choose b such that

$a_1 \notin \operatorname{Supp} t_b^* D \cup \operatorname{Supp} t_{a+b}^* D$. Thus, $a \in \operatorname{Supp} D$ implies that $a \in t_{a_1}^* \operatorname{Supp} D$. Since the divisor D has no components with multiplicity, this means that $t_{a_1}^* D = D$. In particular, $a_1 \in K(\mathscr{L})$, hence a_1 has finite order. Let a_1 generate the finite group G. We then have an étale morphism $\pi : A \to A/G$, and $D_1 = \pi(\operatorname{Supp} D)$ is a closed subset pure of codimension one in A/G, which we may consider as a divisor with all components of multiplicity one. Since π is étale, $\pi^*(D_1)$ is again a divisor with components of multiplicity one and the same support as D, so $D = \pi^*(D_1)$ and $\mathscr{L} \simeq \pi^*(\mathscr{O}_{A/G}(D_1))$. But note that

$$\begin{aligned} \dim H^0(A, \mathscr{L}) &= \chi(\mathscr{L}) \\ &= (\#G)\chi(\mathscr{O}_{A/G}(D_1)) \\ &= (\#G)\dim H^0(A/G, \mathscr{O}_{A/G}(D_1)) \\ &> \dim H^0(A/G, \mathscr{O}_{A/G}(D_1)). \end{aligned}$$

Since the set of all divisors D_1 such that $\mathscr{L} \simeq \pi^*(\mathscr{O}_{A/G}(D_1))$ fall into a finite set of linear equivalence classes, this proves that all sections $s \in \Gamma(A, \mathscr{L})$ either define multiple divisors, or lie in one of a finite number of lower-dimensional subspaces $\pi^*\Gamma(\mathscr{O}_{A/G}(D_1))$. This is a contradiction, so we have proven (1).

Now, suppose (2) is false for some non-zero tangent vector $t \in \operatorname{Lie}(A)$ and let T be the corresponding invariant vector field. If (2) is false for all divisors $t_a^* D + t_b^* D + t_{-a-b}^* D$, it follows immediately that under the $k[\epsilon]/(\epsilon^2)$-valued automorphism of A defined by t, the divisor D is invariant. This implies that the $k[\epsilon]/(\epsilon^2)$-valued point of A is defined by t is in the subgroup $K(\mathscr{L})$ of points leaving $\mathscr{L}$ invariant. Now in characteristic 0, all group schemes are reduced, so $K(\mathscr{L})$ is finite and discrete and this cannot hold unless $t = 0$. On the other hand, in characteristic $p > 0$, let H be the smallest subgroup of $K(\mathscr{L})$ containing t. Then $H \subset A^{(p)}$ and will be determined by the its Lie algebra $\operatorname{Lie}(H)$ which is the span of t and its pth powers. It is easy to see that D will be invariant under translations by all points of H. In final draft, give explanation pg. 165 of mum Let $A' = A/H$ and $D' = D/H$. Since the quotient by a finite group $\pi : A \to A'$ is flat and surjective, D' is a closed subscheme of A' and $D \simeq D' \otimes_{A'} A$. Therefore if $\mathscr{I}'$ is the sheaf of ideals of D'

$$\mathscr{O}_A(-D) \simeq \mathscr{I}' \otimes_{\mathscr{O}_{A'}} \mathscr{O}_A.$$

Since D is a divisor, $\mathscr{O}_A(-D)$ is a locally free sheaf, and thus $\mathscr{I}'$ is too, i.e. D' is a divisor as well. Now, $D = \pi^*(D')$, so we compute as before:

$$\begin{aligned} \dim H^0(A, \mathscr{L}) &= \deg \pi \dim H^0(A/H, \mathscr{O}_{A/H}(D')) \\ &\geq \dim H^0(A/H, \mathscr{O}_{A/H}(D')). \end{aligned}$$

Exactly as before, this implies that all sections $s \in \Gamma(A, \mathscr{L})$ either define multiple divisors, or lie in one of a finite set of proper subspaces, a contradiction. □

Chapter 2

Tate's Conjecture for Finite Fields & Endomorphisms of Abelian Varieties

"*There is enough of a link between Tate and Hodge that Dinakar Ramakrishnan of CalTech... joked that the person solving Tate should get $500,000.*"

In a Blog Post on *The Tate Conjecture over Finite Fields*
J.S. Milne

The purpose of this chapter is to study endomorphisms with the final goal of proving Tate's Conjecture for finite fields.

Theorem 2.0.0.1 (Tate's Conjecture for Finite Fields)**.** For A, B abelian varieties over $k = \mathbb{F}_q$ finite, and $G_k = \mathrm{Gal}(\overline{k}/k)$, then the natural map

$$T_\ell : \mathbb{Z}_\ell \otimes_{\mathbb{Z}} \mathrm{Hom}_k(A, B) \to \mathrm{Hom}_{\mathbb{Z}_\ell[G_k]}(T_\ell A, T_\ell B)$$

is an isomorphism.

This theorem is a (reformulation of a) special case of the much more general Tate conjecture, which is formulated via fixed points of the Galois action on ℓ-adic cohomology groups. Here, using the tools we built in Chapter 1, we will endow the endomorphism rings with a positive involution, allowing us to classify them via a purely algebraic classification. In line with the theme of how rigid the structure of abelian varieties is, in this chapter, we will see how classifiable their endomorphisms are and how they are determined by seemingly much smaller packets of information, in the form of their Tate modules.

Section 2.1

Algebraic Structure of Hom(A, B)

We begin with a short section of motivation that will give us intuition for Tate modules, the tools which will allow us to study homomorphisms of abelian varieties. We will then continue by studying the basic structure of Hom(A, B) for abelian varieties A and B. In this section, we will reduce this study to Hom$(A, A) =$ End(A) which we will then prove, using the duality theory of abelian varieties developed in the previous chapter, that these endomorphism rings are division algebras with positive involutions (called Rosati involutions). This will then allow us to use the classification of division algebras of finite rank over $\mathbf{Q}$ with positive involutions, a theory developed by Abraham Adrian Albert in the 1930s. This classification will prepare us to prove Tate's Conjecture for finite fields in the following section 2.2.

§2.1.1 Étale Coverings over $\overline{k}$

We will begin with a result over $k = \overline{k}$ to motivate our understanding of the Tate modules $T_\ell(A) := \varprojlim A[\ell^n](\overline{k})$ (see Def 1.1.4.14).

> **Theorem 2.1.1.1 (Serre-Lang).** Let A be an abelian variety and B a variety, both over $k = \overline{k}$. If $f : B \to A$ is an étale covering, then B has a structure of an abelian variety such that f becomes a separable isogeny.

Warning. This result does *not* hold over general fields. For example, it is clearly impossible to give a variety without a k-point the structure of an abelian variety over k.

Proof. Let Γ_m be the graph of the multiplication $m : A \times A \to A$ in $A \times A \times A$, and let Γ' be the scheme-theoretic preimage in $B \times B \times B$. Since we have the following commutative diagram

$$\begin{array}{ccc} \Gamma' & \xrightarrow{\text{étale}} & \Gamma + m \\ {\scriptstyle \mathrm{pr}_{12}}\downarrow & & {\scriptstyle \mathrm{pr}_{12}}\downarrow\wr \\ B \times B & \xrightarrow[f \times f]{\text{étale}} & A \times A \end{array}$$

where the top and bottom arrows are étale coverings, and the right arrow is an isomorphism, we must have that $\mathrm{pr}_{12} : \Gamma' \to B \times B$ is an étale covering too. Choose a point $b_0 \in B(k)$ such that $f(b_0) = 0$ (note that the existence of such a b_0 crucially relies on k being algebraically closed), and let Γ be the connected component of Γ' containing (b_0, b_0, b_0) (which belongs to Γ' since $f(b_0) = 0$ and $(0, 0, 0) \in \Gamma_m$). Then the restriction $p : \Gamma \to B \times B$ of pr_{12} is again an étale covering, so that the degree of p equals the number of k-points of *any* fiber of p. We want to show that p is an isomorphism, or equivalently that there is one point of $B \times B$ whose preimage in Γ is again a single point. Let $\sigma_1, \sigma_2 : B \to \Gamma$ be defined by $\sigma_1(b) = (b_0, b, b)$, $\sigma_2(b) = (b, b_0, b)$ (note that since $\sigma_i(B) \subset \Gamma'$ and $(b_0, b_0, b_0) \in \sigma_i(B)$, it follows that $\sigma_i(B) \subset \Gamma$). Then the restriction of p to $\sigma_2(B)$ is a bijection of $\sigma_2(B)$ onto $B \times \{b_0\}$. It therefore suffices to establish that $p^{-1}(B \times \{b_0\}) = \sigma_2(B)$, or equivalently that if $q := \mathrm{pr}_2 |_\Gamma : \Gamma \to B$ then $q^{-1}(b_0) = \sigma_2(B)$. Since $\sigma_2(B)$ is an irreducible component of $q^{-1}(b_0)$, it suffices to show that $q^{-1}(b_0)$ is irreducible. Now, Γ

is non-singular, being étale over $B \times B$ and hence $A \times A$, and since it is also connected, it is irreducible. Further, the morphism $q : \Gamma \to B$ is smooth, being the composite of an étale morphism $\Gamma \to B \times B$ and the projection $\mathrm{pr}_2 : B \times B \to B$. Finally, $\sigma_1 : B \to \Gamma$ is a section for q. Now the assertion that $q^{-1}(b_0)$ is irreducible follows from

> LEMMA 2.1.1.2. Let $f : X \to Y$ be a proper smooth morphism of (irreducible) varieties such that there is a section $\sigma : Y \to X$, $f \circ \sigma = 1_Y$. Then all fibers of f are irreducible.

Returning to our case, we have shown that $\mathrm{pr}_{12} : \Gamma : B \to B$ is an isomorphism, so that $\nu = \mathrm{pr}_3 \circ \mathrm{pr}_{12}^{-1} : B \times B \to B$ is a morphism. Since, as we say, $\Gamma \supset \sigma_1(B)$ and $\Gamma \supset \sigma_2(B)$, it follows that $\nu(b, b_0) = b = \nu(b_0, b)$. Therefore, we may conclude by the following proposition (which notably only works over $\overline{k}$):

> PROPOSITION 2.1.1.3. Let A be a proper variety over $k = \overline{k}$, $e \in A$ a k-point, and
>
> $$m : A \times A \to A$$
>
> a morphism such that $m(a, e_S) = m(e_S, a) = a$ for all $a \in A(S)$. Then A is an abelian variety with group law m and identity e.

Thus, B is an abelian variety with composition law ν and zero element b_0. Since $f(b_0) = 0$, f is a homomorphism of abelian varieties by Corollary 1.1.1.3. □

To finish the proof we need to settle the

Proof of Lemma 2.1.1.2. We may assume $Y = \operatorname{Spec} R$ affine. Let $S = \Gamma(X, \mathscr{O}_X)$, so that S is an R-algebra which is a domain since X is irreducible and a finite R-module since f is proper. The morphism f factorizes as $X \xrightarrow{g} \operatorname{Spec} S \xrightarrow{h} Y$ where $\operatorname{Spec} S$ is again an irreducible variety. But, $g \circ \sigma$ is a section h, and since $\dim(\operatorname{Spec} S) = \dim Y$, $g \circ \sigma$ is surjective, hence h is an isomorphism, and $R = S$.

Since f is smooth, its fibers are non-singular, and it suffices to show that they are connected. Let $0 \to K_0 \to K_1 \to \cdots$ be a complex of finite free R-modules giving the direct images of $\mathscr{O}_X$ universally, so that by the above, we have an exact sequence

$$0 \to A \to K_0 \to K_1.$$

Let y be a closed point of Y, $\mathfrak{m}$ its maximal ideal in R. Since the completion with respect to the $\mathfrak{m}$-adic topology is an exact functor, we have an exact sequence

$$0 \to \widehat{R} \to \widehat{K_0} \to \widehat{K_1}$$

of the completions, so that

$$\widehat{A} \simeq \varprojlim \ker(K_0/\mathfrak{m}^n K_0 \to K_1/\mathfrak{m}^n K_1).$$

But now,

$$\ker(K_0/\mathfrak{m}^n K_0 \to K_1/\mathfrak{m}^n K_1) = H^0(f^{-1}(y), \mathscr{O}_X/\mathfrak{m}^n \mathscr{O}_X),$$

so the natural map

$$\widehat{A} \xrightarrow{\sim} \varprojlim H^0(f^{-1}(y), \mathscr{O}_X/\mathfrak{m}^n \mathscr{O}_X)$$

is a ring isomorphism. Suppose $f^{-1}(y)$ were not connected, and $f^{-1}(y) = Z_1 \cup Z_2$, for Z_i distinct connected components. We can find a unique $f_n \in H^0(f^{-1}(y), \mathscr{O}_X/\mathfrak{m}^n\mathscr{O}_X)$ which reduces to 1 on Z_1 and 0 on Z_2 and $\{f_n\}$ defines an element $f \in \varprojlim H^0(\mathscr{O}_X/\mathfrak{m}^n\mathscr{O}_X) \simeq \widehat{R}$ with $f^2 = f$ and $f \neq 0, 1$. But this is impossible since $\widehat{R}$ is a local ring. □

Additionally, we will settle the

Proof of Proposition 2.1.1.3. For $a, b \in A(S)$, we shall denote $m(a, b)$ by ab. Let

$$\psi : A \times A \to A \times A$$

given by $\psi(a, b) = (ab, b)$. Then $\psi^{-1}(e, e) = \{(e, e)\}$, so that by the dimension theorem, $\dim(\mathrm{Im}(\psi)) = \dim(A \times A)$. Since $A \times A$ is proper, this implies that ψ is surjective. In particular, since $k = \overline{k}$, given any $a \in A(k)$, there is some $a' \in A(k)$ such that $aa' = e$. Thus, if $\Gamma' = m^{-1}(e)$, and pr_i is the ith projection of $A \times A$, $\mathrm{pr}_2(\Gamma) = A$. Choose an irreducible component Γ of Γ'. Note that $\dim\Gamma \geq \dim A$. If $\mathrm{pr}_i' = \mathrm{pr}_i|_\Gamma$, $(\mathrm{pr}_i')^{-1}(e) = \{(e, e)\}$, so that again by the dimension theorem, $\dim(\mathrm{Im}(p_1')) = \dim A$. Since Γ is proper, this implies that pr_1' is surjective.

Let $\varphi : \Gamma \times A \to A$ be defined by $\varphi((a', a), b) = a'(ab)$. Then $\varphi(\Gamma \times \{e\}) = \{e\}$, so by the Rigidity Theorem 1.1.1.1, $\varphi((a', a), b)) = \varphi((e, e), b)) = b$, that is,

$$a'(ab) = b \tag{2.1.1-I}$$

for all $(a', a) \in \Gamma(S)$, $b \in A(S)$. In particular, if $(a', a) \in \Gamma(S)$, then $a'(aa') = a'$. Choose a pair $(a'', a') \in \Gamma(S)$, and multiply the last equation on the left by a'', to obtain

$$a''(a'(aa')) = a''a'.$$

But $a''a' = e$, and by equation 2.1.1-I, $a''(a'(aa')) = aa'$. Therefore if $(a', a) \in \Gamma(S)$, then not only is $a'a = e$, but also $aa' = e$.

Let $\chi : \Gamma \times X \times X \to X$ be the map defined functorially by $\chi((a', a), b, c) = a((a'b)c)$. Since $\chi(\Gamma \times \{e\} \times \{e\}) = e$, by the Rigidity Theorem 1.1.1.1,

$$a((a'b)c) = e((eb)c) = bc.$$

Multiplying on the left by a' and using equation 2.1.1-I, we get

$$(a'b)c = a'(a((a'b)c)) = a'(bc).$$

Since pr_1' is surjective on k-points, and equality of maps may be checked on geometric points, this shows that multiplication is associative (as it is simply equality of maps in a commutative diagram). Thus, A is an algebraic group with group law m. By homogeneity considerations, since there is a smooth point on any variety over an algebraically closed field, we deduce that A is smooth. Since for any $a_0 \in A(k)$, the translation t_{a_0} is an automorphism, so ψ is an isomorphism. The inverse of ψ is given by $(a, b) \mapsto (ab^{-1}, b)$ on S-points, so $b \mapsto b^{-1}$ is also a morphism, hence A is an abelian variety. □

Note that in general we have:

PROPOSITION 2.1.1.4. Let A, B be abelian varieties and let $f : A \to B$ be an isogeny of degree n. Then there exists a map $g : B \to A$ such that $[n]_A$ factors as:

$$[n]_A : A \xrightarrow{f} B \xrightarrow{g} A.$$

Proof. Because $\#(\ker f) = n$, we have that $[n]_B$ kills $\ker f$ (see [TO70, §1, Deligne's Theorem]), so $\ker f \subset A[n]$. The universal property of quotients then gives the map g. □

Thus, we can now combine Theorem 2.1.1.1 and Proposition 2.1.1.4 to determine the étale fundamental group. Recall that if X is a connected locally noetherian scheme and $x_0 \in X(\overline{k})$ a geometric point. Let $\mathcal{C}$ be the category of pairs (Y, f) such that $f : Y \to X$ is a finite étale morphism from a scheme Y (where morphisms are morphisms $Y \to Y'$ as schemes over X). This category has a natural functor $F : \mathcal{C} \to \text{Set}$ given by

$$F(Y) = \text{Hom}_X(x, Y)$$

which geometrically corresponds to the fiber of $f : Y \to X$ over x, which is represented by x as an X-scheme. F is typically not representable in $\mathcal{C}$, however it is pro-representable in $\mathcal{C}$ by Galois covers of X. I.e., we have a projective system

$$\{X_j \to X_i \mid i < j \in I\}$$

in $\mathcal{C}$, indexed by a directed set I, where each X_i is a Galois cover of X (i.e., a finite étale scheme over X such that $\#\,\text{Aut}_X(X_i) = \deg(X_i/X)$) such that

$$F(Y) = \varinjlim \text{Hom}_{\mathcal{C}}(X_i, Y).$$

For two such X_i, X_j, the map $X_j \to X_i$ induces a group homomorphism $\text{Aut}_X(X_j) \to \text{Aut}_X(X_i)$ which gives us a projective system of automorphism groups. We then make the following definition:

DEFINITION 2.1.1.5. The **étale fundamental group** of X is the

$$\pi_1^{\text{ét}}(X, x) = \varprojlim_{i \in I} \text{Aut}_X(X_i),$$

with the inverse limit topology. We sometimes write $\pi_1(X, x)$ when there is no ambiguity about which fundamental group we are talking about.

Returning to the case of $X = A$ an abelian variety over $k = \overline{k}$, with geometric point $x_0 = e$, all such X_i's are abelian varieties by Theorem 2.1.1.1, and for $f_i : X_i \to A$, $\text{Aut}_A(X_i) = \ker(f_i)(k)$ acting by translations. In particular, we see that $\pi_1(A, e)$ is abelian. To examine it more carefully, it is convenient to break it up into the product of its ℓ-primary pieces for different primes ℓ.

First, suppose $\ell \neq \text{char}\, k =: p$. By Prop. 2.1.1.4 and Thm. 2.1.1.1, the set of étale coverings

$$A \xrightarrow{[\ell^n]} A$$

is cofinal in the set of all étale coverings

$$X_i \xrightarrow{\pi} A$$

such that $\#\ker \pi = \ell^m$ for some m. Therefore, the ℓ-adic component of $\pi_1(A, e)$ is $\varprojlim A[\ell^n](k) = T_\ell \simeq \mathbb{Z}_\ell^{2g}$, the ℓ-adic Tate module of A (see Def 1.1.4.14).

On the other hand, if $\ell = p$, we have that

$$A[p^n] = A[p^n]^0 \times A[p^n]_{\text{red}}$$

where $A[p^n]^0$ is local and $A[p^n]_{\text{red}}$ is reduced (Def 1.3.2.2). Then, as before, by Proposition 2.1.1.4, the set of étale coverings π_n in the diagrams

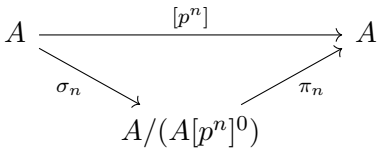

is cofinal in the set of all étale coverings of A whose degree is a power of p. But, $\ker \pi_n \simeq A[p^n]_{\text{red}}$ whose k-points recover the p-adic Tate module $\varprojlim A[p^n]_{\text{red}}(k) = T_p(A) \simeq \mathbb{Z}_p^i$ for some $0 \leq i \leq g$ where i is the p-rank (from §1.3.2).

Combining this gives us that

$$\pi_1(A, e) \simeq \prod_{\text{primes } \ell} T_\ell(A).$$

EXAMPLE 2.1.1.6. If $k = \mathbb{C}$, we have that $A = V/\Lambda$ for V a complex vector space and Λ a full rank $\mathbb{Z}$-lattice. Then the ordinary topological fundamental group $\pi_1^{\text{top}} \simeq \Lambda$ canonically . On the other hand,

$$A[\ell^n] = \frac{1}{\ell^n}\Lambda/\Lambda \subset V/\Lambda = A$$

and it follows that

$$T_\ell(A) \simeq \varprojlim \frac{1}{\ell^n}\Lambda/\Lambda$$

with maps $\frac{1}{\ell^{n+1}}\Lambda/\Lambda \xrightarrow{\ell} \frac{1}{\ell^n}\Lambda/\Lambda$. Or equivalently

$$T_\ell(A) \simeq \varprojlim \Lambda/\ell^n\Lambda$$

with maps $\Lambda/\ell^{n+1}\Lambda \xrightarrow{q} \Lambda/\ell^n\Lambda$ the quotient map. In other words, $T_\ell(X)$ is the ℓ-adic completion of $U = \pi_1^{\text{top}}(A)$, and

$$\begin{aligned}
\pi_1^{\text{ét}}(A) &= \prod_{\text{primes } \ell} T_\ell(A) \\
&= \varprojlim_n \Lambda/n!\Lambda \\
&= \text{pro-finite completion of } \Lambda \\
&= \widehat{\pi_1^{\text{top}}(A)}
\end{aligned}$$

Thus, we may think of the Tate modules as containing information similar to the topological fundamental group $\pi_1^{\text{top}}(A, e) = H_1(A, \mathbb{Z})$.

This example gives us a general framework through which to view Tate modules: they contain information similar to $\pi_1^{\text{top}}(A, e) = H_1(A, \mathbb{Z})$ in the $k = \mathbb{C}$ case. In fact, since this information completely determines the abelian variety over $\mathbb{C}$, it begs the question: how much information can we get about A from $T_\ell(A)$? The rest of this chapter will be dedicated to answering this question. It turns out that, as we will see, we

can deduce a lot of information about an abelian variety from its Tate modules!

§2.1.2 Structure of $\mathrm{Hom}(A, B)$

This section will be dedicated to studying the homomorphisms between abelian varieties. The key idea is to reduce to the category of abelian varieties "up to isogeny." We will show that we may then reduce the study of homomorphisms to those between "simple" abelian varieties, then reduce further to studying the *ring* of endomorphisms (thereby giving us more structure to work with). At the end, we will build up some general structure of associative simple algebras to further study these endomorphisms up to isogeny.

Let A, B be abelian varieties. We let $\mathrm{Hom}(A, B)$ be the group of homomorphisms of $A \to B$ and $\mathrm{End}(A)$ the ring $\mathrm{Hom}(A, A)$. Further, we let $\mathrm{Hom}^0(A, B) := \mathbf{Q} \otimes_{\mathbf{Z}} \mathrm{Hom}(A, B)$ and $\mathrm{End}^0(A) = \mathbf{Q} \otimes_{\mathbf{Z}} \mathrm{End}(A)$.

Now, because composition of homomorphisms extends uniquely to a $\mathbf{Q}$-bilinear map $\mathrm{Hom}^0(A, A') \times \mathrm{Hom}^0(A', A'') \to \mathrm{Hom}^0(A, A'')$, we may form a category whose objects are abelian varieties, and morphisms from $A \to B$ are elements of $\mathrm{Hom}^0(A, B)$. We call this category "abelian varieties up to isogeny." Proposition 2.1.1.4 says exactly that in this new category, isogenies are isomorphisms (in the theorem, f and g are inverses). Henceforth, given an isogeny $f : B \to A$, we let $f^{-1} \in \mathrm{Hom}^0(A, B)$ be its inverse. It is now clear that we may give an alternative equivalent definition:

$$\mathrm{Hom}^0(A, B) = \varinjlim_{\left(\substack{\text{Isogenies} \\ A' \to A}\right)} \mathrm{Hom}(A', B).$$

> THEOREM 2.1.2.1 (Poincaré's complete reducibility theorem). If A is an abelian variety and A' is an abelian subvariety, there is an abelian subvariety A'' such that $A \cap A''$ is finite and $A' + A'' = A$. In other words, such that A is isogenous to $A' \times A''$.

Proof. Choose an ample line bundle $\mathscr{L}$ on A. For any $f : B \to A$, we have

$$\widehat{f} \circ \phi_{\mathscr{L}} \circ f = \phi_{f^*\mathscr{L}} \tag{2.1.2-I}$$

i.e., the diagram commutes

$$\begin{array}{ccc} B & \xrightarrow{f} & A \\ {\scriptstyle \phi_{f^*\mathscr{L}}}\downarrow & & \downarrow{\scriptstyle \phi_{\mathscr{L}}} \\ \widehat{B} & \xleftarrow[\widehat{f}]{} & \widehat{A}. \end{array}$$

This is because for S-points $b \in B(S)$,

$$\begin{aligned} t_b^*(f_S^*\mathscr{L}_S) \otimes f_S^*\mathscr{L}^{-1} &= f_S^* t_{f(b)}^* \mathscr{L}_S \otimes f_S^*\mathscr{L}^{-1} \\ &= f_S^*(t_{f(b)}^*\mathscr{L}_S \otimes \mathscr{L}^{-1}) \\ &= \widehat{f}_S(t_{f(b)}^*\mathscr{L}_S \otimes \mathscr{L}^{-1}) \\ &= (\widehat{f}_S \circ \phi_{\mathscr{L}} \circ f_S)(b) \end{aligned}$$

Applying this to the map $i : A' \hookrightarrow A$, we get the commutative diagram

$$\begin{array}{ccc} A' & \xrightarrow{i} & A \\ {\scriptstyle \phi_{f^*\mathscr{L}}}\downarrow & & \downarrow{\scriptstyle \phi_{\mathscr{L}}} \\ \widehat{A'} & \xleftarrow[\widehat{i}]{} & \widehat{A}. \end{array}$$

Consider $\ker(\widehat{i} \circ \phi_{\mathscr{L}})$. By Theorem 1.2.6.9, we know that $A'' := \ker(\widehat{i} \circ \phi_{\mathscr{L}})^0_{\text{red}}$ is an abelian variety. Since $A'' \cap A' \subset \ker \phi_{i^*\mathscr{L}}$, it is finite since $\ker \phi_{i^*\mathscr{L}}$ is. So, we just need to check that $\dim A'' \geq \dim A - \dim A'$.

Since $\phi_{\mathscr{L}}$ is finite and surjective, it doesn't affect dimension, so we have

$$\dim A'' = \dim \ker \widehat{i} \geq \dim \widehat{A} - \dim \widehat{A'} = \dim A - \dim A',$$

as desired. □

This shows that the category of abelian varieties up to isogeny is a semisimple abelian category, all of whose objects have finite length. We will concretely display this in the following corollaries.

DEFINITION 2.1.2.2. An abelian variety is **k-simple** if it has no abelian subvarieties other than itself and zero over k. Oftentimes, when k is clear, we will just call an abelian variety simple.

WARNING. k-simple does not imply $\overline{k}$-simple.

COROLLARY 2.1.2.3. Any abelian variety A is isogenous to a product

$$A_1^{n_1} \times \cdots \times A_m^{n_m}$$

where each A_i is simple and not isogenous to any A_j, $j \neq i$. Further, the isogeny type of the A_i and the integers n_i are uniquely determined.

Proof. We will prove this by induction on $\dim A$. The base case $\dim A = 0$ is trivial. So, assume this holds for each abelian variety of dimension $< \dim A$. If A is simple, this is trivial, so we may assume A contains a nontrivial abelian subvariety. By Theorem 2.1.2.1, A is isogenous to a product $A' \times A''$ of abelian varieties of smaller dimension. Thus, since the property of being isogenous is transitive, we get our first result. The uniqueness of the isogeny type of the A_i's and n_i's is clear. □

COROLLARY 2.1.2.4. For A simple, the ring $\text{End}^0 A$ is a division algebra. Thus, for any abelian variety

$$A = A_1^{n_1} \times \cdots \times A_m^{n_m}$$

with A_i simple and pairwise non-isogenous, and $D_i = \text{End}^0 A_i$, then

$$\text{End}^0 A = \text{Mat}_{n_1}(D_1) \oplus \cdots \oplus \text{Mat}_{n_m}(D_m)$$

(where $\text{Mat}_m(R)$ is the ring of $m \times m$ matrices over R).

Proof. For A simple, any non-zero endomorphism of A is an isogeny (since the image is an abelian variety), hence an invertible element of $\text{End}^0 A$, which proves the first assertion. As for the second, $\text{Hom}(A_i, A_j) = 0$

for $i \neq j$ (as if the image of a homomorphism weren't 0, it would be an isogeny). Thus,

$$\operatorname{End}^0 A = \bigoplus_{i=1}^{m} \operatorname{End}^0(A_i^{n_i}).$$

It is clear that $\operatorname{End}^0(A_i^{n_i}) = \operatorname{Mat}_{n_i}(D_i)$. □

DEFINITION 2.1.2.5. We say that a function φ defined on a vector space V is a **polynomial function** of degree n if when restricted to any finite-dimensional subspace, it is a polynomial function of degree n or equivalently, if for any $v_0, v_1 \in V$, $\varphi(x_0 v_0 + x_1 v_1)$ is a polynomial in x_0, x_1 of degree n. We define a **homogeneous polynomial function** of degree n similarly, i.e., a polynomial function such that $\varphi(av) = a^n \varphi(v)$ for any $v \in V$.

For instance, in Theorem 1.3.3.1, we have seen that $\chi(\mathscr{L})$ extends to a homogeneous polynomial function of degree g on the vector space $\mathrm{NS}(A) \otimes_{\mathbb{Z}} \mathbb{Q}$.

THEOREM 2.1.2.6. The function $\varphi \mapsto \deg \varphi$ on $\operatorname{End} A$ extends to a homogeneous polynomial function of degree $2g$ on $\operatorname{End}^0 A$.

Proof. Since for $\varphi \in \operatorname{End} A$ and $n \in \mathbb{Z}$,

$$\deg n\varphi = \deg[n] \cdot \deg \varphi = n^{2g} \deg \varphi,$$

it suffices to show that for $\varphi, \psi \in \operatorname{End} A$, the function $P(n) = \deg(n\varphi + \psi)$ is a polynomial function. If $\mathscr{L}$ is an ample line bundle, by 1.2.6.7 we have that

$$\deg(n\varphi + \psi) = \frac{\chi((n\varphi + \psi)^* \mathscr{L})}{\chi(\mathscr{L})}.$$

Therefore it suffices to show that $\chi((n\varphi + \psi)^* \mathscr{L})$ is a polynomial in n. Define $\mathscr{L}_{(n)} = (n\varphi + \psi)^* \mathscr{L}$ and applying Corollary 1.1.3.4 to the morphisms $n\varphi + \psi$, φ, and φ respectively, we get that

$$\mathscr{L}_{(n+2)} \otimes \mathscr{L}_{(n+1)}^{-2} \otimes \mathscr{L}_{(n)} \otimes (2\varphi)^* \mathscr{L}^{-1} \otimes \varphi^* \mathscr{L} \otimes \varphi^* \mathscr{L} = \mathscr{O}_A$$

from which it follows by induction on n that for suitable line bundles $\mathscr{L}_1, \mathscr{L}_2$, and $\mathscr{L}_3$ on A,

$$\mathscr{L}_{(n)} = \mathscr{L}_1^{n(n-1)/2} \otimes \mathscr{L}_2^n \otimes \mathscr{L}_3.$$

Since $\chi(\mathscr{L})$ is a polynomial function of $\mathscr{L}$, $\chi(\mathscr{L}_{(n)})$ is a polynomial in n. □

We would like to prove that $\dim_{\mathbb{Q}} \operatorname{Hom}^0(A, B)$ is finite.

EXAMPLE 2.1.2.7. Taking inspiration from the complex analytic case, if $k = \mathbb{C}$, we know $A_1 = V_1/\Lambda_1$ and $A_2 = V_2/\Lambda_2$ for V_1, V_2 complex vector spaces of dimension g_1, g_2 respectively and Λ_1, Λ_2 full rank $\mathbb{Z}$-lattices. A map $\varphi : V_1 \to V_2$ induces an algebraic homomorphism $f : A_1 \to A_2$ if and only if $f(\Lambda_1) \subset \Lambda_2$. In particular, we see that the morphisms φ are determined by their restrictions to Λ_1, so we have an injection

$$T : \operatorname{Hom}_k(A_1, A_2) \to \operatorname{Hom}_{\mathbb{Z}}(\Lambda_1, \Lambda_2).$$

Since $U_i = \pi_1^{\text{top}}(A_i) = H_1(A_i)$, the map T is just the functorial representation between spaces via maps between π_1's (or H_1's). We have a faithful representation of $\mathrm{Hom}(A_1, A_2)$ by $2g_1 \times 2g_2$-integral matrices. In particular $\mathrm{Hom}(A_1, A_2)$ is a free abelian group of at most rank $4g_1g_2$.

In accordance with the analogy given in example 2.1.1.6, we will in fact be able to replace the role of $H_1(A)$ with $T_\ell(A)$ for $\ell \neq \mathrm{char}\, k$. If A_1 and A_2 are two abelian varieties, every homomorphism $f : A_1 \to A_2$ restricts to maps $f : A_1[\ell^n] \to A_2[\ell^n]$ and hence induces a map:

$$T_\ell(f) : T_\ell(A_1) \to T_\ell(A_2).$$

This functor then gives a canonical homomorphism:

DEFINITION 2.1.2.8. We call the map

$$T_\ell : \mathrm{Hom}_k(A_1, A_2) \to \mathrm{Hom}_{\mathbb{Z}_\ell}(T_\ell(A_1), T_\ell(A_2)),$$

induced by $f \mapsto T_\ell(f)$, the **ℓ-adic representation.**

In fact, in terms of bases of $T_\ell(A_i)$ over $\mathbb{Z}_\ell$, this represents homomorphisms f by $2g_1 \times 2g_2$ matrices with coefficients in $\mathbb{Z}_\ell$ (where $g_i = \dim A_i$). We now show:

THEOREM 2.1.2.9. For any pair of abelian varieties A and B, $\mathrm{Hom}(A, B)$ is finitely generated free abelian group, and for any prime $\ell \neq \mathrm{char}\, k$, the natural map

$$\mathbb{Z}_\ell \otimes_{\mathbb{Z}} \mathrm{Hom}(A, B) \hookrightarrow \mathrm{Hom}_{\mathbb{Z}_\ell}(T_\ell(A), T_\ell(B)) \tag{2.1.2-II}$$

induced by the ℓ-adic representation (Def 2.1.2.8) is injective.

REMARK 2.1.2.10. There is a version for $\ell = p = \mathrm{char}\, k > 0$ using p-divisible groups in place of Tate modules and similarly the ℓ-adic case can be recast in terms of ℓ-divisible groups thus treating all ℓ and characteristics uniformly. However, the injectivity result for $\ell = p$ requires already knowing $\mathbb{Z}$-finiteness of $\mathrm{Hom}(A, B)$, so this step is indeed necessary.

Proof. Without loss of generality, we may assume $k = \overline{k}$. Note that since $\mathrm{Hom}(A, B)$ is torsion-free, we have an inclusion $\mathrm{Hom}(A, B) \subset \mathrm{Hom}^0(A, B)$.

Step I. We claim the following: for any finitely generated subgroup M of $\mathrm{Hom}(A, B)$,

$$\mathbb{Q}M \cap \mathrm{Hom}(A, B) = \{\varphi \in \mathrm{Hom}(A, B) \mid n\varphi \in M, \text{ some } n \neq 0\}$$

is again finitely generated.

To prove this, choose isogenies $\prod A_i \to A$ and $B \to \prod B_j$ where A_i, B_j are simple abelian varieties. Then $\mathrm{Hom}(A, B)$ gets mapped injectively into $\prod_{i,j} \mathrm{Hom}(A_i, B_j)$, so that it suffices to prove this result for A and B simple. If A and B are not isogenous, then $\mathrm{Hom}(A, B) = 0$, so that we may assume they are. In this case, we have an injection $\mathrm{Hom}(A, B) \hookrightarrow \mathrm{End}(A)$ induced by an isogeny $B \to A$. Thus, we are reduced to the case $A = B$ and A simple. By Theorem 2.1.2.6, there is a homogeneous polynomial function P on $\mathrm{End}^0 A$ such that for any $\phi \in \mathrm{End}\, A$, $P(\phi) \deg \phi \in \mathbb{Z}$. Since any $\phi \neq 0$ is an isogeny, $P(\phi) \geq 1$ if $\phi \in \mathrm{End}\, A - 0$. Now, $\mathbb{Q}M$ is a finite dimensional space, and $|P(\phi)| < 1$ is a neighborhood U of 0 in this space. Therefore, $U \cap \mathrm{End}(A) = 0$, so $\mathrm{End}\, A \cap \mathbb{Q}M$ is discrete in $\mathbb{Q}M$ and hence finitely generated.

Step II. We now prove that for any $\ell \neq p$, the map 2.1.2-II is injective. In fact, it suffices to show, in view of Step I, that for any finitely generated (hence free) submodule M of $\mathrm{Hom}(A, B)$ such that $M = \mathbb{Q}M \cap \mathrm{Hom}(A, B)$,

$$\mathbb{Z}_\ell \otimes_{\mathbb{Z}} M \to \mathrm{Hom}_{\mathbb{Z}_\ell}(T_\ell(A), T_\ell(B))$$

is injective. Let $f_1, \ldots, f_m$ be a $\mathbb{Z}$-base of M. If this map is not injective, since the right side is $\mathbb{Z}_\ell$-free, we can find $\alpha_i \in \mathbb{Z}_\ell$ with at least one α_i a unit such that $\sum \alpha_i f_i \mapsto 0$. Hence we can find integers n_i $(1 \leq i \leq m)$ not all 0 mod ℓ such that $T_\ell(\sum_{i=1}^m n_i f_i)$ maps $T_\ell(A)$ into $\ell T_\ell(B)$. By the definition of $T_\ell(f)$, this means that $\sum_{i=1}^m n_i f_i =: f$ maps A_ℓ into 0. But this implies that f factorizes $f : A \xrightarrow{\ell} A \xrightarrow{g} B$, and since $g \in \mathbb{Q}M \cap \mathrm{Hom}(A, B) = M$, $g = \sum_{i=1}^m m_i f_i$. Thus, $\sum n_i f_i = \ell \sum m_i f_i$, but because $f_1, \ldots, f_p$ is a basis of M, this implies that $\ell \mid n_i$ for all i, a contradiction.

The theorem now follows. By the injectivity of the map 2.1.2-II, $\mathrm{Hom}^0(A, B)$ is finite-dimensional over $\mathbb{Q}$ and because of Step I, $\mathrm{Hom}(A, B)$ is finitely generated, and being torsion-free, is also free. □

This proof of finite rank actually gives us a crude bound on the rank:

Corollary 2.1.2.11. $\mathrm{Hom}(A, B) \simeq \mathbb{Z}^\rho$ with $\rho \leq 4\dim A \dim B$.

Proof. The rank is at most that of $\mathrm{Hom}_{\mathbb{Z}_\ell}(T_\ell(A), T_\ell(B)$ which is $4\dim A \dim B$. □

We can also tie up the loose end mentioned at the end of Section 1.2.5.

Corollary 2.1.2.12. For any abelian variety A, the group $\mathrm{NS}(A) = \mathrm{Pic}_{A/k}(\overline{k})/\mathrm{Pic}^0_{A/k}(\overline{k})$ is free of finite rank (called the **base number** of A).

Proof. The homomorphism $\mathscr{L} \mapsto \phi_{\mathscr{L}}$ induces an injection of $\mathrm{NS}(A) \hookrightarrow \mathrm{Hom}(A, \widehat{A})$ by Theorem 1.2.3.5. Note that this also gives a bound of rank $4\dim A^2$. □

Corollary 2.1.2.13. $\mathrm{End}^0 A$ is a finite-dimensional semisimple algebra over $\mathbb{Q}$.

Proof. The "finite-dimensional" part is the only new addition and follows from Theorem 2.1.2.9. The rest follows from our isogeny decomposition of abelian varieties into simple abelian varieties. □

Now that we know that for A simple, $\mathrm{End}^0 A$ is a simple division algebra over k, we seek to generally understand such algebras so we can deduce more about the structure of $\mathrm{End}^0 A$ specifically. Let R be a finite-dimensional associative algebra over a field K.

Definition 2.1.2.14. A **trace form** on R is a K-linear form

$$T : R \to K$$

such that $T(XY) = T(YX)$ for all $X, Y \in R$.

A **norm form** on R is a non-zero polynomial function (see Def 2.1.2.5)

$$N : R \to K$$

such that $N(XY) = N(X)N(Y)$ for all $X, Y \in R$.

The following lemma is well known, but we will give a proof for completeness.

LEMMA 2.1.2.15. Let R be a finite-dimensional associative simple algebra over a field K with center F, separable over K. Then there is a canonical norm form Nm^0 and a canonical trace form Tr^0 of R over K such that any norm form (resp. trace form) of R over K is of the type $(\mathrm{Nm}_{F/K} \circ \mathrm{Nm}^0)^k$ with $k \in \mathbb{Z}_{\geq 0}$ (resp. $\phi \circ \mathrm{Tr}^0$ where $\phi : F \to K$ is a K-linear form). If $[R : K] = d^2$, Nm^0 is homogeneous of degree d.

Proof. When $K = K_s$ is separably closed, $F = K$ (as F/K is separable), then $R \simeq \mathrm{Mat}_d(K)$ for some d. In this case, the elements $[X, Y] = XY - YX$ span the vector subspace of matrices of trace 0, and any norm form gives rise to a rational homomorphism of algebraic groups $\mathrm{GL}_d \to \mathbf{G}_m$. This shows the lemma for $\mathrm{Nm}^0 = \det$, the matrix determinant, and $\mathrm{Tr}^0 = \mathrm{Tr}$, the matrix trace.

In the general case, let K_s be the separable closure of K, $\sigma_i : F \to K_s$ $(1 \leq i \leq [F : K])$ the various embeddings of F in K_s over K, and $K_s^{(i)}$ the field K_s considered as an F-algebra through σ_i. We have an isomorphism of K_s-algebras

$$R \otimes_K K_s \simeq R \otimes_F (F \otimes_K K_s) \simeq \prod_i R \otimes_F K_s^{(i)}.$$

Denote the image of $\alpha \in A \otimes_K K_s$ under this isomorphism by $(\phi_i(\alpha))_i$. If Nm is any norm form on R over K, it extends to a norm form of $R \otimes_K K_s$ over K_s, which defines norm forms Nm_i on each $R \otimes_K K_s^{(i)}$ by

$$\mathrm{Nm}_i(\xi) = \mathrm{Nm}(1, \ldots, 1, \xi, 1, \ldots, 1)$$

(note that a polynomial on a vector space is determined by its points since K_s is infinite). By what we have seen, in the separably closed case, each $\mathrm{Nm}_i = (\mathrm{Nm}_i^0)^{n_i}$, where Nm_i^0 is the reduced norm of $R \otimes_F K_s^{(i)}$ over $K_s^{(i)}$ so that we get

$$\mathrm{Nm}(\alpha) = \prod_i \mathrm{Nm}_i^0(\phi_i(\alpha))^{n_i}.$$

We shall show that the norm $\alpha \mapsto \prod_i \mathrm{Nm}_i^0(\phi_i(\alpha))^{n_i}$ of $R \otimes_K K_s$ over K_s comes from a norm of R over K by base extension if and only if all the n_i are equal. Since K_s is the separable closure of K, Nm comes from K if and only if for all $\gamma \in \mathrm{Gal}(K_s/K)$, $\mathrm{Nm}((1 \otimes \gamma)\alpha) = \gamma(\mathrm{Nm}\,\alpha)$, i.e.,

$$\prod_i \mathrm{Nm}_i^0(\phi_i((1 \otimes \gamma)\alpha))^{n_i} = \gamma \prod_i \mathrm{Nm}_i^0(\phi_i(\alpha))^{n_i}.$$

Now, $\gamma \circ \sigma_i = \sigma_{\pi(i)}$ for some permutation π of $1, \ldots, [F : K]$, and we have the commutative diagram

$$\begin{array}{ccc} R \otimes_K K_s & \xrightarrow{\phi_i} & R \otimes_K K_s^{(i)} \\ {\scriptstyle 1\otimes\gamma}\downarrow & & \downarrow{\scriptstyle 1\otimes\gamma} \\ R \otimes_K K_s & \xrightarrow[\phi_{\pi(i)}]{} & R \otimes_K K_s^{\pi(i)} \end{array}$$

so that (since the second vertical is an isomorphism of simple algebras over the isomorphisms γ of separably closed base fields) we get $\mathrm{Nm}_{\pi(i)}^0\left(\phi_{\pi(i)}((1 \otimes \gamma)(\alpha))\right) = \gamma\,\mathrm{Nm}_i^0(\phi_i(\alpha))$, and on substitution, we see that we must have $n_{\pi(i)} = n_i$ for all i. Now, $\mathrm{Gal}(K_s/K)$ acts transitively on the embeddings over $F \hookrightarrow K_s$, so that we must have all the n_i equal. Thus, we see that we may take $\mathrm{Nm}^0(\alpha) = \prod_i \mathrm{Nm}_i^0(\phi_i(\alpha))$ in the lemma,

and thus, $\mathrm{Nm}(\alpha) = \mathrm{Nm}_{F/K}(\mathrm{Nm}^0_{R/K}(\alpha))^{n_i}$.

For the assertion about the trace, the same Galois descent argument works. □

DEFINITION 2.1.2.16. We define the reduced norm $\mathrm{Nrd}_{R/K} := \mathrm{Nm}_{F/K} \circ \mathrm{Nm}^0$ is called the reduced norm of R over K and similarly, we call $\mathrm{Trd}_{R/K} := \mathrm{Tr}_{F/K} \circ \mathrm{Tr}^0$ the reduced trace of R over K

Now return to our study of the endomorphisms of A with the following important

THEOREM 2.1.2.17. Let $f : A \to A$ be an endomorphism, and $T_\ell(f)$ the induced endomorphism fo $T_\ell(A)$ $(\ell \neq p)$. Then

$$\deg f = \det T_\ell(f)$$

and hence

$$\deg([n]_A - f) = P(n),$$

where $P(t)$ is the characteristic polynomial of $T_\ell(f)$. Furthermore, the polynomial P is monic of degree $2g$, has integer coefficients, and $P(f) = 0$.

Proof. The functions $f \mapsto \deg f$ and $f \mapsto \det T_\ell(f)$ both extend uniquely to norm forms Nm_1 and Nm_2 respectively of degree $2g$ on the semi-simple $\mathbb{Q}_\ell \otimes_{\mathbb{Z}} \mathrm{End}\, A$. If $|\cdot|$ denotes the ℓ-adic absolute value, we claim that $|\mathrm{Nm}_1 \alpha| = |\mathrm{Nm}_2 \alpha|$ for all $\alpha \in \mathbb{Q}_\ell \otimes_{\mathbb{Z}} \mathrm{End}\, A$. It suffices to show this for $\alpha \in \mathbb{Z}_\ell \otimes_{\mathbb{Z}} \mathrm{End}\, A$ by homogeneity of our norms and for $\alpha \in \mathrm{End}\, A$ by continuity (since $\mathbb{Z} \subset \mathbb{Z}_\ell$ is dense). Thus, we have reduced to showing that $|\deg f| = |\det T_\ell(f)|$ for $f \in \mathrm{End}\, A$. The power of ℓ dividing $\deg f$ is the order of the kernel of $A[\ell^n] \xrightarrow{f} A[\ell^n]$ for n sufficiently large. Upon passing to the limit, it is also the order of the kernel of $T_\ell(f)$, which is the order of ℓ occurring in $\det T_\ell(f)$.

Now let $\mathbb{Q}_\ell \otimes_{\mathbb{Z}} \mathrm{End}\, A \simeq \prod_{j=1}^r R_j$ be the decomposition of $\mathbb{Q}_\ell \otimes_{\mathbb{Z}} \mathrm{End}\, A$ into a product of simple algebras. The norms Nm_1 and Nm_2 go over into norms on $\prod_j R_j$, and by Lemma 2.1.2.15, into products

$$\mathrm{Nm}_i(\alpha_1, \ldots, \alpha_r) = \prod \mathrm{Nm}^0_j(\alpha_j)^{\nu_{ij}},$$

for $i = 1, 2$, where Nm^0_j are the norm forms on R_j over $\mathbb{Q}_\ell$ of lowest degree, as in Lemma 2.1.2.15. By taking $\alpha_j = 1$ for all $j \neq j_0$ for some j_0, we see that $|\mathrm{Nm}^0_{j_0}(\alpha_{j_0})|^{\nu_{1j_0} - \nu_{2j_0}} = 1$ for all $\alpha_{j_0} \in R_{j_0}$. Since $\mathrm{Nm}^0_{j_0}$ is homogeneous of positive degree, we see (by multiplying α_{j_0} by ℓ) that $\nu_{1j_0} = \nu_{2j_0}$. Since this holds for all j_0, we deduce that $\mathrm{Nm}_1 = \mathrm{Nm}_2$. Thus, we are done with the first claim of the theorem.

The second claim follows by substituting in $[n]_A - f$ for f and using $T_\ell([n]_A - f) = n \cdot \mathrm{id}_{T_\ell(X)} - T_\ell(f)$. P has to be monic of degree $2g$ and since $P(n) \in \mathbb{Z}$ for all $n \in \mathbb{Z}$, its coefficients are all rational. Further, since $\mathrm{End}\, A$ is a finite $\mathbb{Z}$-module, f is integral over $\mathbb{Z}$, so f and hence $T_\ell(f)$ satisfies a monic polynomial over $\mathbb{Z}$. Hence, the eigenvalues of $T_\ell(f)$ are all algebraic integers, so its characteristic polynomial has coefficients which are algebraic integers. But, since they are rational, they must be in $\mathbb{Z}$. Hence $P(f)$ is a well-defined element of $\mathrm{End}\, A$, and finally, we have that $T_\ell(P(f)) = P(T_\ell(f)) = 0$, so since T_ℓ is injective (by Thm 2.1.2.9), we have that $P(f) = 0$.

□

DEFINITION 2.1.2.18. The above polynomial $P(t) \in \mathbb{Z}[t]$ (and is notably independent of ℓ) is called the **characteristic polynomial of** f. If $P(t) = t^g + a_{g-1}t^{g-1} + \cdots + a_0$, then we define $\operatorname{Nm} f := (-1)^g a_0$ and $\operatorname{Tr} f := -a_{g-1}$ and they are called the **norm** and **trace** of f respectively.

By Lemma 2.1.2.15, we see that if $\operatorname{End}^0 A = R_1 \times \cdots \times R_k$ where R_i are simple algebras over $\mathbb{Q}$, and we denote the components of an $f \in \operatorname{End}^0 A$ in A_i by f_i, then we have

$$\operatorname{Nm}_{\operatorname{End}^0 A/\mathbb{Q}} f = \prod_{i=1}^{k} (\operatorname{Nrd}_{R_i/\mathbb{Q}} f_i)^{n_i},$$

$$\operatorname{Tr}_{\operatorname{End}^0 A/\mathbb{Q}} f = \sum_{i=1}^{k} n_i \operatorname{Trd}_{R_i/\mathbb{Q}} f_i$$

where $n_i \in \mathbb{Z}_{>0}$.

COROLLARY 2.1.2.19. Let A be a simple abelian variety of dimension g, K the center of $\operatorname{End}^0 A$, $[K : \mathbb{Q}] = e$ and $[\operatorname{End}^0 A : K] = d^2$. Then $de \mid 2g$.

Proof. By Lemma 2.1.2.15, $\operatorname{Nm}_{\operatorname{End}^0 A/\mathbb{Q}} = (\operatorname{Nrd}_{\operatorname{End}^0 A/\mathbb{Q}})^n$ for some n. But we know by Theorem 2.1.2.17 is a polynomial of degree $2g$, and again by Lemma 2.1.2.15, we know that $\operatorname{Nrd}_{\operatorname{End}^0 A/\mathbb{Q}}$ is a polynomial of degree de. □

REMARK 2.1.2.20. When $\operatorname{char} k = 0$, we can further say that $d^2e \mid 2g$. For $k = \overline{k}$, by the Lefschetz principle it suffices to treat the case $k = \mathbb{C}$, so in this case, we have $A = V/\Lambda$ for V a complex vector space and Λ a full-rank $\mathbb{Z}$-lattice. Then the division algebra $\operatorname{End}^0 A$ admits a faithful representation in the rational vector space $\Lambda \otimes \mathbb{Q}$, so $\Lambda \otimes \mathbb{Q}$ becomes a module over $\operatorname{End}^0 A$. Hence since $\operatorname{End}^0 A$ is a division algebra, it is free and thus $\dim_{\mathbb{Q}} \Lambda \otimes \mathbb{Q} = 2g$ must be divisible by $\dim_{\mathbb{Q}} \operatorname{End}^0 A = d^2e$. Because this is true for $\mathbb{C}$, by the Lefschetz principle, it is true for general algebraically closed fields of characteristic 0. Now since $\operatorname{End}^0 A \hookrightarrow \operatorname{End}^0 A_{\overline{k}}$ is an injective algebra map, we have that $\dim_{\mathbb{Q}} \operatorname{End}^0 A$ divides $\dim_{\mathbb{Q}} \operatorname{End}^0 A_{\overline{k}}$ and thus this result holds for all fields of characteristic 0.

This is false in characteristic p, even for algebraically closed fields. For all $p > 0$, we will see in §2.1.9 that there exists an elliptic curve E with p-rank 0 and that for such a curve $\operatorname{End}^0 A$ is non-commutative of rank 4 with center $\mathbb{Q}$. Thus, in this case, $de = 2 = 2g$.

DEFINITION 2.1.2.21. We say that a simple abelian variety A is **of CM-type** if $de = 2g$ where as above, for K the center of $\operatorname{End}^0 A$, $d^2 = [\operatorname{End}^0 A : K]$, $e = [K : \mathbb{Q}]$, and $g = \dim A$. (CM stands for complex multiplication).

In a division algebra D of rank d^2 over its center K, it is well known ([Stacks, Lemma 074V]) that all maximal commutative subfields have degree d over K. Thus, a simple abelian variety A is of CM-type if and only if $\operatorname{End}^0 A$ admits a subfield of degree $2g$. Note also that in the $k = \mathbb{C}$ case, by remark 2.1.2.20, this implies that $d = 1$, so if $A/\mathbb{C}$ is CM-type, $\operatorname{End}^0 A$ is a field of degree $2g$ over $\mathbb{Q}$.

§2.1.3 Riemann Forms

Riemann forms give us an invaluable way to study the structure of Tate modules, which by Theorem 2.1.2.9 will allow us to give structure to endomorphism rings. The key idea is to use the Weil pairing introduced in §1.3.1, together with the ϕ construction to give us a bilinear form on Tate modules.

Let $\ell \neq p$ be a prime different from $\operatorname{char} k$, let μ_{ℓ^n} be the group scheme of ℓ^nth roots of unity, and as in §1.3.1, let $\mathbb{Z}_\ell(1) = \varprojlim \mu_{\ell^n}$ (the projective system is given by ℓth power maps). Define $\mathbf{M}_\ell = \mathbb{Z}_\ell(1)(\overline{k})$, which we note is (*non-canonically*) isomorphic to $\mathbb{Z}_\ell$ as the algebraic closure contains all ℓ^nth roots of unity. In §1.3.1, we introduced the Weil Paring $\langle \cdot, \cdot \rangle_{A,n} : A[n] \times \widehat{A}[n] \to \mu_n$, and extending to the inverse limit (and restricting our view to $\overline{k}$-points), we get a $\mathbb{Z}_\ell$-bilinear non-degenerate pairing

$$\langle \cdot, \cdot \rangle_{A,\ell^\infty} : T_\ell(A) \times T_\ell(\widehat{A}) \to \mathbf{M}_\ell \simeq \mathbb{Z}_\ell.$$

We may now give the

Definition 2.1.3.1. Let $\mathscr{L}$ be a line bundle on an abelian variety A, and $\ell \neq p$ a prime number. We define the **Riemann form** $E^{\mathscr{L}}$ of $\mathscr{L}$ to be the $\mathbb{Z}_\ell$-bilinear form

$$E^{\mathscr{L}} : T_\ell(A) \times T_\ell(A) \to \mathbf{M}_\ell$$

given by $E^{\mathscr{L}}(a, b) = \langle a, T_\ell(\phi_{\mathscr{L}})(b) \rangle_{A,\ell^\infty}$.

We will prove the general result

Theorem 2.1.3.2. Let $f : A \to \widehat{A}$ be a homomorphism and define $e_{f,\ell^\infty} : T_\ell(A) \times T_\ell(A) \to \mathbf{M}_\ell$ defined by $e_{f,\ell^\infty}(a, b) = \langle a, f(b) \rangle_{A,\ell^\infty}$. Then f is symmetric if and only if $e_{f,\mathscr{L}}$ is skew-symmetric.

With obvious corollary (using Prop. 1.2.6.15)

Corollary 2.1.3.3. The Riemann form $E^{\mathscr{L}}$ of any line bundle $\mathscr{L}$ is skew-symmetric.

Proof. By Prop 1.2.6.15, $\phi_{\mathscr{L}}$ is symmetric. Since $E^{\mathscr{L}} = e_{\phi_{\mathscr{L}},\ell^\infty}$, we are done. □

This result will essentially fall out from the following two lemmas. We begin with a fact we mentioned (without proof) in the warning after Theorem 1.3.1.2. We prove it here:

Lemma 2.1.3.4. Let $f : A \to B$ be an isogeny,

$$\langle \cdot, \cdot \rangle_f : \ker f \times \ker \widehat{f} \to \mathbf{G}_m$$

and

$$\langle \cdot, \cdot \rangle_{\widehat{f}} : \ker \widehat{f} \times \ker \widehat{\widehat{f}} \to \mathbf{G}_m$$

be the duality pairings given in Theorem 1.3.1.2, let $\sigma : \ker f \times \ker \widehat{f} \xrightarrow{\sim} \ker \widehat{f} \times \ker f$ be the swap isomorphism, and $\iota_A : A \to \widehat{\widehat{A}}$ be the canonical isomorphism given by $\mathscr{P}_A$. Then for all k-schemes T

and any $a \in (\ker f)(T)$ and $\widehat{b} \in (\ker \widehat{f})(T)$,

$$\langle a, \widehat{b} \rangle_f = \langle \widehat{b}, \iota_A(a) \rangle_{\widehat{f}}^{-1}.$$

Proof. Consider $\mathscr{Q} = (f \times 1)^* \mathscr{P}_B \simeq (1 \times \widehat{f})^* \mathscr{P}_A$ (the isomorphism given in Lem 1.2.6.5). We first note that $\mathscr{Q}|_{A \times \ker \widehat{f}} \simeq \mathscr{O}_{A \times \widehat{B}}$ as we saw in the proof of Lemma 1.2.6.5. Fix a trivializing section σ along $A \times \ker \widehat{f}$. For any k-scheme T, let $a \in (\ker f)(T)$ and $\widehat{b} \in (\ker \widehat{f})(T)$, under the canonical isomorphism $h_{a,A} : t_{(a,0)}^*(\mathscr{Q}_T) \simeq \mathscr{Q}_T$ over $A_T \times \widehat{B}_T$ has pullback along $1 \times \widehat{b} : A_T \to A_T \times \widehat{B}_T$ is given by $t_{(a,0)}^*(\sigma) \mapsto \langle a, \widehat{b} \rangle_f \cdot \sigma_T$. This is because $(\mathscr{Q}_T, \sigma_T)$ pulls back along $1 \times \widehat{b}$ to $(\mathscr{L}_{\widehat{b},T}, \sigma_T|_{A \times \{\widehat{b}\}})$ (the line bundle on B_T associated to the point $\widehat{b}$). Then, the canonical isomorphism $t_{(a,0)}^*(\mathscr{Q}_T) \simeq \mathscr{Q}_T$ pulls back to the canonical isomorphism $t_a^*(\mathscr{L}_{\widehat{b},T}) \simeq \mathscr{L}_{\widehat{b},T}$ which acts by $t_a^*(\sigma_T|_{A \times \{\widehat{b}\}}) \mapsto \langle a, \widehat{b} \rangle_f \cdot \sigma_T|_{A \times \{\widehat{b}\}}$ by the explicit description of the pairing (equation 1.3.1-I).

We will now show that the following diagram of isomorphisms of line bundles over $(A \times \widehat{B})_T$ is commutative

$$\begin{array}{ccccc}
\mathscr{Q}_T & \xleftarrow{h_{A,a}} & t_{(a,0)}^*(\mathscr{Q}_T) & \xleftarrow{t_{(a,0)}^*(h_{\widehat{B},\widehat{b}})} & t_{(a,0)}^* t_{(0,\widehat{b})}^*(\mathscr{Q}_T) \\
\| & & & & \uparrow{\scriptstyle \langle a,\widehat{b}\rangle_f} \\
\mathscr{Q}_T & \xleftarrow{h_{\widehat{B},\widehat{b}}} & t_{(0,\widehat{b})}^*(\mathscr{Q}_T) & \xleftarrow{t_{(0,\widehat{b})}^*(h_{A,a})} & t_{(0,\widehat{b})}^* t_{(a,0)}^*(\mathscr{Q}_T)
\end{array}$$

It suffices to work over $(A \times \ker \widehat{f})_T$ because each of the homomorphisms preserves this subgroup and are therefore determined by what it does on this subgroup by translation. Then, pulling back over $1 \times y$ for any $y \in (\ker \widehat{f})(T)$ makes the lower right map into multiplication by the unit $\langle a, \widehat{b} + y \rangle_f$ relative to σ because the pullback $(1 \times y)^* t_{(0,\widehat{b})}^* = (1 \times (y + \widehat{b}))^*$, and as we saw above, the pullback of $h_{A,a}$ along $(1 \times (y + \widehat{b}))^*$ makes it into multiplication by $\langle a, \widehat{b} + y \rangle_f$ relative to σ.

The upper right map similarly becomes $(1 \times y)^* t_{(a,0)}^* h_{\widehat{B},\widehat{b}} = t_{(a,0)}^* (1 \times y)^* h_{\widehat{B},\widehat{b}})$ and the lower left pulls back to $(1 \times y)^* h_{\widehat{B},\widehat{b}}$. Since they both act by multipliers on σ and the pullback along t_a^* doesn't affect this multiplier, they give the same multiplier on σ and hence cancel out. Finally, the upper left map gives $(1 \times y)^* h_{A,a} = \langle a, y \rangle_f$ on sigma, and by the bimultiplicativity of the pairing $\langle \cdot, \cdot \rangle_f$, we have that the composite effect of the maps is the same. In summary we have that the following diagram of effects on sections commutes

$$\begin{array}{ccccc}
\sigma_T & \xleftarrow{\langle a,y\rangle_f} & t_{(a,0)}^*(\sigma_T) & \xleftarrow{t_{(a,0)}^*((1\times y)^* h_{\widehat{B},\widehat{b}})} & t_{(a,0)}^* t_{(0,\widehat{b})}^*(\sigma_T) \\
\| & & & & \uparrow{\scriptstyle \langle a,\widehat{b}\rangle_f} \\
\sigma_T & \xleftarrow{(1\times y)^* h_{\widehat{B},\widehat{b}}} & t_{(0,\widehat{b})}^*(\sigma_T) & \xleftarrow{\langle a,\widehat{b}+y\rangle_f} & t_{(0,\widehat{b})}^* t_{(a,0)}^*(\sigma_T)
\end{array}$$

Finally, replacing $f : A \to B$ by $\widehat{f} : \widehat{B} \to \widehat{A}$, we let that $\mathscr{Q}' = (\widehat{f} \times 1)^*(\mathscr{P}_{\widehat{A}}) \simeq (1 \times \widehat{\widehat{f}})^*(\mathscr{P}_{\widehat{B}})$ on $\widehat{\widehat{B}} \times \widehat{A}$. Under the identification of $\widehat{\widehat{B}} \times \widehat{A}$ with $B \times \widehat{A} \simeq (A \times \widehat{B})^\wedge$ and $\mathscr{Q}$ and $\mathscr{Q}'$ are duals under this identification,

and looking at the right map in the above diagram gives us

$$\begin{array}{c} t^*_{(a,0)} t^*_{(0,\widehat{b})}(\mathscr{Q}_T) \\ \uparrow \langle a,\widehat{b}\rangle_f \\ t^*_{(0,\widehat{b})} t^*_{(a,0)}(\mathscr{Q}_T), \end{array}$$

but dualizing the maps in the above commutative diagrams, we get that the right vertical map on the right becomes

$$\begin{array}{c} t^*_{(\iota_A(a),0)} t^*_{(0,\widehat{b})}(\mathscr{Q}'_T) \\ \downarrow \langle \widehat{b},\iota_A(a)\rangle_f \\ t^*_{(0,\widehat{b})} t^*_{(\iota_A(a),0)}(\mathscr{Q}'_T). \end{array}$$

Thus, since these maps are equal (by the fact that the larger diagram has to commute, and units are unaffected by dualizing), and the maps go the opposite way, we must have that $\langle a, \widehat{b}\rangle_f = \langle \widehat{b}, \iota_A(a)\rangle^{-1}_{\widehat{f}}$. □

LEMMA 2.1.3.5. Let $h : A \to B$ be a homomorphism. Then, $T_\ell(h)$ and $T_\ell(\widehat{h})$ are adjoint with respect to the Weil pairings $\langle\cdot,\cdot\rangle_{A,\ell^\infty}$ and $\langle\cdot,\cdot\rangle_{B,\ell^\infty}$

Proof. We are trying to prove that for all $a \in T_\ell(A)$ and $\widehat{b} \in T_\ell(\widehat{B})$,

$$\langle T_\ell(h)(a), \widehat{b}\rangle_{B,\ell^\infty} = \langle a, T_\ell(\widehat{h})(\widehat{b})\rangle_{A,\ell^\infty}.$$

To see this, it suffices to check for each $n \geq 1$ and $a \in A[\ell^n]$, $\widehat{b} \in \widehat{B}[\ell^n]$ that

$$\langle h(a), \widehat{b}\rangle_{B,\ell^n} = \langle a, \widehat{h}(\widehat{b})\rangle_{A,\ell^n}.$$

But by the explicit description of the Weil pairing (equation 1.3.1-I), for the pair $(\mathscr{L}, \sigma)$, $\mathscr{L} \in \mathrm{Pic}_{B/k}(k)$ together with trivialization $\sigma : e^*\mathscr{L} \simeq \mathscr{O}_{\mathrm{Spec}\, k})$ corresponding to $\widehat{b}$, we get

$$\langle h(a), \widehat{b}\rangle_{B,\ell^n} = \frac{t_{h(a)^*(\sigma)}}{\sigma},$$

and since $\widehat{h}(\widehat{b})$ corresponds to $(h^*\mathscr{L}, h^*\sigma)$,

$$\langle a, \widehat{h}(\widehat{b})\rangle_{A,\ell^n} = \frac{t^*_x(h^*\sigma)}{h^*\sigma}$$

but $t^*_x(h^*\sigma) = h^* t^*_{h(a)}\sigma$, so $\frac{t^*_x(h^*\sigma)}{h^*\sigma} = \frac{t^*_{h(a)}\sigma}{\sigma}$ because h^* acts on a trivialization (which is just a choice of basis for $\mathscr{L}(e)$ as we are working over $\mathrm{Spec}\, k$) by some scalar $\alpha \in k^\times$. Thus, we have proven the adjunction of $T_\ell(h)$ and $T_\ell(\widehat{h})$ □

We can now return to the

Proof of Theorem 2.1.3.2. First suppose h is symmetric. We would like to prove that e_{h,ℓ^∞} is skew symmetric, i.e., that for all $a, b \in T_\ell(A)$,

$$e_{h,\ell^\infty}(a,b) := \langle a, T_\ell(h)(b)\rangle_{A,\ell^\infty} = \langle b, T_\ell(h)(a)\rangle^{-1}_{A,\ell^\infty} =: e_{h,\ell^\infty}(b,a)^{-1}.$$

Note that it suffices to check this on the Weil pairings for ℓ^n for all n large enough so that $a,b \in A[\ell^n](\overline{k})$. By the definition of symmetric (1.2.6.13), $h = \widehat{h} \circ \iota_A$. Thus,

$$\langle a, h(b)\rangle_{A,\ell^n} = \langle a, \widehat{h}(\iota_A(b))\rangle_{A,\ell^n}.$$

By Lemma 2.1.3.5, $T_\ell(h)$ and $T_\ell(\widehat{h})$ are adjoint to one another, so

$$\langle a, \widehat{h}(\iota_A(b))\rangle_{A,\ell^n} = \langle h(a), \iota_A(b)\rangle_{\widehat{A},\ell^n}.$$

But, applying Lemma 2.1.3.4 to $[m]_A : A \to A$, for any $a \in A[m](T)$ and $\widehat{b} \in \widehat{A}[m](T)$ for a k-scheme T we have that the Weil pairing

$$\langle a, \widehat{b}\rangle_{A,m}^{-1} = \langle \widehat{b}, \iota_A(a)\rangle_{\widehat{A},m}.$$

So, returning to our case, we have

$$\langle h(a), \iota_A(b)\rangle_{\widehat{A},\ell^n} = \langle b, h(a)\rangle_{A,\ell^n}^{-1},$$

and passing back to the limit we have

$$\langle a, T_\ell(h)(b)\rangle_{A,\ell^\infty} = \langle b, T_\ell(h)(a)\rangle_{A,\ell^\infty}^{-1},$$

as desired.

Conversely, suppose e_{h,ℓ^∞} is anti-symmetric, so for all $a,b \in T_\ell(A)$,

$$e_{h,\ell^\infty}(a,b) := \langle a, h(b)\rangle_{A,\ell^\infty} = \langle b, h(a)\rangle_{A,\ell^\infty}^{-1} =: e_{h,\ell^\infty}(b,a)^{-1}.$$

Running the argument backwards gives us that for all $a,b \in T_\ell(A)$,

$$\langle a, T_\ell(h)(b)\rangle_{A,\ell^\infty} = \langle a, T_\ell(\widehat{h})(T_\ell(\iota_A)(b))\rangle_{A,\ell^\infty},$$

but since the pairing is perfect for $\ell \neq \mathrm{char}\, k$, this implies that $T_\ell(h) = T_\ell(\widehat{h}) \circ T_\ell(\iota_A) = T_\ell(\widehat{h} \circ \iota_A)$ by functoriality of T_ℓ. By Theorem 2.1.2.9, this implies that $h = \widehat{h} \circ \iota_A$, i.e., that h is symmetric. □

Returning to the specific case of $E^{\mathscr{L}}$, we have now that $\mathscr{L} \mapsto E^{\mathscr{L}}$ induces a map:

$$\mathrm{Pic}_{A/k}(k) \Big/ \mathrm{Pic}^0_{A/k}(k) \to \left\{ \begin{array}{c} \text{Alternating 2-forms} \\ T_\ell(A) \times T_\ell(A) \to \mathbf{M}_\ell \end{array} \right\}$$

which is injective since $E^{\mathscr{L}} = 0$ implies $\phi_{\mathscr{L}} = 0$ (because $\langle\cdot,\cdot\rangle_{A,\ell^\infty}$ is nondegenerate), and $\phi_{\mathscr{L}} = 0$ if and only if $\mathscr{L} \in \mathrm{Pic}^0_{A/k}(k)$ by Theorem 1.2.3.5. Note that for $k = \overline{k}$, this gives an injection of $\mathrm{NS}(A)$ into the right set since $T_\ell(A) = T_\ell(A_{\overline{k}})$.

We will now record some useful properties of $E^{\mathscr{L}}$:

Proposition 2.1.3.6. Let $\mathscr{L}$ be a line bundle on A, $E^{\mathscr{L}}$ its Riemann form, $a,b \in T_\ell(A)$, and $\widehat{a},\widehat{b} \in T_\ell(\widehat{A})$. Then we have

$$E^{f^*\mathscr{L}}(a,b) = E^{\mathscr{L}}(T_\ell f(a), T_\ell f(b)). \tag{2.1.3-I}$$

Considering the Poincaré bundle $\mathscr{P}_A$ and identifying $T_\ell(A \times \widehat{A})$ with $T_\ell(A) \times T_\ell(\widehat{A})$, we have

$$E^{\mathscr{P}_A}((a,\widehat{a}),(b,\widehat{b})) = \langle a, \widehat{b}\rangle_{A,\ell^\infty} - \langle b, \widehat{a}\rangle_{A,\ell^\infty}. \tag{2.1.3-II}$$

Proof. For equation 2.1.3-I, we have

$$\begin{aligned} E^{f^*\mathscr{L}}(a,b) &= \langle a, T_\ell(\phi_{f^*\mathscr{L}})b\rangle_{A,\ell^\infty} \\ &= \langle a, T_\ell(\widehat{f}) \circ T_\ell(\phi_{\mathscr{L}}) \circ T_\ell(f)(b)\rangle_{A,\ell^\infty} \\ &= \langle T_\ell(f)(a), T_\ell(\phi_{\mathscr{L}}) \circ T_\ell(f)(b)\rangle_{A,\ell^\infty} \\ &= E^{\mathscr{L}}(T_\ell(f)(a), T_\ell(f)(b)). \end{aligned}$$

Where the second equality is by equation 2.1.2-I and the third is by Lemma 2.1.3.5.

We may now move on to showing equation 2.1.3-II. By skew-symmetry of $E^{\mathscr{P}_A}$ (Cor 2.1.3.3) and linearity, it suffices to show that

$$E^{\mathscr{P}_A}((a,0),(b,0)) = E^{\mathscr{P}_A}((0,\widehat{a}),(0,\widehat{b})) = 0$$

and

$$E^{\mathscr{P}_A}((a,0),(0,\widehat{b})) = \langle a, \widehat{b}\rangle_{A,\ell^\infty}.$$

Using the functoriality property of E for the inclusion of $A\times\{0\}$ in $A\times\widehat{A}$ and the triviality of the restriction of $\mathscr{P}_A$ to $A \times \{0\}$, it follows that $E^{\mathscr{P}_A}((a,0),(b,0)) = 0$ and similarly $E^{\mathscr{P}_A}((0,\widehat{a}),(0,\widehat{b})) = 0$. To prove the last assertion, first note that we have an isomorphism $(A \times B)^\wedge \xrightarrow{\sim} \widehat{A} \times \widehat{B}$ which is given by the map of line bundles $\mathscr{L} \mapsto (\mathscr{L}|_{A\times\{0\}}, \mathscr{L}|_{\{0\}\times B})$ for $\mathscr{L} \in \mathrm{Pic}^0(A \times B)$. In particular, taking $B = \widehat{A}$, we have an identification of $(A \times \widehat{A})^\wedge$ with $\widehat{A} \times \widehat{\widehat{A}}$.

Now, for any $(a, \widehat{a}) \in A \times \widehat{A}$, $\phi_{\mathscr{P}_A}((a,\widehat{a}))$ is given by the line bundle

$$\left(T^*_{(a,\widehat{a})}\mathscr{P}_A \otimes \mathscr{P}^{-1}|_{A\times\{0\}}, T^*_{(a,\widehat{a})}\mathscr{P}_A \otimes \mathscr{P}^{-1}|_{\{0\}\times\widehat{A}}\right) \simeq \left(\mathscr{P}_A|_{A\times\{\widehat{a}\}}, \mathscr{P}_A|_{\{a\}\times\widehat{A}}\right).$$

Therefore, $\phi_{\mathscr{P}_A}((a,\widehat{a})) = (\widehat{a}, \iota_A(a))$. Thus, we obtain

$$E^{\mathscr{P}_A}((a,0),(0,\widehat{b})) = \langle (a,0),(\widehat{b},0)\rangle_{A,\ell^\infty} = \langle a,\widehat{b}\rangle_{A,\ell^\infty}.$$

□

The following is a partial converse to Corollary 2.1.3.3 (and by equivalence to Proposition 1.2.6.15):

> PROPOSITION 2.1.3.7. Let A be an abelian variety and $\phi : A \to \widehat{A}$ a homomorphism. Then the bilinear form e_{ϕ,ℓ^∞} is skew-symmetric or equivalently, ϕ is symmetric if and only if there is a line bundle $\mathscr{L}$ on A such that $2\phi = \phi_{\mathscr{L}}$

Proof. If $2\phi = \phi_{\mathscr{L}}$, then $2e_{\phi,\ell^\infty} = e_{\phi_{\mathscr{L}},\ell^\infty}$ is skew-symmetric and hence so is e_{ϕ,ℓ^∞}. Conversely, assume e_{ϕ,ℓ^∞} is skew symmetric. Let $\mathscr{L} = (1,\phi)^*\mathscr{P}_A$ where $(1,\phi)^*$ is the composition $A \xrightarrow{\Delta} A \times A \xrightarrow{1\times\phi} A \times \widehat{A}$. We claim that $2\phi = \phi_{\mathscr{L}}$. By the non-degeneracy of $\langle\cdot,\cdot\rangle_{A,\ell^\infty}$, it suffices to show that $2e_{\phi,\ell^\infty} = e_{\phi_{\mathscr{L}},\ell^\infty}$. For any $a, b \in T_\ell(A)$, we have

$$\begin{aligned} \langle a, \phi_{\mathscr{L}}(b)\rangle &= E^{\mathscr{L}}(a,b) \\ &= E^{\mathscr{P}_A}((1,\phi)(a),(1,\phi)(b) \\ &= \langle a,\phi(b)\rangle_{A,\ell^\infty} - \langle b,\phi(a)\rangle_{A,\ell^\infty} \\ &= 2\langle a,\phi(b)\rangle_{A,\ell^\infty}, \end{aligned}$$

by equations 2.1.3-I and 2.1.3-II and the skew-symmetry of $\langle a, \phi(b)\rangle_{A,\ell^\infty}$. □

§2.1.4 The Rosati Involution

Now that we have built the theory of Riemann forms, we may now introduce the Rosati involution which will be adjoint with respect to the Riemann form. We will prove that this is totally positive, which (in section 2.1.5) will allow us to completely classify up-to-isogeny endomorphism rings.

In this section, we fix an ample line bundle $\mathscr{L}$ on the abelian variety A, so that $\phi_{\mathscr{L}} : A \to \widehat{A}$ is an isogeny.

DEFINITION 2.1.4.1. The **Rosati involution** on the algebra $\mathrm{End}^0 A$ with respect to $\mathscr{L}$ is the involution

$$\phi^\dagger := \phi_{\mathscr{L}}^{-1} \circ \widehat{\phi} \circ \phi_{\mathscr{L}}$$

for all $\phi \in \mathrm{End}^0 A$.

One has the following properties of the map:

(1) For $\phi, \psi \in \mathrm{End}^0 A$,

$$(c\phi)^\dagger = c\phi^\dagger, c \in \mathbb{Q}$$
$$(\phi + \psi)^\dagger = \phi^\dagger + \psi^\dagger$$
$$(\phi\psi)^\dagger = \psi^\dagger \phi^\dagger$$

.

Proof. Firstly,

$$(c\phi)^\dagger = \phi_{\mathscr{L}}^{-1} \circ \widehat{(c\phi)} \circ \phi_{\mathscr{L}} = c\phi_{\mathscr{L}}^{-1} \circ \widehat{\phi} \circ \phi_{\mathscr{L}} = c\phi^\dagger,$$

for the next,

$$(\phi + \psi)^\dagger = \phi_{\mathscr{L}}^{-1} \circ \widehat{(\phi + \psi)} \circ \phi_{\mathscr{L}} = \phi_{\mathscr{L}}^{-1} \circ (\widehat{\phi} + \widehat{\psi}) \circ \phi_{\mathscr{L}} = \phi^\dagger + \psi^\dagger,$$

and finally,

$$(\phi\psi)^\dagger = \phi_{\mathscr{L}}^{-1} \circ \widehat{(\phi\psi)} \circ \phi_{\mathscr{L}} = \phi_{\mathscr{L}}^{-1} \circ \widehat{\psi}\widehat{\phi} \circ \phi_{\mathscr{L}} = \psi^\dagger \phi^\dagger.$$

□

(2) Extend $T_\ell : \mathrm{End}\, A \hookrightarrow \mathrm{End}_{\mathbb{Z}_\ell} T_\ell(A)$ to a homomorphism $\mathrm{End}^0 A \hookrightarrow \mathrm{End}_{\mathbb{Q}_\ell}(\mathbb{Q}_\ell \otimes_{\mathbb{Z}_\ell} T_\ell(A))$ and denote this extended map again by T_ℓ. Then for any $\phi \in \mathrm{End}^0 A$, $T_\ell(\phi^\dagger)$ is the adjoint of $T_\ell(\phi)$ for the nondegenerate bilinear form $E^{\mathscr{L}}$, that is, we have

$$E^{\mathscr{L}}(\phi(a), b) = E^{\mathscr{L}}(a, \phi^\dagger(b)). \tag{2.1.4-I}$$

In particular, $(\phi^\dagger)^\dagger = \phi$, i.e., $\phi \mapsto \phi^\dagger$ is an involution.

Proof. We have

$$E^{\mathscr{L}}(a, \phi^\dagger b) = \langle a, \phi_{\mathscr{L}} \circ \phi_{\mathscr{L}}^{-1} \circ \widehat{\phi} \circ \phi_{\mathscr{L}}(b)\rangle_{A,\ell^\infty}$$

$$= \langle a, \widehat{\phi} \circ \phi_{\mathscr{L}}(b) \rangle_{A,\ell^\infty}$$
$$= \langle \phi(a), \phi_{\mathscr{L}}(b) \rangle_{A,\ell^\infty}$$
$$= E^{\mathscr{L}}(\phi(a), b)$$

which proves the equation. □

(3) If we identify $\mathbb{Q} \otimes_{\mathbb{Z}} \mathrm{Pic}_{A/k}(k)$ with a subspace of $\mathrm{Hom}^0(A, \widehat{A})$ via the map $\mathscr{M} \mapsto \phi_{\mathscr{M}}$, then under the isomorphism $\mathrm{Hom}^0(A, \widehat{A}) \xrightarrow{\sim} \mathrm{End}^0 A$ given by $\psi : \phi_{\mathscr{L}}^{-1} \circ \psi$ (which is an isomorphism because $\phi_{\mathscr{L}}$ is an isogeny),

$$\mathbb{Q} \otimes_{\mathbb{Z}} \mathrm{Pic}_{A/k}(k) / \mathrm{Pic}^0_{A/k}(k) \xrightarrow{\sim} \{\psi \in \mathrm{End}^0 A \mid \psi^\dagger = \psi\}, \tag{2.1.4-II}$$

the set of symmetric elements of $\mathrm{End}^0 A$ with respect to the Rosati involution.

Proof. By Corollary 2.1.3.3, an element $\psi \in \mathrm{End}^0 A$ belongs to this subspace if and only if $\phi = \phi_\ell \circ \psi$, we have that $\langle a, \phi(b) \rangle_{A,\ell^\infty} = -\langle b, \phi(a) \rangle_{A,\ell^\infty}$. But this means that $E^{\mathscr{L}}(a, \psi(b)) = E^{\mathscr{L}}(b, \psi(a))$, that is

$$E^{\mathscr{L}}(a, \psi(b)) = E^{\mathscr{L}}(\psi(a), b) = E^{\mathscr{L}}(a, \psi^\dagger(b))$$

for all $a, b \in T_\ell(A)$ by the adjoint property of $\psi^\dagger$ (equation 2.1.4-I). But, $E^{\mathscr{L}}$ is non-degenerate and T_ℓ is faithful, so we conclude $\psi^\dagger = \psi$. □

THEOREM 2.1.4.2. Let A be an abelian variety. Then there is a generator

$$v \in \mathrm{Hom}_{\mathbb{Z}_\ell}\left(\bigwedge^{2g} T_\ell(A), \mathbf{M}_\ell^g\right)$$

with the following property: for all divisors $D_1, \ldots, D_g$ on A, let $\mathscr{L}_i = \mathscr{O}_A(D_i)$ and let $E_i = E^{\mathscr{L}_i}$ be the associated Riemann forms. Then

$$E_1 \wedge \cdots \wedge E_g = (D_1 \cdots D_g) \cdot v.$$

Proof. Firstly, we may assume without loss of generality that $k = \overline{k}$. This is because intersection numbers are invariant under change of base field and so is $E^{\mathscr{L}}$ since the Weil pairing is, and $\phi_{\mathscr{L}}$ acts the same as $\phi_{\mathscr{L}_{\overline{k}}}$ on $\overline{k}$-points by definition. Additionally, since both the left and right-hand sides depend polynomially on the images of $\mathrm{NS}(A)$ (since intersection numbers are locally constant by local constancy of the Hilbert polynomial), this formula results by polarization from the formula with $\mathscr{L}_1 = \cdots = \mathscr{L}_g$, $D_1 = \cdots = D_g$. Using the fact that $\chi(\mathscr{L}) = \frac{(D^g)}{g!}$ (Riemann-Roch, Thm. 1.3.3.1), we are reduced to proving

$$(E^{\mathscr{L}})^{\wedge g} = g! \chi(\mathscr{L}) \cdot v.$$

Since we are working over $k = \overline{k}$, there is a (non-canonical!) isomorphism $\mathbf{M}_\ell \simeq \mathbb{Z}_\ell$. Fix one and choose a basis for $T_\ell(A)$ over $\mathbb{Z}_\ell$. Then $\wedge^{2g} T_\ell(A)$ becomes isomorphic to $\mathbb{Z}_\ell$ by using this basis, and so $(E^{\mathscr{L}})^{\wedge g}$ becomes a scalar and as does v. Further, by using this basis, $E^{\mathscr{L}}$ has an associated matrix. We assert that

$$((E^{\mathscr{L}})^{\wedge g})^2 = (\det E^{\mathscr{L}})(g!)^2.$$

In fact, this equation remains invariant under change of basis for $\mathbb{Q}_\ell \otimes T_\ell(A)$, and it follows by a simple computation on choosing a basis so that $E^{\mathscr{L}}$ takes the standard form

$$\begin{pmatrix} \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} & & 0 \\ & \ddots & \\ 0 & & \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \end{pmatrix}.$$

Again, since both sides of the equality of the theorem are polynomials in $\mathscr{L}$, we are reduced to showing $\chi(\mathscr{L})^2 \stackrel{?}{=} c \det E^{\mathscr{L}}$, and by the Riemann-Roch theorem (1.3.3.1), this is equivalent to $\deg \phi_{\mathscr{L}} \stackrel{?}{=} c \det E^{\mathscr{L}}$, where c is an ℓ-adic unit. Now, $E^{\mathscr{L}}(a, b) = \langle a, \phi_{\mathscr{L}}(b) \rangle_{A,\ell^\infty}$, and since $\langle \cdot, \cdot \rangle_{A,\ell^\infty}$ is a non-degenerate pairing over $\mathbb{Z}_\ell$, we see that $\det E^{\mathscr{L}} = \det T_\ell(\phi_{\mathscr{L}})$ if we define the matrix representation of $T_\ell(\phi_{\mathscr{L}})$ using the dual basis of $T_\ell(\widehat{A})$.

Choose and fix an isogeny $\psi : \widehat{A} \to A$. We then have that

$$\begin{aligned} \deg \psi \cdot \deg \phi_{\mathscr{L}} &= \deg(\psi \circ \phi_{\mathscr{L}}) \\ &= \det(T_\ell(\psi) \circ T_\ell(\phi_{\mathscr{L}})) \\ &= \det(T_\ell(\psi)) \det(T_\ell(\phi_{\mathscr{L}})) \end{aligned}$$

by Theorem 2.1.2.17. Thus, to complete the proof of the theorem, we only have to show that $\frac{\deg \psi}{\det T_\ell(\psi)}$ is an ℓ-adic unit, but this follows from the fact that the largest power of ℓ dividing $\deg \psi$ is the order of the kernel, or equivalently cokernel, of $\psi|_{A[\ell^n]} : A[\ell^n] \to A[\ell^n]$ for n sufficiently large, but this is the same as the largest power of ℓ dividing $\det T_\ell(\psi)$. □

> COROLLARY 2.1.4.3. Let A be a simple abelian variety of dimension g and $K \subset \mathrm{End}^0 A$ a $\mathbb{Q}$-subalgebra such that $\phi = \phi^\dagger$ for all $\phi \in K$. Then $[K : \mathbb{Q}]$ divides g.

Proof. Since $\phi\psi = \phi^\dagger\psi^\dagger = (\psi\phi)^\dagger = \psi\phi$ for all $\phi, \psi \in K$, K is a subfield of $\mathrm{End}^0 A$. Further, since K consists of symmetric elements, it is contained in the image of $\mathbb{Q} \otimes_{\mathbb{Z}} \mathrm{Pic}_{A/k}(k)/\mathrm{Pic}^0_{A/k}(k)$ under the map $\mathscr{M} \mapsto \phi_{\mathscr{L}}^{-1}$ by property 2.1.4-II. Now, $\chi(\mathscr{M})$ depends only on the endomorphism $\phi_{\mathscr{L}}^{-1}\phi_{\mathscr{M}}$ and it extends to a homogeneous polynomial function of degree g on the space of symmetric elements of $\mathrm{End}^0 A$. We claim that the restriction of $\frac{\chi(\mathscr{M})}{\chi(\mathscr{L})}$ is a norm function on K. It is clear (since it is already polynomial) that if its square $\frac{\chi^2(\mathscr{M})}{\chi^2(\mathscr{L})}$ is multiplicative on K, then it is multiplicative as well. We have

$$\frac{\chi^2(\mathscr{M})}{\chi^2(\mathscr{L})} = \frac{\deg \phi_{\mathscr{M}}}{\deg \phi_{\mathscr{L}}} = \deg(\phi_{\mathscr{L}}^{-1} \circ \phi_{\mathscr{M}})$$

by Riemann-Roch (Thm. 1.3.3.1), which is multiplicative in $\phi_{\mathscr{L}}^{-1} \circ \phi_{\mathscr{M}}$. Thus, we get a norm function of degree g on K, and K being a field of degree $[K : \mathbb{Q}]$, the corollary follows. □

We will now prove the positivity of the Rosati involution, which will allow us to classify the endomorphism rings of abelian varieties.

THEOREM 2.1.4.4. Let D be an ample divisor on an abelian variety, $\mathscr{L} = \mathscr{O}_A(D)$ the associated line bundle, and † the Rosati involution of $\mathrm{End}^0 A$ with respect to $\mathscr{L}$. Then for any $\phi \in \mathrm{End}\, A$, we have

$$\mathrm{Tr}(\phi\phi^\dagger) = \frac{2g}{(D^g)}(D^{g-1} \cdot \phi^*(D))$$

where $(\cdot)$ denotes intersection numbers and $\mathrm{Tr} = \mathrm{Tr}_{\mathrm{End}^0 A/\mathbb{Q}}$ for ease of notation. In particular, $\phi \mapsto \mathrm{Tr}(\phi\phi^\dagger)$ is a positive definite quadratic form on $\mathrm{End}^0 A$.

Proof. The first assertion clearly implies the second, since for any effective divisor H and ample divisor D, $(D^{g-1} \cdot H) > 0$. It suffices therefore to prove the first statement. Choose and fix bases for $T_\ell(A)$ and $\mathbf{M}_\ell$. Applying Theorem 2.1.4.2, we get

$$(E^{\mathscr{L}})^{\wedge g} = c(D^g),$$
$$(E^{\mathscr{L}})^{\wedge(g-1)} \wedge E^{\phi^*\mathscr{L}} = c(D^{g-1} \cdot \phi^*(D)),$$

for some ℓ-adic constant c. Therefore,

$$\frac{(E^{\mathscr{L}})^{\wedge(g-1)} \wedge E^{\phi^*\mathscr{L}}}{(E^{\mathscr{L}})^{\wedge g}} = \frac{D^{g-1} \cdot \phi^* D}{(D^g)}.$$

Since we have $E^{\phi^*\mathscr{L}} = E^{\mathscr{L}} \circ (\phi \times \phi)$ (see equation 2.1.3-I and note the abuse of notation treating ϕ as $T_\ell(\phi)$), we are reduced to proving the equation

$$\frac{(E^{\mathscr{L}})^{\wedge(g-1)} \wedge E^{\phi^*\mathscr{L}}}{(E^{\mathscr{L}})^{\wedge g}} = \frac{1}{2g}\mathrm{Tr}(\phi\phi^\dagger).$$

where we note that $\phi^\dagger$ is the transpose of ϕ with respect to $E^{\mathscr{L}}$ by 2.1.4-I. This is now purely a problem on linear algebra. Choose a basis $\{e_1, \ldots, e_{2g}\}$ of $\mathbb{Q}_\ell \otimes T_\ell(A)$ such that $E^{\mathscr{L}}(e_{2i-1}, e_{2i}) = 1$ and $E^{\mathscr{L}}(e_{2i-1}, e_j) = E^{\mathscr{L}}(e_{2i}, e_j) = 0$ if $j \neq 2i$ or $2i-1$. Then the left side becomes (by definition of the exterior product)

$$\sum_{i_1,\ldots,i_g \text{ odd}} E^{\mathscr{L}}(\phi(e_{i_g}), \phi(e_{i_g+1})) \Big/ \sum_{i_1,\ldots,i_g \text{ odd}} 1 = \frac{(g-1)!}{g!} \sum_{i \text{ odd}} E^{\mathscr{L}}(\phi(e-i), \phi(e_{i+1}))$$
$$= \frac{1}{2g} \sum_{i \text{ odd}} \left(E^{\mathscr{L}}(e_i, \phi\phi^\dagger(e_{i+1})) + E^{\mathscr{L}}(\phi^\dagger\phi(e_i), e_{i+1})\right)$$
$$= \frac{1}{2g}\mathrm{Tr}(\phi^\dagger\phi).$$

□

§2.1.5 The Structure of $\mathrm{End}^0 A$ for Simple A

We may now begin our classification of the possible structures of $\mathrm{End}^0 A$ for simple A. We have seen that for a simple abelian variety A, $D = \mathrm{End}^0 A$ is a division algebra of finite rank over $\mathbb{Q}$ with an involution $x \mapsto x^\dagger$ such that if $x \neq 0$, $\mathrm{Tr}(xx^\dagger) > 0$ (i.e. † is a positive involution) where the $\mathrm{Tr} = \mathrm{Trd}_{\mathrm{End}^0 A/\mathbb{Q}}$ (or any positive multiple of it).

Before giving the classification of such algebras, we first define the Brauer group and give a general result about it which is important to the proof the classification (which we will not give) and will also be mentioned later in the paper.

DEFINITION 2.1.5.1. Two rings R, S are said to be **Morita equivalent** if their categories of modules are additively equivalent.

The **Brauer group** of a field K is an abelian group whose elements consist of Morita equivalence classes of central simple algebras over K, with group structure given by tensor product of algebras.

THEOREM 2.1.5.2. We have the following:

(i) The Brauer group of a p-adic field is canonically isomorphic to $\mathbb{Q}/\mathbb{Z}$. If $L \supset K$ are two p-adic fields with $[L : K] = n$, the induced map $\mathrm{Br}(K) \to \mathrm{Br}(L)$ goes over by means of the above isomorphisms into multiplication by n in $\mathbb{Q}/\mathbb{Z}$.

(ii) The Brauer group of $\mathbb{R}$ is cyclic of order two, and we identify it with the unique cyclic subgroup of order two in $\mathbb{Q}/\mathbb{Z}$.

(iii) The Brauer group of $\mathbb{C}$ is trivial.

(iv) For any central simple algebra D over an algebraic number field K, and any finite or infinite place v of K, if K_v denotes the completion of K at v, let $\mathrm{Inv}_v(K)$ denote the element of $\mathbb{Q}/\mathbb{Z}$ corresponding to the class of $D \otimes_K K_v$ in $\mathrm{Br}(K_v)$. Then we have an exact sequence

$$0 \to \mathrm{Br}(K) \to \prod_v \mathrm{Br}(K) \to \mathbb{Q}/\mathbb{Z} \to 0$$

where the second map is the sum of the elements of $\mathbb{Q}/\mathbb{Z}$.

Proof. See [Mil20, Ch. III, Thm. 2.1] for (i)-(iii), and [Hür92] or [Mil20, Ch. VIII, Thm. 4.2] for (iv). □

We shall now state the classification, due to Albert:

THEOREM 2.1.5.3. Let D be a finite-dimensional division algebra over $\mathbb{Q}$ with positive involution †. Let $K = Z(D)$ be the center, $K_0 := K^{\langle\dagger\rangle}$ the subfield of elements fixed by †, and $e_0 := [K_0 : \mathbb{Q}]$. Then K_0 is a totally real number field and one of the four following possibilities

(I) $D = K = K_0$ and † is the identity;

(II) $K = K_0$ and D is a quaternion algebra over K_0 such that

$$D \otimes_{\mathbb{Q}} \mathbb{R} \simeq \mathrm{Mat}_2(\mathbb{R})^{e_0},$$

and there exists $\mu \in D^\times$ such that $\mu^2 \in K_0^\times$ is totally negative and $x^\dagger = \mu^{-1}\overline{x}\mu$ for all $x \in D$, where $\overline{x} = \mathrm{Tr}\, x - x$ is the standard involution;

(III) $K = K_0$ is a quaternion algebra over K_0 such that

$$D \otimes_{\mathbb{Q}} \mathbb{R} \simeq \mathbb{H}^{e_0},$$

where $\mathbb{H}$ is the standard quaternion algebra and † acts by the standard involution on each factor; or

(IV) K_0 is a totally imaginary quadratic extension of K_0 and

$$D \otimes_{\mathbb{Q}} \mathbb{R} \simeq \mathrm{Mat}_d(\mathbb{C})^{e_0}$$

for some $d \geq 1$, and † extends to the conjugate transpose on each factor $\mathrm{Mat}_d(\mathbb{C})$.

Proof. See [BL04b, §§5.3-5.5] or [Mum74, §21, Application I]. □

Applying this to $D = \mathrm{End}^0 A$ where A is a simple g-dimensional abelian variety, if we let $d^2 = [D : K]$, $e = [K : \mathbb{Q}]$, $e_0 = [K_0 : \mathbb{Q}]$, and $\eta = \frac{\dim_{\mathbb{Q}} S}{\dim_{\mathbb{Q}} D}$ where $S = D^{\langle\dagger\rangle}$, then we have the following table:

Type	e	d	η	Restriction in char $k = 0$	Restriction in char $k = p > 0$
I	e_0	1	1	$e \mid g$	$e \mid g$
II	e_0	2	$\frac{3}{4}$	$2e \mid g$	$2e \mid g$
III	e_0	2	$\frac{1}{4}$	$2e \mid g$	$e \mid g$
IV	$2e_0$	d	$\frac{1}{2}$	$e_0 d^2 \mid g$	$e_0 d \mid g$

The results about e, d, η are all contained in Theorem 2.1.5.3. For the restrictions by characteristic, these follow from Corollary 2.1.2.19 in characteristic p, remark 2.1.2.20 in characteristic 0, and Corollary 2.1.4.3.

§2.1.6 The Riemann Hypothesis

Our goal of this section is to count the number of rational points on an abelian variety over finite fields. This will be important for proving Tate's Isogeny Theorem in Section 2.2, but will also be interesting in its own right as it will be the analogue for the Riemann Hypothesis in the case of abelian varieties over finite fields.

We will first prove

PROPOSITION 2.1.6.1. Let A be an abelian variety, † the Rosati involution on $\mathrm{End}^0 A$ defined by some ample line bundle, and $\alpha \in \mathrm{End}\, A$ such that $\alpha^\dagger \alpha = a \in \mathbb{Z}$. Let $\omega_1, \ldots, \omega_{2g}$ be the roots (in $\mathbb{C}$) of the characteristic polynomial $P(t)$ of α. Then the subalgebra $\mathbb{Q}[\alpha] \subset \mathrm{End}\, A$ generated by α is semisimple, and

(i) $|\omega_i|^2 = a$ for all i;

(ii) the map $\omega_i \to \frac{a}{\omega_i}$ is a permutation of the roots ω_i.

Proof. Note that (ii) is an immediate consequence of (i) since $\frac{a}{\omega_i} = \overline{\omega}_i$ and P is an integral polynomial. Now, let $Q(t)$ be the minimal polynomial over $\mathbb{Q}$ of α. We claim that P and Q have the same complex roots. In fact since P has integral coefficients and $P(\alpha) = 0$, we must have $Q \mid P$. But P is also the characteristic polynomial of $T_\ell(\alpha)$ (see def. 2.1.2.18) in the ℓ-adic matrix representation

$$T_\ell : \mathrm{End}(A) \to \mathrm{End}(T_\ell(A)).$$

If $\omega \in \overline{\mathbb{Q}}_\ell$ is a root of P, then ω is an eigenvalue of $T_\ell(\alpha)$, hence $Q(\omega)$ is an eigenvalue of $T_\ell(Q(\alpha))$. But $T_\ell(Q(\alpha)) = 0$, so $Q(\omega) = 0$, i.e. all roots of P in $\overline{\mathbb{Q}}_\ell$ are roots of Q. Therefore $P \mid Q^n$ for some n, so P and Q have the same complex roots as well.

The restriction S of the trace on $\mathrm{End}^0 A$ to $R := \mathbb{Q}[\alpha]$ is a trace form on R satisfying $S(xx^\dagger) > 0$ if $x \in R$, $x \neq 0$. Further, since α is invertible in $\mathrm{End}^0 A$ and R is finite dimensional, it follows that $\alpha^{-1} \in R$. Hence $\alpha^\dagger = a/\alpha \in R$, so R is stable under the involution. If $\mathfrak{a} \subset R$ is any ideal in R, and $\mathfrak{b}$ is the orthogonal complement in R for the quadratic form $S(xx^\dagger)$, $\mathfrak{b}$ is again an ideal and $\mathfrak{a} \cap \mathfrak{b} = (0)$, $\mathfrak{a} \oplus \mathfrak{b} = R$. Thus, R is semisimple, hence isomorphic to $K_1 \times \cdots \times K_m$ where K_i are algebraic number fields. The involution, being an automorphism of $\mathbb{Q}[\alpha]$, permutes the factors K_i. But since $S(xx^\dagger) > 0$ for all $x \neq 0$, the involution must take each K_i onto itself, and therefore, S is a trace form on each K_i over $\mathbb{Q}$ with $S(xx^\dagger) > 0$ if $x \neq 0$. Hence, each K_i is either totally real with identity involution or is a totally imaginary quadratic extension of a totally real subfield with complex conjugation as its involution. Now, the roots ω of the P are precisely the image of α under the various embeddings $\sigma_{ij} : K_i \hookrightarrow \mathbb{C}$, since $\sigma_{ij}(x^\dagger) = \overline{\sigma_{ij}(x)}$ for all $x \in \mathbb{Q}[\alpha]$, it follows that

$$a = \sigma_{ij}(a) = \sigma_{ij}(\alpha^\dagger \alpha) = |\sigma_{ij}(\alpha)|^2.$$

□

We will apply this proposition to prove the Riemann hypothesis on abelian varieties over finite fields. Let $\mathbb{F} = \mathbb{F}_q$ be a finite field with $q = p^f$ elements, and A an abelian variety over $\mathbb{F}$. We define the Frobenius morphism $\pi : A \to A$ to be the morphism defined by the identity on the underlying topological space and the map of structure sheaves $\mathscr{O}_A \to \mathscr{O}_A$ given by $f \mapsto f^q$. By Corollary 1.1.1.3, π is the translate of an endomorphism, but since π is the identity on the underlying topological space, we must have that π is in fact an *endomorphism*.

Let $P_\pi(t)$ be the characteristic polynomial of π and suppose

$$P_\pi(t) = \prod_{i=1}^{2g} (t - \omega_i).$$

Then we have the following:

THEOREM 2.1.6.2 (Weil). Let A be an abelian variety. Let $N_m = \#A(\mathbb{F}_{q^m})$. Then

$$N_m = \prod_{i=1}^{2g} (1 - \omega_i^m)$$

and the $\omega_i \in \mathbb{C}$ satisfy

(i) (The Riemann Hypothesis) $|\omega_i| = q^{1/2}$ and

(ii) $\omega_{\sigma(i)} = q/\omega_i$ for some permutation $\sigma \in S_{2g}$.

Proof. We first note that since π^m fixes 0, by Corollary 1.1.1.3, π^m is an endomorphism. Moreover, since $\mathrm{Lie}(\pi) = 0$, $1 - \pi^m$ is the identity on $\mathrm{Lie}(A)$, so π^n is furthermore a separable endomorphism. Hence we obtain:

$$N_m = \#(\ker(1 - \pi^n)) = \deg(1 - \pi^n).$$

But then the characteristic polynomial $P_m(t)$ of π^m for all m is $\prod_{i=1}^{2g}(t - \omega_i^m)$. But since $\deg(1 - \pi^n) = P_m(1)$ by Theorem 2.1.2.17, the first claim follows.

We will now show claims (i) and (ii). Let $\mathscr{L}$ be an ample line bundle on A. Let † be the Rosati

involution with respect to $\mathscr{L}$. We claim that

$$\pi^\dagger \circ \pi \stackrel{?}{=} [q]_A,$$

so that the proposition applies. But by the definition of †, this means that

$$\widehat{\pi} \circ \phi_{\mathscr{L}} \circ \pi = [q]_A \circ \phi_{\mathscr{L}}.$$

But π acts on $\mathscr{O}_A$ by $f \mapsto f^q$, so it follows that $\pi^*\mathscr{L} = \mathscr{L}^q$ and for any $a \in A(S)$

$$\begin{aligned}\pi^*(T^*_{\pi(a)}\mathscr{L}_S \otimes \mathscr{L}_S^{-1}) &\simeq T_a^*\pi^*\mathscr{L}_S \otimes (\pi^*\mathscr{L}_S)^{-1} \\ &\simeq (T_x^*\mathscr{L} \otimes \mathscr{L}^{-1})^{\otimes q}.\end{aligned}$$

Since the line bundle on the left represents $\widehat{\pi}(\phi_{\mathscr{L}}(\pi(a)))$ and the right represents $[q]_A(\phi_{\mathscr{L}}(a))$, we have that $\pi^\dagger \circ \pi = [q]_A$. Thus, by Proposition 2.1.6.1, (i) and (ii) follow. □

COROLLARY 2.1.6.3. For some constant C, $|N_m - q^{mg}| \leq Cq^{n(g-\frac{1}{2})}$ for all n.

Another application of Proposition 2.1.6.1 is

THEOREM 2.1.6.4 (Serre). For any $n \geq 3$, and any $\mathscr{L}$ ample on an abelian variety A, the restriction homomorphism

$$\left\{\alpha \in \operatorname{Aut} A \mid \alpha^*\mathscr{L} \simeq \mathscr{L} \bmod \operatorname{Pic}^0 A\right\} \to \operatorname{Aut}(A[n])$$

is injective.

Proof. If $\alpha^*\mathscr{L} \simeq \mathscr{L} \bmod \operatorname{Pic}^0 A$, then $\phi_{\alpha^*\mathscr{L}} = \phi_{\mathscr{L}}$, hence $\widehat{\alpha} \circ \phi_{\mathscr{L}} \circ \alpha = \phi_{\mathscr{L}}$. This means that for the Rosati involution defined by $\mathscr{L}$, $\alpha^\dagger\alpha = 1$. Hence by Proposition 2.1.6.1. Hence by the proposition the roots of the characteristic polynomial of α are all algebraic integers all of whose conjugates have absolute value 1, and hence all are roots of unity.

Suppose now that α restricts to the identity on some $A[n]$ $(n \geq 3)$. Then the restriction of $\alpha - 1$ to $A[n]$ is 0, so that $(\alpha - 1) = n\beta$ for some $\beta \in \operatorname{End} A$. We deduce that if ω is any characteristic root of α, $\omega - 1 = n\eta$ where η is an algebraic integer. We how have

LEMMA 2.1.6.5. If ω is a root of unity such that $\omega = 1 + n\eta$ where $n \in \mathbb{Z}$, $n \geq 3$, and η an algebraic integer, then $\omega = 1$.

Proof. If not, by raising ω to a suitable power, we may assume that ω is a primitive pth root of unity for a prime p. Taking norms over $\mathbb{Q}$ in the equation $\omega - 1 = n\eta$, we obtain

$$\prod_{i=1}^{p-1}(1 - \omega^i) = n^{p-1}N$$

where $N = (-1)^{p-1}\operatorname{Nm}(\eta)$ is a integer. But the left side is the derivative of $t^p - 1 = \prod_{i=0}^{p-1}(t - \omega^i)$ at $t = 1$, that is p. Hence n^{p-1} divides p, which is impossible if $n \geq 3$. □

Applying this, we get that the characteristic roots of α are all 1, so $1 - \alpha$ is nilpotent, but by Proposition 2.1.6.1, $\mathbb{Q}[\alpha]$ is semisimple, so it has no non-zero nilpotent elements. Thus, $\alpha = 1$. □

§2.1.7 The Structure of $\mathrm{NS}^0(A)$

Another interesting application of what we have done so far is that it allows us to add structure to the Néron-Severi group and classify the "up-to-isogeny" version of it. Let A be an abelian variety and let $\mathrm{NS}^0(A) = \mathrm{NS}(A) \otimes \mathbb{Q}$. Fix ample $\mathscr{L}$ on A and let † be its associated Rosati involution, then we can identify

$$\rho : \mathrm{NS}^0(A) \xrightarrow{\sim} \{\alpha \in \mathrm{End}^0 A | \alpha^\dagger = \alpha\}$$

(see 2.1.4-II). In particular, $\mathrm{NS}^0(A)$ has a natural structure of a Jordan algebra over $\mathbb{Q}$ if we define

$$\alpha \circ \beta = \frac{1}{2}\rho^{-1}(\rho(\alpha)\rho(\beta) + \rho(\beta)\rho(\alpha))$$

for $\alpha, \beta \in \mathrm{NS}^0(A)$, using composition in $\mathrm{End}^0 A$. What can we say about this Jordan algebra? Firstly, since $\mathrm{Tr}(\rho(\alpha)^2) \geq 0$ for all $\alpha \in \mathrm{NS}^0(A)$, $\alpha \neq 0$, $\mathrm{NS}^0(A)$ is formally real, i.e.,

$$\sum_{i=1}^{n} \alpha_i \circ \alpha_i = 0 \implies \alpha_1 = \cdots = \alpha_n = 0.$$

The formally real Jordan algebras over $\mathbb{R}$ has been classified. (See [Mcc03, §1] for a general discussion of Jordan algebras, and §§1.11 for this specific classification). In our case, $\mathrm{NS}^0(A) \otimes \mathbb{R}$ does not go through all possible types of Jordan algebras. In fact, we have

THEOREM 2.1.7.1. $\mathrm{NS}^0(A) \otimes \mathbb{R}$ is isomorphic to a product of Jordan algebras of the types:

(I) $\mathscr{H}_r(\mathbb{R}) = r \times r$ symmetric real matrices;

(II) $\mathscr{H}_r(\mathbb{C}) = r \times r$ Hermitian complex matrices;

(III) $\mathscr{H}_r(\mathbb{H}) = r \times r$ Hermitian quaternionic matrices, i.e., $\overline{X}^\top = X$ where $x \to \overline{x}$ is the standard involution on $\mathbb{H}$.

Proof. Decompose $\mathrm{End}^0(A) \otimes \mathbb{R}$ into a product of copies of $\mathrm{Mat}_n(\mathbb{R})$, $\mathrm{Mat}_n(\mathbb{C})$, and $\mathrm{Mat}_n(\mathbb{H})$ using Corollary 2.1.2.4 and Theorem 2.1.5.3. Then $\mathrm{NS}^0(A) \otimes \mathbb{R}$ is isomorphic to the set of fixed points under a positive involution. But it is easy to check that every such involution (a) fixes each of the factors $\mathrm{Mat}_n(\mathbb{R})$, $\mathrm{Mat}_n(\mathbb{C})$, and $\mathrm{Mat}_n(\mathbb{H})$, and (b) can be put in the standard form $X \mapsto \overline{X}^\top$ by inner automorphisms. □

Fix isomorphisms ϕ and ψ:

$$\begin{array}{ccc} \mathrm{End}^0(A) \otimes \mathbb{R} & \xrightarrow[\phi]{\sim} & \prod \mathrm{Mat}_{r_i}(\mathbb{R}) \times \mathrm{Mat}_{s_i}(\mathbb{C}) \times \mathrm{Mat}_{t_i}(\mathbb{H}) \\ \uparrow \rho & & \uparrow \\ \mathrm{NS}^0(X) \otimes \mathbb{R} & \xrightarrow[\psi]{\sim} & \prod \mathscr{H}_{r_i}(\mathbb{R}) \times \mathscr{H}_{s_i}(\mathbb{C}) \times \mathscr{H}_{t_i}(\mathbb{H}) \end{array} \tag{2.1.7-I}$$

What happens to the polynomial Euler characteristic function $\chi : \mathrm{NS}^0(A) \otimes \mathbb{R} \to \mathbb{R}$? For any $x \in \mathrm{End}^0(A) \otimes \mathbb{R}$, let $\phi_{i,1}(x), \phi_{i,2}(x), \phi_{i,3}(x)$ denote the components of $\phi(x)$ in the above decomposition. Then we know that the degree function can be written

$$\deg(x) = \prod \det(\phi_{i,1}(x))^{a_i} \cdot \prod |\det(\phi_{i,2}(x))|^{2b_i} \cdot \prod \mathrm{Nrd}_{\mathrm{Mat}_t(\mathbb{H})/\mathbb{R}}(\phi_{i,3}(x))^{c_i} \tag{2.1.7-II}$$

Now NS(A), $\deg(\rho x) = \alpha\chi(x)^2$ for some constant α. It follows that the a_i are even and that the function χ can be written

$$\chi(x) = C \cdot \prod \det(\psi_{i,1}(x))^{a_i/2} \cdot \prod \det(\psi_{i,2}(x))^{b_i} \cdot \prod \mathrm{HNm}(\psi_{i,3}(x))^{c_i} \tag{2.1.7-III}$$

where C is a constant for all $x \in \mathrm{NS}^0(A) \otimes \mathbb{R}$. Note that $\psi_{i,2}(x)$ is Hermitian, so its complex determinant is real. Here HNm is the "Haupt norm," a polynomial function of degree t from $\mathscr{H}_t(\mathbb{H}) \to \mathbb{R}$. If $\lambda_0 \in \mathrm{NS}^0(A) \otimes \mathbb{R}$ is the point defined by the ample $\mathscr{L}_{\overline{k}}$ on $A_{\overline{k}}$ used to set up ρ, then $\rho(\lambda_0) = 1$, so $\psi_{i,j}(\lambda_0) = I$, and the constant in the above formula is $\chi(\lambda_0)$. Using Theorem 1.3.3.3, it follows that if $\chi(x) \neq 0$ then the index of the associated line bundle $\mathscr{L}_x$

$$i(x) = \sum \frac{a_i}{2} \left(\begin{smallmatrix}\#\text{neg. eigenvalues} \\ \text{of } \psi_{i,1}(x)\end{smallmatrix}\right) + \sum b_i \left(\begin{smallmatrix}\#\text{neg. eigenvalues} \\ \text{of } \psi_{i,2}(x)\end{smallmatrix}\right) + \sum c_i \left(\begin{smallmatrix}\#\text{neg. eigenvalues} \\ \text{of } \psi_{i,3}(x)\end{smallmatrix}\right). \tag{2.1.7-IV}$$

(note that the eigenvalues of a quaternionic Hermitian matrix H are defined as the entries of a diagonal matrix D such that $H = UDU^{-1}$, and $\overline{U}^{\top} = U^{-1}$). Since ample line bundles $\mathscr{L}$ are characterized by $\chi(\mathscr{L}) = 0$, $i(\mathscr{L}) = 0$, it follows from equations 2.1.7-III and 2.1.7-IV that the images of the ample line bundles in NS(A) are exactly the totally positive elements of the formally real Jordan algebra $\mathrm{NS}^0(A) \otimes \mathbb{R}$.

§2.1.8 CM Abelian Varieties of characteristic 0

§2.1.9 Revisiting Elliptic Curves in Characteristic $p > 0$

This needs to be written better – This should just be revisiting elliptic curves In this section, we will explore a very important example, the one dimensional case of elliptic curves, in characteristic $p > 0$. All the facts in this section will be an immediate consequence of the general theory which we have proved so far in this chapter.

Let A be an abelian variety of dimension one, i.e., an elliptic curve. We, will denote the divisor corresponding to a closed point $P \in A$ by $[P]$. Then for any divisor on A, we have $\chi(\mathscr{O}_A(D)) = \deg D$, and if further $\deg D > 0$, then $\chi(\mathscr{O}_A(D)) = \dim H^0(\mathscr{O}_A(D))$ and $H^1(\mathscr{O}_A(D)) = 0$. A divisor D belongs to $\mathrm{Pic}^0 A$ if and only if $\deg D = 0$. A divisor $D = \sum n_i[P_i]$ of degree 0 is linearly equivalent to zero if and only if $\sum n_i P_i = 0$ on A. Any divisor D of degree ≥ 3 is very ample (see 1.3.3.10).

Suppose now that the characteristic not 2 (allowing characteristic 0). Let 0 be the identity element of the group A and let P_1, P_2, P_3 be the (geometric) points of order two on A. Since $\dim H^0(A, \mathscr{O}_A(2[0])) = 2$, we can choose a non-constant function $x \in H^0(A, \mathscr{O}_A(2[0]))$ having a double pole at 0 and regular elsewhere. Subtracting a constant from x, we may assume $x(P_1) = 0$, and since the sum of the zeroes with multiplicity is 0 and there are exactly two zeros, we deduce that P_1 is a double zero of x and there are no other zeros. Thus, dividing by a constant, we may assume that $x(P_2) = 1$ and $x(P_3) = \lambda \in k^{\times}$. By applying the above argument to $x - 1$, we deduce that $\lambda \neq 0, 1$.

Since $H^0(\mathscr{O}_A(3[0]))$ is of dimension 3, we can find a function y having a triple pole at 0 and regular everywhere else. By subtracting a suitable linear combination $ax + b$ from y, we may assume that $y(P_1) = y(P_2) = 0$, and since the number of zeroes is 3 (with multiplicity) and the sum of zeroes is 0, we deduce that y has simple zeros at P_1, P_2, and $P_3 = -P_1 - P_2$. Both the functions y^2 and $x(x-1)(x-\lambda)$ have poles of order 6 and 0 and double zeroes at P_1, P_2, P_3 and no other zeroes or poles anywhere, so that they

differ by a nonzero scalar factor. Replacing y by a non-zero scalar multiple, we arrive at an equation

$$A_\lambda : y^2 = x(x-1)(x-\lambda) \tag{2.1.9-I}$$

for $A - \{0\}$ in $\mathbb{A}^2_k$. Conversely, the projective curve $y^2t = x(x-t)(x-\lambda t)$ has no singularities and is of genus 1 for $\lambda \neq 0, 1$, and hence defines an elliptic curve A_λ in $\mathbb{P}^2_k$.

We wish to find all possible values of λ for which A_λ is of p-rank 0, for characteristics $p > 2$. We know that the p-rank is 0 if and only if the Frobenius map in $H^1(A, \mathscr{O}_A)$ is trivial (see Thm. 1.3.2.13). The meromorphic form dx on A is regular $A - \{0\}$ and vanishes at $P_1 = (0,0)$, $P_2 = (1,0)$, and $P_3 = (\lambda, 0)$ to the first order, and nowhere else, since $dx(P) = 0$ and $x(P) = \alpha$ implies that $x - \alpha$ vanishes to the second order at P, hence $2P$ must be 0. Thus, the form $\omega = dx/y$ is regular and nowhere vanishing in $A - \{0\}$. It follows that ω must be regular and non-vanishing at 0 also. If we let $U_0 = A - \{0\}$, $U_1 = A - \{P_1\}$, then $\mathcal{U} = (U_0, U_1)$ is an affine covering of A. A 1-cocycle for this covering is a regular function f in $U_0 \cap U_1 = A - \{0\} - \{P_1\}$, and this is a coboundary if and only if $f = g - h$ with g regular on U_0 and h regular on U_1.

Consider the linear form on $C^1(\mathcal{U}, \mathscr{O}_A) = \Gamma(U_0 \cap U_1, \mathscr{O}_{U_0 \cap U_1})$ defined by $f \mapsto \mathrm{Res}_{P_1}(f\omega)$. Since the residue at any point of a meromorphic form is zero by the residue theorem, we see that $\mathrm{Res}_{P_1}(f\omega) = 0$ if f is a coboundary. On the other hand, the function y/x is regular on $U_0 \cap U_1$ has a simple pole at P_1, so that $\mathrm{Res}_{P_1}(y/x, \omega) \neq 0$. Since $\dim H^1(A, \mathscr{O}_A) = 1$, we deduce that the above linear form induces an isomorphism $H^1(A, \mathscr{O}_A) \xrightarrow{\sim} k$, and also that $y/x \in \Gamma(U_0 \cap U_1, \mathscr{O}_{U_0 \cap U_1})$ defines a nonzero cohomology class in $H^1(A, \mathscr{O}_A)$. Hence, the Frobenius map on $H^1(A, \mathscr{O}_A)$ is trivial if and only if $y^p/x^p \in \Gamma(U_0 \cap U_1, \mathscr{O}_{U_0 \cap U_1})$ is a coboundary, hence if and only if $\mathrm{Res}_{P_1}\left(\frac{y^p}{x^p}\frac{dx}{y}\right) = 0$. Now, let c_1 coefficient of x^{p-1} in $(y^2)^{\frac{p-1}{2}} = x^{\frac{p-1}{2}}(x-1)^{\frac{p-1}{2}}(x-\lambda)^{\frac{p-1}{2}}$ and c_2 be the coefficient of $x^{\frac{p-1}{2}}$ in $(x-1)^{\frac{p-1}{2}}(x-\lambda)^{\frac{p-1}{2}}$. We have

$$\begin{aligned}
\mathrm{Res}_{P_1}\left(\frac{y^p}{x^p}\frac{dx}{y}\right) &= \mathrm{Res}_{P_1}\left(\frac{(y^2)^{\frac{p-1}{2}}}{x^{p-1}}\frac{dx}{x}\right) \\
&= 2c_1 = 2c_2 \\
&= \pm 2\left\{\sum_{\nu=0}^{\frac{(p-1)}{2}} \binom{\frac{p-1}{2}}{\nu} \cdot \lambda_\nu\right\} \\
&= \pm 2\Phi(\lambda),
\end{aligned}$$

where

$$\Phi(\lambda) = \sum_{\nu=0}^{\frac{p-1}{2}} \binom{\frac{p-1}{2}}{\nu}^2 \lambda_\nu.$$

Now, $\Phi(0)$ is the coefficient of $x^{\frac{p-1}{2}}$ in $x^{\frac{p-1}{2}}(x-1)^{\frac{p-1}{2}}$, and $\Phi(1)$ is the coefficient of $x^{\frac{p-1}{2}}$ in $(x-1)^{p-1} = \frac{x^p-1}{x-1} = 1 + x + \cdots + x^{p-1}$, so that $\Phi(0) \neq 0$, $\Phi(1) \neq 0$. Thus, every root of Φ defines an elliptic cure A_λ of p-rank 0, and these are the only elliptic curves of p-rank 0 up to isomorphism.

We call an elliptic curve in characteristic $p > 0$ **supersingular** if its p-rank is 0. We have then shown that in characteristic $p > 2$, any supersingular curve is isomorphic to one of the A_λ where λ is a root of $\Phi(t)$. If $p = 2$, it is not hard to see that there is exactly one supersingular elliptic curve, namely $y^2 + y = x^3$ (which we will not prove here). Prove this in final draft Therefore, in any positive characteristic, there

are only finitely many supersingular curves up to isomorphism.

We now study the algebra $\mathrm{End}^0 A$ of an elliptic curve over a field of characteristic $p > 0$.

THEOREM 2.1.9.1 (Deuring). Let A/k be an elliptic curve in characteristic $p > 0$. We have the following equivalences:

(1) A does not descend to an abelian variety over some finite field $k_0 \iff \mathrm{End}^0 A = \mathbb{Q}$.

(2) Suppose A descends to an elliptic curve over a finite field k_0. Then,

(i) $\mathrm{End}^0 A$ is imaginary quadratic over $\mathbb{Q} \iff$ the p-rank of A is $1 \iff \pi^n \neq [p]_A^m$ for any integers n, m, where π is the Frobenius morphism over k_0.

(ii) If the p-rank is 0, then $\mathrm{End}^0 A$ is the (unique up to isomorphism) quaternion division algebra $\mathbb{H}_{(p)}$ over $\mathbb{Q}$ which satisfies $\mathrm{Inv}_\ell \mathbb{H}_{(p)} = 0$ if $\ell \neq p$ is finte and $\mathrm{Inv}_p \mathbb{H}_{(p)} = \mathrm{Inv}_\infty \mathbb{H}_{(p)} = \frac{1}{2}$.

Finally, there are at most finitely many A (up to isomorphism) in each characteristic for which (ii) holds.

REMARK 2.1.9.2. Note that the condition of descending to an elliptic curve over a finite field k_0 is equivalent to descending to any scheme over that field. An argument for this fact is given in self-contained Lemma 2.2.1.4 (when it is more relevant).

Proof. To begin, we will show that the following three statements are equivalent:

(A) A is of p-rank 0.

(B) $\mathrm{End}^0 A$ is non-commutative.

(C) A is defined over a finite field, and if π is the Frobenius morphism over this field, $\pi^n = [p]_A^m$ for some $m, n > 0$.

We will start with (A) $\iff$ (B) and then prove (A) $\iff$ (C).

Suppose that (B) holds. Looking at the table after Theorem 2.1.5.3, we see that $\mathrm{End}^0 A$ is a central simple quaternion algebra over $\mathbb{Q}$, hence $\mathbb{Q}_p \otimes_{\mathbb{Q}} \mathrm{End}^0 A$ is also a central algebra over $\mathbb{Q}_p$. If the p-rank of A were one, $\mathbb{Q}_p \otimes_{\mathbb{Z}_p} T_p(A)$ would be a one-dimensional $\mathbb{Q}_p$-vector space in which $\mathbb{Q}_p \otimes_{\mathbb{Q}} \mathrm{End}^0 A$ admits a representation, which is impossible. Hence A has p-rank 0, proving (A).

Next, suppose (A) holds and suppose $\mathrm{End}\, A$ is commutative, so that $\mathrm{End}^0 A = K$ is an algebraic number field. Since every elliptic curve isogenous to A is again of p-rank 0, and there are only finitely many isomorphism classes of curves of p-rank 0, if R is the ring of integers of K, we can find an integer $N > 0$ such that for every A' isogenous to A, we have $NR \subset \mathrm{End}\, A'$. Choose a prime ℓ not dividing pN such that ℓR is a prime ideal in R (which exists by class field theory). Let α be a non-zero element of $T_\ell(A)$ not divisible by ℓ and let C_n be the Zariski closure of the cyclic subgroup generated by the image of α under the natural homomorphism $T_\ell(A) \to A[\ell^n]$. Then $C_n \subsetneq C_{n+1}$ and by the finiteness of the number of isomorphism classes of curves of p-rank 0, we can find integers $m > n$ such that $C_n \subsetneq C_m$ and there is an isomorphism

$$\xi : A/C_m \xrightarrow{\sim} A/C_n.$$

Thus, if $A/C_n \to A/C_m$ is the natural homomorphism induced by the inclusion $C_n \subset C_m$, we get an endomorphism $\alpha = \eta \circ \xi \in \operatorname{End} A'$ where $A' = A/C_m$, such that α has cyclic kernel of order ℓ^j, $j > 0$. Since $\deg \alpha$ and hence $\operatorname{Nm}_{K/\mathbb{Q}} \alpha$ is a power of ℓ and ℓR is a prime ideal in R, we must have $\alpha = \ell^r u$ where u is a unit in R and $r > 0$. So, $N\alpha = \ell^r N u$, and $Nu \in \operatorname{End} A'$. Now, the degree of Nu is N^2 since u is a unit in R, so that the ℓ-primary part of $\ker(\ell^r Nu)$ is $(\mathbb{Z}/\ell^r\mathbb{Z})^2$. On the other hand, the ℓ-primary part of $\ker(N\alpha)$ is isomorphic to $\ker \alpha$, which is cyclic. This is a contradiction, so we have (B).

We will now prove (A) $\iff$ (C). (A) implies that A is defined over a finite field because taking a large enough finite field $\mathbb{F}_q$, we can always find a root of $\Phi(\lambda)$. Let π be the Frobenius morphism over this field. Since $\operatorname{End} A$ is finitely generated, we can find a finite extension of degree n, say, such that every element of $\operatorname{End} A$ is defined over this extension. Hence, π^n commutes with $\operatorname{End} A$, so π^n belongs to the center of $\operatorname{End} A$. Since (A) holds, so does (B), so that $\operatorname{End}^0 A$ is a quaternion algebra with center $\mathbb{Q}$. Hence π^n is an integer, and by consideration of degree, $\pi^n = \pm p^m$ for some $m > 0$, so $\pi^{2n} = p^{2m}$. Conversely, if $\pi^n = [p]_A^m$ for some $n, m > 0$, then since π is bijective, $[p]_A$ is also bijective and A has p-rank 0.

We have thus shown that (A), (B), and (C) are equivalent. Next, if $\operatorname{End}^0 A$ is non-commutative, it is a quaternion division algebra over $\mathbb{Q}$, and since for $\ell \neq p$, $\mathbb{Q}_\ell \otimes_{\mathbb{Q}} \operatorname{End}^0 A \to \operatorname{End}_{\mathbb{Q}_\ell}(\mathbb{Q}_\ell \otimes_{\mathbb{Z}_\ell} T_\ell(A))$ is injective (Thm 2.1.2.9) and both sides have dimension 4, it is an isomorphism. Therefore $\operatorname{Inv}_\ell(\operatorname{End}^0 A) = 0$ if ℓ is finite and $\ell \neq p$. Since $\sum_v \operatorname{Inv}_v(\operatorname{End}^0 A) = 0$, the sum being over the finite and infinite places of $\mathbb{Q}$, $\operatorname{Inv}_p(\operatorname{End}^0 A) = \operatorname{Inv}_\infty(\operatorname{End}^0 A) = \frac{1}{2}$. This establishes (b)(ii).

Next, we show that if A is defined over a finite field, $\operatorname{End}^0 A \neq \mathbb{Q}$. We have proved this for A of p-rank 0. Suppose then that the p-rank of A is 1. As before, π is an endomorphism of A which is bijective and of degree a power of p, so that it cannot equal $[m]_A$ for any integer m. Thus, $\pi \in \operatorname{End} A$, $\pi \notin \mathbb{Z}$. It follows from the table of possibilities following Theorem 2.1.5.3 that $\operatorname{End}^0 A$ is an imaginary quadratic extension of $\mathbb{Q}$. This proves (b) (i).

We have now established that if $\operatorname{End}^0 A = \mathbb{Q}$, A does not descend to a finite field. Suppose finally that A is not defined over a finite field. We may assume that A is in the normal form (2.1.9-I) in characteristic $p > 2$, with λ transcendental. (If $p = 2$, the argument still works, only with a different normal form not defined in this paper). Define in final draft Call this curve A_λ. Since any two transcendental elements λ and λ' over the prime field are conjugate over the prime field, we see that if $\operatorname{End} A_\lambda = D$, then $\operatorname{End}_\mu \simeq D$ for any other transcendental μ over the prime field. Now, $\operatorname{End}^0 A$ must be either $\mathbb{Q}$ or an imaginary quadratic extension of $\mathbb{Q}$, since the only other possibility is that of a non-commutative division algebra, in which case, by the implication (B) $\implies$ (A) above, A_λ has p-rank 0 and λ must be algebraic.

Suppose then that $\operatorname{End}^0 A$ is an imaginary quadratic extension of $\mathbb{Q}$. Then we can find an element $\alpha \in D$ such that $\alpha^2 = N \in \mathbb{Z}$, $N < 0$. Hence for any transcendental μ over the prime field, there is an $\alpha_\mu \in \operatorname{End} A_\mu$ such that $\alpha_\mu^2 = N$. Suppose ℓ is a prime not dividing pN and $\xi \in T_\ell(A_\lambda)$. Let C_n be the Zariski closure of the cyclic group generated by the image of ξ under the map $T_\ell(A_\lambda) \to A_\lambda[\ell^n]$ and let $q : A_\lambda \to A_\lambda/C_n$ be the natural map. A_λ/C_n also cannot be descend to a finite field since if it did, then since q is an isogeny, there is some $[n]$ such that it factors $[n] : A_\lambda \xrightarrow{q'} A_\lambda/C_n \xrightarrow{q} A_\lambda$ which expresses A_λ/C_n as a quotient of A_λ by the finite subgroup $\ker q'$ and since finite subgroups descend, the quotients do too. Thus, we would have a contradiction. Hence, A_λ/C_n is of the form A_μ for some μ transcendental over the prime field.

With α_μ as above, since the map

$$\mathrm{End}^0 A_\mu \to \mathrm{End}^0 A_\lambda$$

given by $\alpha \mapsto q^{-1} \circ \alpha \circ q$ is an isomorphism, we deduce that $(q^{-1} \circ \alpha_\mu \circ q)^2 = [N]_A$, and since $\alpha_\lambda^2 = [N]_A$ and $\mathrm{End}^0 A_\lambda$ is a commutative field, $q^{-1} \circ \alpha_\mu \circ q = \pm\alpha_\lambda$ and $\alpha_\mu \circ q = \pm q \circ \alpha_\lambda$. Thus, $\alpha_\lambda(\ker q) \subset \ker q$, that is $\alpha_\lambda(C_n) \subset C_n$. Since this holds for every n, we deduce that $\alpha_\lambda(\xi) = a\xi$ for some $a \in \mathbb{Z}_\ell$. Thus, for α_λ acting as an endomorphism of $T_\ell(A_\lambda)$, every vector is an eigenvector, so that α_λ acts as a scalar on $T_\ell(A_\lambda)$. Since the characteristic polynomial of α_λ has integer coefficients, this scalar must be rational, and its square cannot be negative. This contradiction shows that if A is not defined over a finite field, $\mathrm{End}^0 A = \mathbb{Q}$, thus proving (a). □

We will later give a more complete description (Theorem 2.2.2.2) which we state part of here:

THEOREM 2.1.9.3. If A is an abelian variety over a finite field k, π is the Frobenius morphism over k, then

$$\mathrm{End}_k(A) \otimes \mathbb{Q}_\ell = Z_{\mathrm{Hom}_{T_\ell}(T_\ell(A), T_\ell(A))}(T_\ell(\pi)),$$

the centralizer of the image of π under T_ℓ.

Section 2.2

Tate's Isogeny Theorem

We have seen that for A, B abelian varieties over k and a prime $\ell \neq \operatorname{char} k$, the natural homomorphism

$$T_\ell : \mathbb{Z}_\ell \otimes_Z \operatorname{Hom}(A, B) \to \operatorname{Hom}_{\mathbb{Z}_\ell}(T_\ell(A), T_\ell(B))$$

is injective (see Thm. 2.1.2.9). Note also that $G_k = \operatorname{Gal}(\overline{k}/k)$ acts on $A_{\overline{k}}$ and thus $T_\ell(A)$ and similarly for B. Thus, we have an action of G_k on $\operatorname{Hom}_{\mathbb{Z}_\ell}(T_\ell(A), T_\ell(B))$ via $(g\phi)(a) = g(\phi(g^{-1}a))$. We note that if $\phi \in \operatorname{Hom}(A, B)$ then ϕ is G_k-equivariant since A and B are defined over k. Thus, we have a map

$$T_\ell : \mathbb{Z}_\ell \otimes_{\mathbb{Z}} \operatorname{Hom}_k(A, B) \to \operatorname{Hom}_{\mathbb{Z}_\ell[G_k]}(T_\ell A, T_\ell B).$$

The goal of this section is to prove

> Theorem 2.2.0.1 (Tate's Conjecture for Finite Fields). For A, B abelian varieties over $k = \mathbb{F}_q$ finite, and $G_k = \operatorname{Gal}(\overline{k}/k)$, then the natural map
>
> $$T_\ell : \mathbb{Z}_\ell \otimes_{\mathbb{Z}} \operatorname{Hom}_k(A, B) \to \operatorname{Hom}_{\mathbb{Z}_\ell[G_k]}(T_\ell A, T_\ell B) \tag{2.2.0-I}$$
>
> is an isomorphism.

Remark 2.2.0.2. The Tate conjecture for abelian varieties is that 2.2.0-I is an isomorphism over *all* fields k of finite type over the prime field. This was proved by Tate for finite fields, Zarhin for general characteristic p and Faltings in characteristic 0. Note also that these results also hold for the $\ell = p$ case using p-divisible groups (as mentioned in remark 2.1.2.10 for injectivity). Most of the the proof of Tate's part of this theorem will actually be generalizable. There is a crucial part which requires the finiteness of k and we will note exactly where that occurs (Hypothesis (k, A, ℓ)).

Then as a consequence, we will show

> Theorem 2.2.0.3 (Tate's Isogeny Theorem for Finite Fields). Abelian varieties A and B (over $k = \mathbb{F}_q$) are isogenous if and only if their ℓ-adic Tate modules are isomorphic as Galois representations for some $\ell \neq \operatorname{char} k$.

§2.2.1 The Proof of Tate's Conjecture for Finite Fields

This entire section will be devoted to the proof of Theorem 2.2.0.1, which we will break down into steps.

Step I. It suffices to show that the homomorphism

$$\mathbb{Q}_\ell \otimes_{\mathbb{Z}} \operatorname{Hom}(A, B) = \mathbb{Q}_\ell \otimes_{\mathbb{Q}} \operatorname{Hom}^0(A, B) \to \operatorname{Hom}_{\mathbb{Q}_\ell[G_k]}(V_\ell(A), V_\ell(B)) \tag{2.2.1-I}$$

where $V_\ell(A) = \mathbb{Q}_\ell \otimes_{Z_\ell} T_\ell(A)$ is an isomorphism.

Proof. In fact, if this were true, the image of the homomorphism 2.2.0-I would be a $\mathbb{Z}_\ell$-submodule of maximal rank in $\mathrm{Hom}_{\mathbb{Z}_\ell[G_k]}(T_\ell(A), T_\ell(B))$. On the other hand, if we let $M = \mathrm{Im}(T_\ell)$, then we claim that $\mathrm{Hom}_{\mathbb{Z}_\ell}(T_\ell(A), T_\ell(B))/M$ has no torsion. Suppose $\phi \in \mathrm{Hom}_{\mathbb{Z}_\ell}(T_\ell(A), T_\ell(B))$ and $\ell\phi \in M$, i.e., ϕ is torsion mod M. Then we can find $\psi_n \in \mathrm{Hom}(A, B)$ such that $T_\ell(\phi_n) \to \ell\phi$, so that for all sufficiently large n, $T_\ell(\psi_n) = \ell\phi_n$, $\phi_n \in \mathrm{Hom}_{\mathbb{Z}_\ell}(T_\ell(A), T_\ell(B))$. Thus, $\mathrm{Im}(T_\ell(\psi_n)) \subset \ell T_\ell(B)$ and ψ_n vanishes on $A[\ell]$. Hence, ψ_n admits a factorization

$$\psi_n = \psi_n' \circ [\ell]_A = \ell\psi_n'$$

and there exists some $\psi' \in \mathbb{Z}_\ell \otimes_{\mathbb{Z}} \mathrm{Hom}(A, B)$ such that $\psi_n' \to \psi'$ and $T_\ell(\psi') = \phi$, so $\phi \in M$. This proves our claim, and thus establishes that

$$M = \mathrm{Hom}_{\mathbb{Z}_\ell[G_k]}(T_\ell(A), T_\ell(B)).$$

□

Step II. It suffices to show that for any abelian variety A over $k = \mathbb{F}_q$,

$$\lambda : \mathbb{Q}_\ell \otimes_{\mathbb{Q}} \mathrm{End}^0 A \to \mathrm{End}_{\mathbb{Q}_\ell[G_k]}(V_\ell(A)) \tag{2.2.1-II}$$

is an isomorphism.

Proof. In fact if 2.2.1-II is an isomorphism with $A \times B$ instead of A, we have the direct sum decomposition as G_k-modules

$$\begin{array}{ccccccccc} \mathbb{Q}_\ell \otimes \mathrm{End}^0(A \times B) & \tilde{\leftarrow} & (\mathbb{Q}_\ell \otimes \mathrm{End}^0 A) & \bigoplus & (\mathbb{Q}_\ell \otimes \mathrm{End}^0 B) & \bigoplus & (\mathbb{Q}_\ell \otimes \mathrm{Hom}^0(A, B)) & \bigoplus & (\mathbb{Q}_\ell \otimes \mathrm{Hom}^0(B, A)) \\ \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\ \mathrm{End}(V_\ell(A \times B)) & \xleftarrow{\sim} & \mathrm{End}(V_\ell(A)) & \bigoplus & \mathrm{End}(V_\ell(A)) & \bigoplus & \mathrm{Hom}(V_\ell(A), V_\ell(B)) & \bigoplus & \mathrm{Hom}(V_\ell(B), V_\ell(A)) \end{array}$$

the above diagram is commutative and the first vertical arrow is an isomorphism, so equation 2.2.1-I is an isomorphism. □

Thus, henceforth we will restrict ourselves to a single abelian variety A, and prove that the map $2.2.1 - II$ is an isomorphism.

Step III. It suffices to show that λ is an isomorphism for one ℓ and that $\dim_{\mathbb{Q}_\ell} \mathrm{End}_{\mathbb{Q}_\ell[G_k]}(V_\ell(A))$ is independent of $\ell \neq \mathrm{char}\, k$.

Proof. This follows from the fact that $\dim_{\mathbb{Q}_\ell}(\mathbb{Q}_\ell \otimes_{\mathbb{Q}} \mathrm{End}^0 A)$ is independent of ℓ and λ is always injective (Thm 2.1.2.9). □

Step IV To establish the isomorphism 2.2.1-II for an ℓ, it suffices to show the following: Let $E = \mathrm{Im}(\lambda)$, and

$$F = \langle i(G_k) \rangle \subset \mathrm{End}_{\mathbb{Q}_\ell}(V_\ell(A))$$

where $i : G_k \hookrightarrow \mathrm{End}_{\mathbb{Q}_\ell}(V_\ell(A))$ is the map sending an element g to the map induced by g as an endomorphism of A. Then

$$F = E' := \{\phi \in \mathrm{End}_{\mathbb{Q}_\ell}(V_\ell(A)) \mid \phi\psi = \psi\phi \text{ for all } \psi \in E\},$$

the commutant of E in $\mathrm{End}_{\mathbb{Q}_\ell}(V_\ell(A))$.

Proof. Since E is semi-simple, if the above were true, it would follow from von Neumann's double commutant theorem that E is the commutant of F and the fact that the commutant of F is precisely $\mathrm{End}_{\mathbb{Q}_\ell[G_k]}(V_\ell(A))$. □

REMARK 2.2.1.1. Steps I-IV are valid for any field k, and do not use the fact that k is finite. Step V makes use of a certain hypothesis which is easily seen to be true for finite fields (and which can be shown with more difficulty for any global field). We therefore state it as a hypothesis, verify it for a finite fields and deduce its consequences.

HYPOTHESIS (k, A, ℓ). For A an abelian variety, k a field, and $\ell \neq \mathrm{char}\, k$ a prime, let $d \geq 1$ be an integer. Then there exists (up to isomorphism) only a finite number of abelian varieties B such that:

(i) there is an ample line bundle $\mathscr{L}$ on B with $\chi(\mathscr{L}) = d$;

(ii) there exists a k-isogeny $B \to A$ of ℓ-power degree.

For finite fields, we have the following stronger statement:

LEMMA 2.2.1.2. If k is a finite fields, $d > 0$ and $g > 0$, there are (up to isomorphism) only finitely many abelian varieties B defined over k of dimension g and carrying an ample line bundle $\mathscr{L}$ with $\chi(\mathscr{L}) = d$.

Proof. The line bundle $\mathscr{L}^3$ is very ample (Thm 1.3.3.10) and thus gives a projective embedding of $B \hookrightarrow \mathbb{P}_k^{3^g\chi(\mathscr{L})-1} = \mathbb{P}_k^{3^g d-1}$. The degree of this subvariety is the g-fold intersection number of $\mathscr{L}^3$, that is $3^g(\mathscr{L})^g = 3^g g! d$ by Riemann-Roch (Thm. 1.3.3.1). Thus, it corresponds to a k-rational point of a Chow variety of cycles of dimension g and degree $3^g g! d$ in $\mathbb{P}^{3^g d-1}$. Since k is a finite field, there are only finitely many k-rational points on this Chow variety, which proves the lemma. □

Step V. Let A be an abelian variety and $\mathscr{L}$ an ample line bundle. Suppose for a prime $\ell \neq \mathrm{char}\, k$, Hypothesis (k, A, ℓ) holds. Let $W \subset V_\ell(A)$ be a subspace of $V_\ell(A)$ which is G_k-stable and is maximal isotropic for the skew symmetric form $E^{\mathscr{L}}$. Then there is an element $u \in E$ with $u(V_\ell(A)) = W$.

Proof. For ease of notation, set $T = T_\ell(A)$ and $V = V_\ell(A)$ and for each integer $n \geq 0$, let

$$T_n = (T \cap W) = \ell^n T.$$

If $\psi_n : T_n \to A_{\overline{k}}[\ell^n]$ is the natural homomorphism (where in this case, we are treating T_n as a subgroup scheme of the pro-finite group scheme $T_\ell(A)$. Let $K_n = \psi_n(T_n)$ be the scheme-theoretic image, and $B_n = A_{\overline{k}}/K_n$ and let $\pi_n : A_{\overline{k}} \to B_n$ be the natural homomorphism. Then $[\ell^n]_{A_{\overline{k}}}$ factors as

$$A_{\overline{k}} \xrightarrow{\pi_n} B_n \xrightarrow{\lambda_n} A_{\overline{k}}.$$

Since $\pi_n \circ \lambda_n \circ \pi_n = \ell^n \pi_n$, we get that $\pi_n \circ \lambda_n = [\ell^n]_{B_n}$, so that

$$\lambda_n(B_n[\ell^n]) = \ker \pi_n = K_n.$$

Furthermore, for any $m \geq n$, $\lambda_n(B_n[\ell^n]) \supset (\lambda_n \circ \pi_n)(A_{\overline{k}}[\ell^m]) = A_{\overline{k}}[\ell^{m-n}]$, so that

$$T_\ell(\lambda_n)(T_\ell(B_n)) \supset \ell^n T_\ell(A_{\overline{k}}),$$

and

$$T_\ell(\lambda_n)(T_\ell(B_n)) = T_n.$$

We now want to verify that B_n is actually defined over k, i.e. that there exists an abelian variety B'_n over k such that $(B'_n)_{\overline{k}} = B_n$. But since W and hence $(T \cap W) + \ell^n T$ are G_k-stable, K_n is also G_k-stable. Thus, by Galois descent, $\pi_n : A_{\overline{k}} \to B_n$ is defined over k. Thus, B_n descends to some k-scheme B'_n. We want B'_n to have the structure of an abelian variety over k.

We will show this generally. First, we need the following theorem of Lang:

> THEOREM 2.2.1.3 (Lang). Let A_0 be a scheme over $k = \mathbb{F}_q$ such that $A = (A_0)_{\overline{k}}$ is an abelian variety. Then A_0 has at least one point rational over $\mathbb{F}_q$.

Proof. If π is the Frobenius morphism on A, then π is of the form $\pi(a) = a_0 + f(a)$ for some point $a_0 \in A(\overline{k})$ and some $f \in \operatorname{End} A$. Then $1 - f$ is an endomorphism of A. Since π and hence f induce the zero map on $\operatorname{Lie}(A)$, $1 - f$ induces the identity on $\operatorname{Lie}(A)$. Therefore $\ker(1 - f)$ is 0-dimensional and so $1 - f$ is surjective. Then if $(1 - f)(a_1) = a_0$, it follows that $a_1 = a_0 + f(a_1) = \pi(a_1)$, hence a_1 is rational over $\mathbb{F}_q$. □

> LEMMA 2.2.1.4. Suppose $A/\overline{k}$ is an abelian variety over $\overline{k}$ such that $A = (A_0)_{\overline{k}}$ for some k-scheme A_0. Then A_0 is an abelian variety over k.

REMARK 2.2.1.5. This can be generalized to all field extensions K/k by using "spreading out and specialization" arguments if $K \supset \overline{k}$, and using this lemma to assume K is algebraically closed.

Proof. Firstly, by Lang's theorem 2.2.1.3, by translation, we may choose the identity $0 \in A$ to be k-rational. Since $\overline{k} = \varinjlim K$ where K/k is finite, there exists some extension K/k large enough so that $m : A \to A$, $i : A \to A$ are defined over K. Since an element $g \in \operatorname{Gal}(K/k)$ acts by a morphism $g : A \to A$ fixing 0 (since 0 is k-rational), by Corollary 1.1.1.3, $\operatorname{Gal}(K/k)$ acts by homomorphisms. But this precisely means that $m : A \times A \to A$ and $i : A \to A$ are $\operatorname{Gal}(K/k)$-stable. Thus, by Galois descent, they descend to morphisms $m : A_0 \to A_0$ and $i : A_0 \to A_0$ satisfying the necessary commutative diagrams. □

Hence, returning to our case, B'_n is an abelian variety over k. Now, since π_n and $[\ell^n]_{A_{\overline{k}}} = \lambda_n \circ \pi_n$ are

defined over k, so is λ_n. Let d be the degree of the ample line bundle $\mathscr{L}$ on A. We shall produce an ample line bundle of degree $2^g d$ on B_n'. To do so, observe that

(i) if $B = (B_0)_{\overline{k}}$ with B_0 an abelian k-variety, and we construct $\widehat{B} \simeq B/K(\mathscr{L})$ for some ample $\mathscr{L} = (\mathscr{L}_0)_{\overline{k}}$ with $\mathscr{L}_0$ ample on B_0, then via this identification clearly $\widehat{B} = (\widehat{B_0})_{\overline{k}}$ as quotients behave well under extension of base field. Furthermore, the Poincaré bundle $\mathscr{P}_B = (\mathscr{P}_{B_0})_{\overline{k}}$;

(ii) if $\phi_{\lambda_n^*\mathscr{L}} = \ell^n \psi$, since $\lambda_n^*\mathscr{L}$ is descends to k, so does $\phi_{\lambda_n^*\mathscr{L}}$ and hence also ψ; and finally

(iii) if $\mathscr{P}_{B_n}$ is the Poincaré bundle on $B_n \times \widehat{B}_n$ and $(1,\psi) : B_n \to B_n \times \widehat{B}_n$, then $(1,\psi)^*\mathscr{P}_{B_n} = \mathscr{N}_n$ is a line bundle which descends to k with $\phi_{N_n} = 2\psi$ (see the proof of Prop. 2.1.3.7 for details). So,

$$\begin{aligned}\chi(\mathscr{N}_n)^2 &= \deg \phi_{\mathscr{N}_n} \\ &= 2^{2g} \deg \psi \\ &= 2^{2g}\ell^{-2ng} \deg \lambda_n^*(\mathscr{L}) \\ &= 2^{2g}\ell^{-2ng}(\deg \lambda_n)^2 \chi(\mathscr{L})^2 \\ &= 2^{2g}\chi(\mathscr{L})^2.\end{aligned}$$

We deduce from Hypothesis (k, A, ℓ) that there is an infinite set $I \subset \mathbb{Z}_{\geq 0}$ with smallest integer n and isomorphisms $v_i : B_n' \xrightarrow{\sim} B_i$ for all $i \in I$. Consider the elements $u_i = \lambda_i v_i \lambda_n^{-1} \in \mathrm{End}^0 A_{\overline{k}}$ and their images $u_i' \in \mathrm{End}_{\mathbb{Q}_\ell}(V_\ell)$. We have that $u_i'(T_n) = T_i \subset T_n$ for all $i \in I$, and since $\mathrm{End}_{\mathbb{Z}_\ell}(T_n)$ is compact, we can select a subsequence $(u_j')_{j\in J}$ which converges to a limit u'. Since E is closed in $\mathrm{End}_{\mathbb{Q}_\ell}(V_\ell)$ and $u_j' \in E$, u' also belongs to E. Since T_n is compact, $u'(T_n)$ consists of elements of the form $a = \lim x_j$ where $x_j \in u_j(T_n) = T_j$, and since the sets T_j are decreasing, it follows that

$$u(T_n) = \bigcap_{j\in J} u_j(T_n) = \bigcap_{j\in J} T_j = T \cap W.$$

Hence $u(V) = W$. □

Step VI. Suppose that for any finite algebraic extension k_1 of k_0, Hypothesis (k, A, ℓ) holds, and that F is isomorphic as a $\mathbb{Q}_\ell$-algebra to a direct product of copies of $\mathbb{Q}_\ell$. Then the map 2.2.1-II is an isomorphism.

Proof. Let D be the commutant of E in $\mathrm{End}_{\mathbb{Q}_\ell}(V_\ell(A))$, so that $D \supset F$. We first show that any isotropic subspace W for $E^{\mathscr{L}}$ which is F-stable is also D-stable. We proceed by downward induction on $\dim W$. If W is maximal isotropic, i.e., $\dim W = g$, we can by Step V find some $u \in E$ such that $u(V) = W$, and hence

$$DW = DuV = uDV = vV = W,$$

which proves the assertion in the maximal case.

Suppose then that $\dim W = r < g$ and the assertion holds for F-stable isotropic subspaces of dimension $r+1$. The orthogonal complement $W^\perp$ of W with respect to $E^{\mathscr{L}}$ is also F-stable, since $E^{\mathscr{L}}$ is invariant under the action of $\mathrm{Gal}(\overline{k}/k)$. Further, since any simple F-module is one dimensional and

$$\dim W^\perp - \dim W = 2g - 2\dim W = 2(g-r) \geq 2(g-g+1) = 2,$$

we can find F-stable one-dimensional subspaces L_1 and L_2 of W such that the sum $W = L_1 + L_2$ is direct. By the induction hypothesis, $W + L_1$ and $W + L_2$ are D-stable, hence so is their intersection W. This completes the induction.

We deduce that any eigenvector for F in V is also an eigenvector for D. It follows that $D \subset F$ (the decomposition of V into factors V_i corresponding to the simple factors of F reduces this assertion to the evident statement that an endomorphism of V_i for which every element for which every element of V_i is an eigen-vector is a scalar multiplication). Hence $F = D$. □

Step VI. Finishing the proof.

We assume from now on that $k = \mathbb{F}_q$ is finite. Let π be the Frobenius morphism over k. Then π belongs to the center of $\mathrm{End}^0 A$. We shall first show that there are infinitely many primes ℓ for which $\mathbb{Q}_\ell \otimes \mathbb{Q}[\pi]$ is isomorphic as a $\mathbb{Q}_\ell$-algebra to a direct product of copies of $\mathbb{Q}_\ell$ (to use Step V).

In fact, writing $\mathbb{Q}[\pi] = K_1 \times \cdots \times K_r$ where K_i are finite extensions of $\mathbb{Q}$, it suffices to show that for infinitely many primes ℓ, each $\mathbb{Q}_\ell \otimes K_i$ is isomorphic to a product of copies of $\mathbb{Q}_\ell$. Let K be a finite Galois extension of $\mathbb{Q}$ in which all the K_i's are embeddable. Then it suffices to show that for infinitely many ℓ, $K \otimes_{\mathbb{Q}} \mathbb{Q}_\ell$ splits as a product of $\mathbb{Q}_\ell$'s as a $\mathbb{Q}_\ell$-algebra. It suffices for this that there is one simple factor of $K \otimes_{\mathbb{Q}} \mathbb{Q}_\ell$ isomorphic to $\mathbb{Q}_\ell$,. In fact, if

$$K \otimes_{\mathbb{Q}} \mathbb{Q}_\ell \simeq L_1 \times \cdots L_m,$$

$\mathrm{Gal}(K/\mathbb{Q})$ permutes the factors $\mathscr{L}_i$. It also acts transitively on the simple factors. For if not, suppose $L_1 \times \cdots \times L_r$ is π-stable; then the element $(1, \ldots, 1, 0, \ldots, 0) \in L_1 \times \cdots \times L_m$ is π-stable. On the other hand, since $\mathrm{Gal}(K/\mathbb{Q})$ only fixes the elements of $\mathbb{Q}$ in K, it fixes only the element, it fixes only the elements of $\mathbb{Q}_\ell$ in $\mathbb{Q}_\ell \otimes_{\mathbb{Q}} K$, that is, elements of the form $(\alpha, \ldots, \alpha) \in L_1 \times \cdots \times L_k$ with $\alpha \in \mathbb{Q}_\ell$. This proves the assertion.

Now, choose an algebraic integer $\alpha \in K$ generating K over $\mathbb{Q}$, and let $P(t) \in \mathbb{Z}[t]$ be its irreducible monic polynomial over $\mathbb{Q}$. Since $K \otimes_{\mathbb{Q}} \mathbb{Q}_\ell \simeq \mathbb{Q}_\ell[t]/(P(t))$, it is enough to find infinitely many ℓ for which $P(t)$ has a zero in $\mathbb{Q}_\ell$. Let Δ be the discriminant of $P(t)$, and ℓ any prime not dividing Δ such that $P(t) \equiv 0 \bmod \ell$ has a solution $n \in \mathbb{Z}$. Then $F'(n) \not\equiv 0 \bmod \ell$, so that by Hensel's lemma, n can be refined to a root of P in $\mathbb{Z}_\ell$. Thus, we have reduced to proving the following

> LEMMA 2.2.1.6. Let $P(t) \in \mathbb{Z}[t]$ be a non-constant polynomial. Then there are infinitely many primes ℓ for which $P(t) \equiv 0 \bmod \ell$ has a solution in $\mathbb{Z}$.

Proof. Let $P(t) = a_n t^n + a_{n-1} t^{n-1} + \cdots + a_0$. Since 0 is a root if $a_0 = 0$, we may assume $a_0 \neq 0$. Further, by substituting $a_0 t$ for t in P and removing the common factor a_0, we may assume $a_0 = 1$. Let S be a finite set of primes p. If $N = \prod_{p \in S} p$, then

$$P(\nu N) = a_n \nu^n N^n + \cdots + a_{n-1} \nu N + 1 \equiv 1 \bmod N,$$

so that no prime of S divides $P(\nu N)$. On the other hand $F(\nu N) \neq \pm 1$ for ν large, hence has a prime factor ℓ not belonging to S. □

Next, we show that for all $\ell \neq \operatorname{char} k$, the dimension of $\operatorname{End}_{\mathbb{Q}_\ell[G_k]} V_\ell$ (to apply Step III). Since π belongs to the center of $\operatorname{End}^0 A$, $\mathbb{Q}_\ell[\pi]$ is semisimple and V_ℓ is a $\mathbb{Q}_\ell[\pi]$-module, so that the image π' of π in $\operatorname{End}_{\mathbb{Q}_\ell} V_\ell$ is semisimple. The characteristic polynomial $P(t)$ of π' in $\operatorname{End}_{\mathbb{Q}_\ell} V_\ell$ has coefficients in $\mathbb{Z}$ independent of ℓ. Further, $\operatorname{End}_{\mathbb{Q}_\ell[G_k]} V - \ell$ is the commutant of π' in $\operatorname{End}_{\mathbb{Q}_\ell} V_\ell$. Thus, our assertion is a consequence of the following lemma, applied to an algebraic closure of $\mathbb{Q}_\ell$.

LEMMA 2.2.1.7. Let Φ and Ψ be absolutely semisimple endomorphisms of two finite-dimensional k-vector space V and W respectively, with characteristic polynomials P_Φ and P_Ψ. Let

$$P_\Phi = \prod p_i^{m(p_i)},$$
$$P_\Psi = \prod p_i^{n(p_i)}$$

be the decompositions of P_Φ and P_Ψ as products of powers of distinct irreducible monic polynomials p_i. Then the vector space

$$E = \{\phi \in \operatorname{Hom}_k(V, W) \mid \phi\Phi = \Psi\phi\}$$

has dimension $r(P_\Phi, P_\Psi) = \sum m(p_i)n(p_i) \deg p$, and this integer is invariant under any extension of the base field k.

Proof. Make V (resp. W) into a $k[t]$-module by making t act through Φ (resp. Ψ). Denote the $k[t]$-modules $k[t]/(p_i(t))$ by M_{p_i}. Because of our assumption of semisimplicity, we have isomorphisms of $k[t]$-modules

$$V \simeq \prod M_{p_i}^{m(p_i)},$$
$$W \simeq \prod M_{p_i}^{n(p_i)}.$$

Now, the M_p are non-isomorphic simple $k[t]$-modules for distinct p_i, and E is nothing but $\operatorname{Hom}_{k[t]}(V, W)$. Since $\dim_k \operatorname{Hom}_{k[t]}(M_{p_i}, M_{p_i}) \dim_k M_{p_i} = \deg p_i$ and since E clearly commutes with base extension, the lemma follows. □

Theorem 2.2.0.1 is a consequence of what we just proved together with Steps III and V.

§2.2.2 Applications and Tate's Isogeny Theorem

In this section, we will give some easy consequences of Theorem 2.2.0.1, including Tate's Isogeny Theorem (2.2.0.3). We shall consistently use the notation $r(f, g)$ for two polynomials defined in Lemma 2.2.1.7. Fix $k = \mathbb{F}_q$, A an abelian variety over k, and let π be the Frobenius morphism, which we have seen induces an absolutely semisimple endomorphism of $V_\ell(A)$ for all $\ell \neq \operatorname{char} k$. Thus, the structure of $V_\ell(A)$ as a module over $G_k := \operatorname{Gal}(\overline{k}/k)$ is uniquely determined by the characteristic polynomial of π (as described in the proof of Lemma 2.2.1.7.

We now have

THEOREM 2.2.2.1. Let A, B be abelian varieties over $k = \mathbb{F}_q$ and let P_A and P_B be the characteristic polynomials of their Frobenius endomorphism. Then,

(a) we have

$$\operatorname{rank}(\operatorname{Hom}_k(A, B)) = r(P_A, P_B);$$

(b) the following are equivalent:

(b_1) B is isogenous to an abelian subvariety of A defined over k,

(b_2) $V_\ell(B)$ is G_k-isomorphic to a G-subspace of $V_\ell(A)$ for some ℓ;

(b_3) P_B divides P_A;

(c) the following are equivalent:

(c_1) A and B are k-isogenous,

(c_2) $P_A = P_B$,

(c_3) A and B have the same number of k'-rational points for every finite extension k'/k.

Proof. (a) follows from Theorem 2.2.0.1 and Lemma 2.2.1.7 since π generates G_k in a topological sense, so for the E in the lemma, $E = \mathrm{Hom}_{\mathbb{Q}_\ell[G_k]}(T_\ell(A), T_\ell(B))$, which has dimension $r(P_A, P_B)$ by definition.

The implications (b_2) $\implies$ (b_2) is obvious and (b_2) $\iff$ (b_3) is clear since the structure of $V_\ell(A)$ is uniquely determined as a Galois module by the characteristic polynomial of π. Thus, we must show (b_2) $\implies$ (b_1). Suppose (b_2) holds so that we can find an injective G_k-equivariant homomorphism $\varphi : T_\ell(B) \hookrightarrow T_\ell(A)$. Then φ is in the image of $\mathbb{Z}_\ell \otimes_{\mathbb{Z}} \mathrm{Hom}(A, B)$, so that we can find some $\psi \in \mathrm{Hom}(A, B)$ such that $T_\ell(\psi)$ is injective (as $\mathrm{Hom}(A, B)$ is dense). We know that $B' := (\ker \psi)^0_{\mathrm{red}}$ is an abelian variety (see Thm. 1.2.6.9) with $T_\ell(B') = \ker T_\ell(\psi)$, but since we assumed $T_\ell(\psi)$ is injective, $T_\ell(B') = 0$, so B' is trivial and thus ψ is an isogeny.

The equivalence (c_1) $\iff$ (c_2) is a special case of (b_2) $\iff$ (b_3) when $\dim A = \dim B$. Further, we have seen from proving the Riemann Hypothesis (Thm. 2.1.6.2) that if ω_i ($i \in I$), and ω'_j ($j \in J$) are the roots of P_A and P_B respectively and N_m and N'_m are the number of rational points of A and B respectively in the extension of degree n, $\mathbb{F}_{q^n}/k$, we have

$$N_m = \prod_{i \in I} (1 - \omega_i^m),$$
$$N'_m = \prod_{j \in J} (1 - (\omega'_j)^m).$$

Thus, we have to show that $P_A = P_B$ if and only if $\prod_i (1 - \omega_i^m) = \prod (1 - (\omega'_j)^m)$ for all $n \geq 0$. The forward implication is obvious. For the other implication, note that

$$\prod_{i \in I} (1 - \omega_i^m) = \sum_{S \subset I} (-1)^{|S|} \omega_S^m,$$
$$\prod_{j \in J} (1 - (\omega'_j)^m) = \sum_{T \subset I} (-1)^{|T|} (\omega'_T)^m.$$

where $|S|$ and $|T|$ denote the cardinalities, and $\omega_S = \prod_{i \in S} \omega_i$ and $\omega'_T = \prod_{j \in T} \omega'_j$. Multiplying the equations by t^m, where t is a variables, and summing as formal power series, we obtain

$$\sum_{S \subset I} (-1)^{|S|} \frac{1}{1 - t\omega_S} = \sum_{T \subset J} (-1)^{|T|} \frac{1}{t\omega_T}.$$

Since $|\omega_i| = |\omega'_j| = q^{1/2}$ (by the Riemann Hypothesis), comparing the poles on both sides on the circle $|t| = q^{1/2}$, we obtain that there is a bijection $\sigma : I \to J$ such that $\omega_i = \omega'_{\sigma(j)}$. Hence $P_A = P_B$. □

We now see that Tate's Isogeny Theorem falls out as an easy consequence of what we have just proved:

Proof of 2.2.0.3. We see that Tate's Isogeny Theorem (2.2.0.3) is an immediate corollary of part (c) of Theorem 2.2.2.1 by noting that $P_A = P_B$ if and only if $V_\ell(A) \simeq V_\ell(B)$ as Galois modules by Lemma 2.2.1.7. □

Relating our results to our study of endomorphisms in §2.1.2 and §2.1.5, we get

THEOREM 2.2.2.2. Let A be an abelian variety of dimension g over $k = \mathbb{F}_q$. Let π be the Frobenius endomorphism of A and P its characteristic polynomial. Then we have the following:

(a) The algebra $F = \mathbb{Q}[\pi]$ is the center of the semisimple algebra $E = \mathrm{End}^0(A)$;

(b) $\mathrm{End}^0(A)$ contains a semisimple $\mathbb{Q}$-subalgebra A of rank $2g$ which is maximally commutative;

(c) the following are equivalent:

(c_1) $[E : \mathbb{Q}] = 2g$,

(c_2) P has no multiple root,

(c_3) $E = F$,

(c_4) E is commutative;

(d) the following are equivalent:

(d_1) $[E : \mathbb{Q}] = (2g)^2$,

(d_2) P is a power of a linear polynomial,

(d_3) $F = \mathbb{Q}$,

(d_4) $E \simeq \mathrm{Mat}_g(D_p)$ where D_p is the quaternion division algebra over $\mathbb{Q}$ which splits at all places $\ell \neq p, \infty$,

(d_5) A is isogenous to the gth power of a supersingular curve;

(e) A is isogenous to a power of a simple abelian variety if and only if P is a power of a $\mathbb{Q}$-irreducible polynomial. When this is the case, E is a central simple algebra over F which splits at all finite primes of F not dividing p, but does not split at any real prime of F.

Proof. **Part (a).** It follows from the Theorem 2.2.0.1 that $F_\ell = \mathbb{Q}_\ell \otimes_{\mathbb{Q}} F$ is the center of $E_\ell = \mathbb{Q}_\ell \otimes_{\mathbb{Q}} E$, which proves that F is the center of E.

Part (b). Suppose $E = D_1 \times \cdots \times D_r$ is the expression of E as a product of simple algebras D_i with centers K_i (Cor. 2.1.2.4). Let $[K_i : \mathbb{Q}] = e_i$ and $[D_i : K_i] = d_i^2$. We can choose subrings L_i of D_i containing K_i with L_i semisimple and maximal commutative, $[L_i : K_i] = d_i$. Then $L = L_1 \times \cdots \times L_r$ is a semisimple $\mathbb{Q}$-subalgebra of E which is maximal commutative, and $[L : \mathbb{Q}] = \sum_{i=1}^r e_i d_i$. Now, for any ℓ, we have

$$\begin{aligned} E \otimes_{\mathbb{Q}} \mathbb{Q}_\ell &= \prod_{i=1}^r (D_i \otimes_{\mathbb{Q}} \mathbb{Q}_\ell) \\ &= \prod_{i=1}^r (D_i \otimes_{K_i} (K_i \otimes_{\mathbb{Q}} \mathbb{Q}_\ell)) \end{aligned}$$

$$= \prod_{i=1}^{r} \prod_{j=1}^{n_i} D_i \otimes_{K_i} K'_{ij}$$

where $K_i \otimes_{\mathbb{Q}} \mathbb{Q}_\ell = \prod_{j=1}^{n_i} K'_{ij}$ with K'_{ij} fields. On the other hand, if $P = \prod_{i=1}^{s} P_i^{m_i}$ is the decomposition of P over $\mathbb{Q}_\ell$ into a product of powers of irreducible polynomials over $\mathbb{Q}_\ell$, and if we consider $T_\ell(A)$ as a $\mathbb{Q}_\ell[t]$-module by making t act via π, we have an isomorphism of $\mathbb{Q}_\ell[t]$-modules

$$T_\ell(A) \simeq \prod_{\nu=1}^{s} (\mathbb{Q}_\ell[t]/(P_\nu))^{m_\nu} = \prod_{\nu=1}^{s} S_\nu^{m_\nu},$$

where $S_\nu = \mathbb{Q}_\ell[t]/(P_\nu)$ so that $E \otimes_{\mathbb{Q}} \mathbb{Q}_\ell$, being the commutant of π in $\operatorname{End} T_\ell(A)$, is isomorphic to

$$\prod_{\nu=1}^{s} \operatorname{Mat}_{m_\nu}(S_\nu).$$

Comparing the two factorizations of $E \otimes_{\mathbb{Q}} \mathbb{Q}_\ell$ and keeping in mind that K'_{ij} is the center of $D_i \otimes_{K_i} K'_{ij}$, we deduce that first, for any prime $\ell \neq p$ and any prime v of K_i lying over ℓ, D_i splits at v and second, there is a partition of $[1, s]$ into r disjoint subsets $I_1, \ldots, I_r$ such that

$$D_i \otimes_{\mathbb{Q}} \mathbb{Q}_\ell \simeq \prod_{j=1}^{n_i} D_i \otimes_{K_i} K'_{ij} \simeq \prod_{\nu \in I_i} \operatorname{Mat}_{m_\nu}(S_\nu).$$

It follows that $m_\nu = d_i$ for $\nu \in I_i$ and

$$\sum_{\nu \in I_i} [S_\nu : \mathbb{Q}_\ell] = \sum_{j=1}^{n_i} [K'_{ij} : \mathbb{Q}_\ell] = [K_i : \mathbb{Q}] = e_i,$$

so that

$$\begin{aligned} \sum_{i=1}^{r} e_i d_i &= \sum_{i=1}^{r} \sum_{\nu \in I_i} m_\nu [S_\nu : \mathbb{Q}_\ell] \\ &= \sum_{\nu=1}^{s} m_\nu [S_\nu : \mathbb{Q}_\ell] \\ &= \sum_{\nu=1}^{s} m_\nu \deg P_\nu = \deg P = 2g. \end{aligned}$$

Part (c). Since F is the center of E and E contains a maximal commutative subring of rank $2g$, (c_1), (c_3), (c_4) are equivalent, and since is E commutative if and only if E_ℓ is commutative, which by the above notations is equivalent to $m_\nu = 1$, which is equivalent to (c_2).

Part (d). Now $[E : \mathbb{Q}] = (2g)^2$ if and only if $\mathbb{Q}_\ell \otimes E \simeq \operatorname{Mat}_{2g}(\mathbb{Q}_\ell)$, hence if and only if $s = 1$, $S_\nu = \mathbb{Q}_\ell$ or equivalently, P is a power of a linear polynomial. In this case, $\mathbb{Q}_\ell$ is the center of $\mathbb{Q}_\ell \otimes E$, so that $\mathbb{Q}$ is the center of E, and conversely, if this holds, $\mathbb{Q}_\ell \otimes E$ is the commutant of $\mathbb{Q}_\ell$ in $\operatorname{End} V_\ell$, so that it is the whole of $\operatorname{End} V_\ell$. Thus, ($d_1$), ($d_2$) and ($d_3$) are equivalent.

If $\mathbb{Q}_\ell \otimes E = \operatorname{Mat}_{2g}(\mathbb{Q}_\ell)$, E is a central simple algebra over $\mathbb{Q}$ whose invariants in the Brauer group at all finite primes $\ell \neq p$ are 0 (see Theorem 2.1.5.2). Since its invariant at the infinite prime is 0 or $\frac{1}{2}$ and the sum of invariants at all primes is 0, E is either $\operatorname{Mat}_{2g}(\mathbb{Q})$ or $\operatorname{Mat}_g(D_p)$ where D_p is the quaternion algebra over $\mathbb{Q}$ splitting at all finite primes $\ell \neq p$. The first possibility can be ruled out, since A cannot be a product of $2g$ abelian varieties. This proves the equivalence between (d_1) and (d_4). In view of our remarks preceding the theorem, (d_4) is equivalent to saying $A \simeq C^g$ where C is an elliptic curve with $\operatorname{End} C \simeq D_p$.

We have then shown that C is supersingular (see Deuring's Theorem 2.1.9.1).

Part (e). Let Q be the product of distinct irreducible factors of P. This is the minimal polynomial of π over $\mathbb{Q}$ since, considering the action of π as an endomorphism of V_ℓ, any irreducible factor over $\mathbb{Q}_\ell$ of the characteristic polynomial P divides the minimal polynomial of π. Now, A is isogenous to a power of a simple abelian variety if and only if E is simple, hence only if the center $F = \mathbb{Q}[\pi]$ of E is a field. Since $F \simeq \mathbb{Q}[t]/(Q(t))$, F is a field if and only if Q is irreducible, or equivalently P is the power of an irreducible polynomial.

We have shown earlier that E splits at any finite prime v of F not dividing ℓ. Suppose $\sigma : F \hookrightarrow \mathbb{R}$ is a real embedding of F. Since $v(\pi)$ satisfies $P(v(\pi)) = 0$ and the roots of P have absolute value $q^{1/2}$ (by the Riemann Hypothesis 2.1.6.2), we must have that $\sigma(\pi) = \pm q^{1/2}$. If q is a square, $\sigma(\pi) \in Q$ and $F = \mathbb{Q}$, so that the equivalence of (d_3) and (d_4) implies that E does not split at ∞. If q is not square, $F = \mathbb{Q}(\sqrt{p})$ where $p = \operatorname{char} k$. Let k' be the quadratic extension of k and $\pi' = \pi^2$, the Frobenius over k'. Then $\pi^2 \in \mathbb{Q}$, so that the center F' of $E' = \mathrm{End}_{k'}(A_{k'})$ is $\mathbb{Q}$. Using (d), we conclude that $E' = \mathrm{Mat}_g(D_p)$. On the other hand, we have $F' \subset F \subset E \subset E'$, and E is the commutant of F in E'. By a known result on central simple algebras ([Pie82, §12.7, Theorem 1]), E and $F \otimes_{F'} E'$ define the same element of the Brauer group over F. Since both real primes of $\mathbb{Q}(\sqrt{p})$ lie over the real prime ∞ of $\mathbb{Q}$ and E' has invariant $\frac{1}{2}$ at ∞ with respect to $F' = \mathbb{Q}$, E has invariant $\frac{1}{2}$ at either of the real primes of $F = \mathbb{Q}(\sqrt{p})$. □

Now, we conclude with an interesting result about elliptic curves.

> COROLLARY 2.2.2.3. Let C_1, C_2 be elliptic curves over a finite field k such that $\mathrm{End}^0(C_1)_{\overline{k}} = \mathrm{End}^0(C_2)_{\overline{k}}$. Then $(C_1)_K$ is isogenous to $(C_2)_K$ for K/k some finite extension. In particular, any two supersingular elliptic curves over a finite field k are isogenous after base change to K for some finite extension K/k.

Proof. Suppose that C_1 and C_2 are supersingular elliptic curves over k. We can choose a field extension k'/k such that $\mathrm{End}_{k'} C_1$ and $\mathrm{End}_{k'} C_2$ are quaternion algebras over $\mathbb{Q}$ (as seen in the proof of Deuring's Theorem 2.1.9.1), so that they have $\mathbb{Q}$ as a center. Thus, their Frobenius morphisms π_{C_1} and π_{C_2} lie in $\mathbb{Q}$. Since they must both have absolute value $\sqrt{q}$ where $q = \#(k')$, we see that $\pi_{C_1}^2 = \pi_{C_2}^2 = q$. Thus, if k'' is the quadratic extension of k', there is an isomorphism $T_\ell(C_1) \xrightarrow{\sim} T_\ell(C_2)$ carrying the action of $\pi'_{C_1} = \pi_{C_1}^2$ into $\pi'_{C_2} = \pi_{C_2}^2$. By Theorem 2.2.0.3, C_1 and C_2 are isogenous over k''.

Next suppose $\mathrm{End}^0 C_1 \simeq K \simeq \mathrm{End}^0 C_2$ for some imaginary quadratic extension $K/\mathbb{Q}$. Extend the base field such that all the points of order p in C_1 and C_2 are rational. Now, $\mathbb{Q}_p \otimes_Q \mathrm{End}^0 C_1$ admits a one-dimensional representation in $T_p(C_1)$. Hence, $p = \mathfrak{p}\mathfrak{p}'$ in K and $\mathbb{Q}_p \otimes_{\mathbb{Q}} \mathrm{End}^0 C_1 \simeq K_{\mathfrak{p}} \times K_{\mathfrak{p}'}$.

Suppose that $\mathbb{Q}_p \otimes_{\mathbb{Q}} \mathrm{End}^0 C_1$ acts on $T_p(C_1)$ via $K_{\mathfrak{p}}$. It follows that $\pi_{C_1} \equiv 1 \bmod \mathfrak{p}$, and since $\mathrm{Nm}_{K/\mathbb{Q}}(\pi_{C_1})$ is a power of p, π_{C_1} has to be a power of $\mathfrak{p}$ in the ring of integers of K. A similar assertion (possibly swapping the roles of $\mathfrak{p}$ and $\mathfrak{p}'$) holds for π_{C_2}. By altering the isomorphism $\mathrm{End}^0 C_2 \simeq K$ by conjugation of K if necessary, we may assume that π_{C_2} is also a power of $\mathfrak{p}'$. Since $\mathrm{Nm}_{K/\mathbb{Q}}(\pi_{C_1}) = \mathrm{Nm}_{K/\mathbb{Q}}(\pi_{C_2}) = q = p^f$, $\pi_{C_1} = \pi_{C_2} = (\mathfrak{p}')^f$ so that π_{C_1} and π_{C_2} differ by a unit, i.e., a root of unity since K is imaginary quadratic. Thus, $\pi_{C_1}^n = \pi_{C_1}^n$ in K for suitable n, so they have the same minimal equation over $\mathbb{Q}$, of degree 2. Since this is their characteristic polynomial, C_1 and C_2 are isogenous over an extension of degree n of k. □

Chapter 3

Abelian Schemes

As mentioned in the proof of Tate's Conjecture for Finite Fields (Thm. 2.2.0.1), the crucial hypothesis is a finiteness condition on abelian varieties with certain conditions. It turns out that this is exactly a finiteness condition on the points of a level structure on the moduli space of abelian varieties.

We now give a more general definition which will become important when discussing moduli spaces of abelian varieties.

Definition 3.0.0.1. An **abelian scheme** over a base scheme S of relative dimension g is a smooth, proper group scheme whose fibers are (geometrically) connected and of dimension g. In particular every fiber is an abelian variety.

Chapter 4

p-divisible Groups

Part II

Moduli Spaces

Chapter 5

Algebraic Spaces and Stacks: A Primer

"*...humanity used to think of [real manifolds] as sitting in $\mathbb{R}^n$, but there is really very little advantage in sticking them in $\mathbb{R}^n$, other than psychological, so you should give that up.*"

From Notes for 245C in Spring 2022
Ravi Vakil

Algebraic spaces and stacks are categorical generalizations of schemes and sheaves. When one first works with examples of them, it is not always obvious, for example, that there is some algebraic space we care about that is clearly not a scheme. The reason this generalization is convenient then, is (as we will see), many useful constructions turn out to be very cheaply algebraic spaces, but not very obviously schemes. The goal of this chapter will be to show that many schemes are

Section 5.1

Moduli Functors

Recall that any object $X \in \underline{\mathrm{Sch}}$ admits a **functor of points**, which is the contravariant functor

$$\begin{aligned} h_X : \underline{\mathrm{Sch}} &\to \mathrm{Set} \\ T &\mapsto \mathrm{Hom}(T, X), \end{aligned}$$

with morphisms being carried to the expected set morphisms induced by pullback. As a consequence of Yoneda's Lemma, if there exists an isomorphism of functors $h_X \cong h_Y$, then the identification of $\mathrm{Hom}(Y, Y)$ with $\mathrm{Hom}(X, Y)$ and $\mathrm{Hom}(X, X)$ with $\mathrm{Hom}(Y, X)$ yields an isomorphism of schemes. In essence, a scheme is determined up to (canonical) isomorphism by its functor of points.

Definition 5.1.0.1. A contravariant functor $F : \underline{\mathrm{Sch}} \to \mathrm{Set}$ is said to be **representable** if there exists some $X \in \underline{\mathrm{Sch}}$ such that $F \cong h_X$. In such a case X is said to be a **fine moduli space** for F.

Remark 5.1.0.2. One could replace $\underline{\mathrm{Sch}}$ in the above definitions with an arbitrary category $\mathcal{C}$. One could also, by considering $T \mapsto \mathrm{Hom}(X, T)$, replace h_X with the *covariant* functor of points h^X; this is mostly a matter of notational choice and the former convention is most commonly used.

Most categorical constructions with universal properties can actually be thought of as representations of functors without too much extra effort, as in the following example.

Example 5.1.0.3 (Fiber product as a representable functor). Fix a base scheme S and consider the category of S-schemes $\underline{\mathrm{Sch}}/S$. Fix two objects $X, Y \in \underline{\mathrm{Sch}}/S$, with corresponding structure morphisms f_X, f_Y. Then consider the functor

$$\begin{aligned} \underline{\mathrm{Sch}}/S &\to \mathrm{Set} \\ Z &\mapsto \{(f, g) \in \mathrm{Hom}(Z, X) \times \mathrm{Hom}(Z, Y) : f_X \circ f = g_X \circ g\}. \end{aligned}$$

It is not too much work to verify that this is indeed a functor. If this functor is representable by some object $W \in \underline{\mathrm{Sch}}/S$, we know that W is unique up to canonical isomorphism, and by definition maps to $Z \to W$ would be in bijective correspondence with pairs of maps $Z \to X, Y$ that make the relevant square with the structure morphisms commute. But this condition is *precisely* the universal property of the fiber product, thus

$$W := X \times_S Y$$

is exactly the S-scheme representing this functor.

Sanity Check. By the universal property of the fiber product, representability plays nicely with base change, i.e if $F : \underline{\mathrm{Sch}} \to \mathrm{Set}$ is representable by an object X, then the restriction of that functor $F : \underline{\mathrm{Sch}}/S \to \mathrm{Set}$ is representable by X_S.

As one might pessimisticly expect, it is not the case that all functors are representable. We want to understand certain necessary conditions for a functor to be representable.

Definition 5.1.0.4. Given $\mathfrak{M}$ a class of morphisms in $\underline{\mathrm{Sch}}/S$ which contains all isomorphisms and is stable under fiber products and composition, a contravariant functor $F : \underline{\mathrm{Sch}}/S \to \mathrm{Set}$ is an **$\mathfrak{M}$-sheaf** if the following two conditions are satisfied:

1. For any family of schemes $\{T_i\}_{i\in I}$, the canonical morphism $F\left(\coprod T_i\right) \to \prod F(T_i)$ is an isomorphism.
2. For all morphisms $T' \to T$ in $\mathfrak{M}$, the sequence

$$F(T) \to F(T') \rightrightarrows F\left(T' \times_T T'\right)$$

is exact (here the double arrows denote the projections $T' \times_T T' \to T$).

Oftentimes the class of morphisms $\mathfrak{M}$ is taken to be those of the type $\bigsqcup T_i \to T$, where each $T_i \to T$ is an open immersion and $\{T_i\}$ is an open covering of T. In this case we will call F a **Zariski sheaf**, or often just a sheaf. We can also similarly consider covers $\bigsqcup T_i \to T$ that are fpqc, fppf or étale. We will say F is an **fpcq, fppf** or **étale sheaf** if it is a Zariski sheaf and a sheaf with respect to each of the corresponding cover types. The key result is the following.

Proposition 5.1.0.5. Suppose $F : \underline{\mathrm{Sch}} \to \mathrm{Set}$ is a contravariant representable functor. Then F is an fpqc sheaf (and hence an fppf and étale sheaf).

Section 5.2

Sheaves on Sites

Section 5.3

Definitions and Key Properties

Chapter 6

Moduli of Stable Curves

Chapter 7

Moduli of Principally Polarized Abelian Varieties

Part III

Some Key Intermediary Constructions and Theories

Chapter 8

Néron Models

"*We shall state some of the main results using the language of schemes. However, lacking sufficient familiarity with this language, we deemed it more prudent to refrain from employing it systematically,... leaving the task of translation to the specialists.*"

Translation from *Modeles minimaux des varietes abeliennes.*
André Néron

In the 1960s, André Néron constructed the Néron models to study the abelian varieties over number fields. In particular, a natural question to ask is if S is the spectrum of a DVR, then does an abelian variety over its generic point extend smoothly to all of S? In general it is too much to ask for the special fiber to also be an abelian variety, in part because requiring both smoothness and properness is too rigid of a requirement. The idea of Néron was then to relax the properness condition by altering the valuative criterion requirement to something weaker, which is now called the "Néron mapping property." It turned out this construction had enormous use for number theorists and algebraic geometers, and has become a quintessential object in these fields of study.

The key result we will prove in this Chapter will be the following.

Theorem 8.0.0.1 (Abelian Variety over a Global Field admits a Néron Model)**.** Let A_K be an abelian variety over a global field K. Then A_K admits a Néron model over the ring of integers $\mathcal{O}_K$.

We will define the notion of Néron model in the very first subsection. Néron himself didn't employ Grothendieck's language of schemes – what has now become the modern status quo – the work of mathematicians since then has standardized the formulation and usage of the eponymous models in arithmetic geometry. The goal of the famous book "*Néron Models*" (by Bosch, Lütkebohmert & Raynaud) was to go back to the basics of their construction and rework them in the language of schemes and make it widely accessible to the algebraic geometry community at the time. The goal of this chapter is to give a more fleshed out exposition of certain parts of that book.

Outline of the Chapter

§8.1 goes through the definitions and key properties of Néron models, along with the big picture proof of the above theorem. The key intermediary result is Theorem 8.1.4.1. Assuming this result, our Main

Theorem is nothing but a special case of Theorem 8.1.5.4, which is at the end of the first section. All of the remaining sections are dedicated to the proof of 8.1.4.1.

Section 8.1

Models and their Uses

§8.1.1 Motivation and Definitions

In the theory of Néron models, one is usually interested in the constructions over a base scheme S which is a **Dedekind scheme**, i.e. a noetherian normal scheme of dimension ≤ 1. The local rings of such S are either fields or DVRs.

We will talk about the **local case** if S consists of a local scheme (i.e is the spectrum of a DVR or a field), and the general case will be referred to as the **global case**. Any Dedekind scheme S decomposes into a disjoint sum of finitely many irreducible components S_i with a generic point η_i each. We set $K := \bigoplus_i \mathscr{K}(\eta_i)$, so K is the **ring of rational functions on** S. Furthermore, the affine scheme $\mathrm{Spec}(K)$ is referred to as the **scheme of generic points** of S. Note there exists a canonical inclusion map

$$\mathrm{Spec}(K) \to S,$$

so for any X-scheme S we may speak of its **generic fiber** X_K as the base change to $\mathrm{Spec}(K)$.

Warning (Abuse of Notation). We will often write K in instances where we *really* mean $\mathrm{Spec}(K)$, for example in the notatation of the generic fiber above. We will use similar abuses of notation throughout this text, writing R in general when we really mean $\mathrm{Spec}(R)$. There should be virtually no situation where this shorthand causes any real confusion.

Example 8.1.1.1 (Dedekind Schemes of Interest). There are 3 primary examples of Dedekind schemes which are of interest.

- Let K be a number field with ring of integers $\mathcal{O}_K$, then $S := \mathrm{Spec}(\mathcal{O}_K)$ is an (affine) Dedekind scheme.
- Take an algebraic function field K of transcendence degree 1 over some chosen field of constants $\mathbb{F}$, and define S to be the normal proper $\mathbb{F}$-curve associated to K.
- Start with any 2-dimensional normal noetherian local scheme, and remove its closed point. This yields a Dedekind scheme.

We will see later [Insert example reference here] why we restrict ourselves to Dedekind schemes; in short it is to guarantee the existence of Néron models; there are known examples of schemes with certain hypotheses removed for which there does not exist a Néron model.

But before we get ahead of ourselves, we need to discuss the notion of an S-model.

Definition 8.1.1.2. For S a Dedekind scheme with ring of rational functions K, an **S-model** of a K-scheme X_K is an S-scheme Y with generic fiber X_K.

Sanity Check. Verify that X_K is an S-model for itself. Further, verify that that blowup of any S-model of X_K at a point outside the generic fiber is also an S-model.

A simple method for constructing S-models (of finite type) is that of **chasing denominators**. This applies to the case where S is affine (say $S = \mathrm{Spec}(R)$) and X_K is an affine of finite type over K.

Example 8.1.1.3 (S-Model via Chasing Denominators). Let K be the field of fractions of a Dedekind domain R, and let X_K be some finite-type affine K-scheme, i.e we set

$$A_K := K[t_1, \dots t_n]/I_K$$

for some ideal I_K, and $X_K := \mathrm{Spec}(A_K)$. Then I_K is generated by finitely many polynomials, which WLOG have coefficients in R (since we can write each coefficient as a fraction of elements of R, and cross multiply out without affecting the zero set). Thus defining

$$A := R[t_1, \dots t_n]/I,$$

where I is generated by the same polynomials generating I_K, we get an S-model of finite type.

The motivation for the notion of an S-model arises naturally when one is tackling problems of moduli spaces; wanting to extend some K-scheme to a family of schemes over some base S, for example. In order to use such a construction practically though, we would like S-models to have "nice" properties. For one, if X_K is a group scheme we may ask its S-model to have compatible group structure. For another, given a K-scheme with certain properties (e.g smooth, proper, flat, étale), we might ask for its S-model to have those properties over S. Unfortunately this is often too much to ask for; in particular, asking for smoothness *and* properness is too much.

Example 8.1.1.4. Consider the following projective closure of an elliptic curve:

$$X_{\mathbb{Q}} := \overline{\{y^2z - x^3 - 2x^2z - 6z^3\}} \subseteq \mathbb{P}^2_{\mathbb{Q}}.$$

One can compute the Jacobian of this curve to be $[-3x^2 - 4xz,\ 2yz,\ y^2 - 2x^2 - 18z^3]$. We can check that this is full rank everywhere; thus this curve is smooth via the Jacobi criterion [cite here]. Moreover it is also projective (hence proper!) as a closed subscheme of projective space.
The same equation defines a $\mathbb{Z}$-model of X, which is proper (projective even, as a closed subscheme of $\mathbb{P}^2_{\mathbb{Z}}$), but has singular fibers (e.g at the prime 3, this is given by the equation $y^2 = x^3 - x^2$, which is singular at [1:0:0]).
Suppose we wish to throw out these singularities and consider the regular locus $\mathrm{reg}(X)$ (which is the same as the smooth locus over $\mathrm{Spec}(\mathbb{Z})$). Then indeed $\mathrm{reg}(X)$ would be a smooth $\mathbb{Z}$-model for $X_{\mathbb{Q}}$, but it is no longer proper! If it were, then separatedness of $X \to \mathrm{Spec}(\mathbb{Z})$ would imply that the inclusion $\mathrm{reg}(X) \to X$ were itself proper, hence closed. But it is also a dense open subscheme, so that would imply $\mathrm{reg}(X) = X$, a contradiction.

Intuitively, removing or blowing up singularities ruins properness. The fix that Néron came up with was to massage the properness condition to allow for canonically constructible models with properties that were *nice enough*.

Definition 8.1.1.5. Let X_K be a smooth, separated, finite type K-scheme. A **Néron model** of X_K is an S-model X which is smooth, separated, of finite type, and satisfies the *Néron mapping property*:

(*) For each smooth S-scheme Y and for each K-morphism $u_K : Y_K \to X_K$, there is a unique morphism $u : Y \to X$ extending u_K.

Sanity Check. Verify that X_K can *never* be a Néron model of itself – one of the conditions in the above definition is never satisfied by a non-empty K-scheme.

It is difficult to give an explicit, simple example of a Néron model. The canonical motivating construction to keep in mind is that of abelian schemes.

Example 8.1.1.6 (Abelian schemes). Recall an **abelian scheme** over S is a smooth proper S-group scheme with connected geometric fibers. It is indeed true that any abelian S-scheme is a Néron model of its generic fiber, which we will prove in §8.1.5.

Observation (Unnecessary but Interesting). One can define a Néron model of X_K to be the scheme representing the functor

$$Y \mapsto \mathrm{Hom}_K(Y_K, X_K)$$

on the category of smooth S-schemes. Of course, the definition we gave in 8.1.1.5 demands more than just representability of this functor; we want the representing scheme to be separated and of finite type (smoothness is automatic due to the category we are working in). In particular *a Néron model for X_K may not exist even if the above functor is representable.*

The Néron mapping property, in addition to imposing a universal property on X that ensures it is "minimal" in some sense, can also be interpreted as a *weaker version of the valuative criterion for properness.* To see this suppose $S = \mathrm{Spec}(R)$ for R a dvr with fraction field K. Let $Y := \mathrm{Spec}(R')$ for R' an étale R-algebra, and let $K' := \mathrm{Frac}(R')$. The canonical map $Y \to S$ is thus automatically étale (hence smooth). We then have the following commutative diagram:

$$\begin{array}{ccccc} Y'_K = \mathrm{Spec}(K') & \longrightarrow & X_{K'} & \longrightarrow & X \\ \downarrow & & & \nearrow & \downarrow \\ Y = \mathrm{Spec}(R') & & \longrightarrow & & S = \mathrm{Spec}(R) \end{array} \tag{8.1.1-I}$$

The Néron mapping property implies that for any such diagram above there is exactly one dotted arrow making the diagram commute (equivalently, the canonical map $X_R(R') \to X_R(K')$ is bijective). But this is precisely the valuative criterion for properness, restricted to étale R-algebras!

We can formulate this condition locally for general S as in the definition below.

Definition 8.1.1.7. Let X be a scheme over Dedekind S. Then we say that X satisfies the **extension property for étale points** at a closed point $s \in S$ if, for each étale-local $\mathcal{O}_{S,s}$-algebra R' with field of fractions K', the canonical map $X(R') \to X_K(K')$ is surjective.

We note here the useful technical lemma that the extension property for étale points can be checked on strict henselizations of the stalks.

LEMMA 8.1.1.8. The extension property for étale points is equivalent to surjectivity of the (canonical!) map

$$X(\mathcal{O}_{S,s}^{sh}) \to X_K(\mathrm{Frac}(\mathcal{O}_{S,s}^{sh}))$$

on strict henselizations.

Proof. Recall the construction of strict henselization of a local ring R is given by a filtered colimit over triplets $(R', \mathfrak{p}, \alpha)$, where R' is an étale R-algebra, $\mathfrak{p} \trianglelefteq R'$ is a prime lying over the maximal $\mathfrak{m} \trianglelefteq R$, and α is some choice of R-algebra map $\alpha : R' \to (R/\mathfrak{m})_{\mathrm{sep}}$. The arrows in this diagram are given by R-algebra local maps. The map in the above lemma is then (uniquely!) given by the universal property of the colimit. Right-exactness of the colimit then tells us that this map is surjective if and only if the corresponding maps between systems are all surjective, which is exactly the extension property for étale points. □

Since each $\mathcal{O}_{S,s}$ is a dvr (by assumption of S being Dedekind), each étale-local $\mathcal{O}_{S,s}$-algebra is a dvr again, as in our setup. Thus the extension property for étale points says there is *at least* one such dotted arrow in Eqn. 8.1.1-I making the diagram commute. Considering the valuative criterion for separatedness of $X \to S$, we also know that there must be *at most* one such dotted arrow making 8.1.1-I commute. Thus combining these two criteria we get precisely bijectivity condition imposed by the Néron mapping property.

Thus a Néron model of X_K satisfies:

1. Smoothness
2. Separatedness
3. Finite-Type-ness
4. The extension property for étale points

In lieu of the discussion preceding this list, one would hope that the above four properties characterize Néron models. Unfortunately, this is not the case.

NON-EXAMPLE 8.1.1.9. Take $S = \mathrm{Spec}(\mathbb{Z}_{(p)})$ for some prime p, then $K = \mathbb{Q}$. Take $X_K = \mathbb{P}^1_{\mathbb{Q}}$. One can almost immediately verify that $X := \mathbb{P}^1_S$ is a smooth, separated, finite type S-model for X_K, and it even satisfies the extension property for étale points. However it is *not* a Néron model for $\mathbb{P}^1_{\mathbb{Q}}$, since not every map of its fiber extends. In particular taking $Y = X$, then a map $Y_K \to X_K$ is an automorphism of $\mathbb{P}^1_{\mathbb{Q}}$, not all of which extend to $\mathbb{P}^1_S$. For example, we may take the map represented by the matrix

$$\begin{bmatrix} p & 0 \\ 0 & 1 \end{bmatrix} \in PGL_2(\mathbb{Q}).$$

The above example motivates the following definition.

Definition 8.1.1.10. A **weak Néron model** (over S) of a smooth, separated, finite-type K-scheme X_K is a smooth, separated, finite-type S-scheme X that satisfies the extension property for étale points.

(Non)-Example 8.1.1.9 tells us $\mathbb{P}^1$ is a weak Néron model. The failure of a weak Néron model to be an "honest" Néron model comes down to the fact that in some sense the weak Néron model is "too large" for certain maps to the generic fiber to extend.

Key Idea. A Néron model of a K-scheme X_K is a smooth, finite type S-model satisfying the valuative criterion of separatedness, the extension property for étale points (which amounts to a milder variant of properness), and "simple" enough so that any morphism to its generic fiber extends.

§8.1.2 Elementary Properties

Throughout this section we adopt the same notation as in the previous one, namely S is a Dedekind scheme with ring of rational functions K.

Here we list and prove some elementary properties of Néron models that follow immediately from the definition, assuming its existence. We remark that one can use Néron models to great effect using their existence and universal property alone, without knowing anything about their actual structure or construction – a common theme in the study of moduli problems. This was actually standard practice until the book [BLR90a] made Néron's original construction more widely accesible to the larger mathematical community.

Proposition 8.1.2.1 (Elementary Properties of Néron models). Let X be an S-scheme which is a Néron model of its generic fiber X_K. Then:

(a) (Uniqueness) X is uniquely determined by X_K, up to (canonical!) isomorphism.

(b) (Fiber Product) If X'_K is another K-scheme with Néron model X' (over S), then $X \times_S X'$ is the Néron model of $X_K \times_K X'_K$.

(c) (Locality on the Base) X is a Néron model for X_K if and only if X_U is a Néron model for its generic fiber (over U) for any open $U \subseteq S$.

(d) (Group Structure) If X_K is a group scheme, then its Néron model X admits a unique S-group scheme structure compatible with X_K.

(e) (Étale Base Change) The formation of Néron models commutes with étale base change; i.e. if $S' \to S$ is an étale morphism of Dedekind schemes, then $X_{S'}$ is a Néron model (over S') of $X_{K'}$ (naturally K' is the ring of rational functions of S').

Proof. Assertion (a) is an immediate consequence of the Néron mapping property (essentially we may view it as the universal property that defines the Néron model as the *simplest* object with this property).

For assertion (b), we first note that X, X' are both smooth, separated and of finite type over S (by assumption of being Néron models of their generic fibers). Thus $X \times_S X'$ is smooth, separated and finite-type over X', and thus over S, since all these properties are preserved under composition. It remains to verify the Néron mapping property, to which end take an S-scheme Y, and some K-morphism $\mu_K : Y_K \to X_K \times_K X'_K$, giving the following commutative diagram:

$$\begin{array}{ccccc} Y_K & \xrightarrow{\mu_K} & X_K \times_K X'_K & \longrightarrow & X \times_S X' \\ \downarrow & & & & \downarrow \\ Y & & \longrightarrow & & S \end{array}$$

Note composing μ_K in the diagram above with the projection maps onto each component of the fiber product yields unique S-morphisms $Y \to X$ and $Y \to X'$ respectively, by virtue of them both being Néron models. Thus by the universal property of the fiber product we get a unique map $Y \to X \times_S X'$, which, by construction, must commute with the diagram above, as desired.

For assertion (c), we recall that smoothness, separatedness, and being of finite-type are all local conditions, so it only remains to verify the Néron mapping property. Supposing X is a Néron model for X_K over S, let $U \subseteq S$ be an open, and let F be the ring of rational functions of U. Note that $F \subseteq K$ (this is visibly true on every connected integral component, and thus true of the direct sum), thus we get a canonical chain of morphisms

$$\mathrm{Spec}(K) \to \mathrm{Spec}(F) \to U \hookrightarrow S$$

Thus given any U-scheme Y (which we may also view as an S-scheme by composing the structure morphism with the canonical inclusion), and an F-morphisms $Y_F \to X_F$, one can base change to obtain a K-morphism $Y_K \to X_K$, from which one can apply the Néron mapping property to obtain an S-morphism $Y \to X$, which we then base-change to get a unique map $Y \to X_U$, extending our F-morphism, as desired.
The reverse direction is slightly less obvious. Suppose we have an open cover $S = \bigcup_i U_i$, and that each X_{U_i} is a Néron model for X_{F_i}; we wish to show that a morphism $Y_K \to X_K$ extends uniquely. Note we have by construction that $\mathrm{Spec}(F_i) = U_i \cap \mathrm{Spec}(K)$, thus we can base change the map $Y_K \to X_K$ to any $Y_{F_i} \to X_{F_i}$. Each such F_i-morphism $Y_{F_i} \to X_{F_i}$ extends uniquely to a U_i-morphism $Y_{U_i} \to X_{U_i}$, which must agree on intersections by uniqueness. Thus we may glue these to obtain a unique morphism $Y \to X$ extending the map on generic fibers, hence verifying the Néron mapping property.

For assertion (d) we start with extending all the relevant morphisms from X_K to X via the Néron mapping property. For example, suppose $m_K : X_K \times_K X_K \to X_K$ is the multiplication morphism, applying the Néron mapping property yields a unique map m making the following diagram commute.

$$\begin{array}{ccccc} X_K \times_K X_K & \xrightarrow{m_K} & X_K & \longrightarrow & X \\ \downarrow & \overset{m}{\nearrow} & & & \downarrow \\ X \times_S X & & \longrightarrow & & S \end{array}$$

The inversion morphism $\iota_K : X_K \to X_K$ and unit morphism $\epsilon_K : \mathrm{Spec}(K) \to X_K$ extend similarly. From here one we show that the relevant diagrams for veirfying group axioms all commute. For example, letting $\pi : X \to S$ denote the structure morphism, we take the inverse diagram:

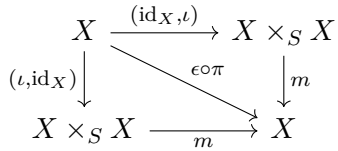

By virtue of X and $X \times_S X$ being Néron models for their respective generic fibers, each of the maps $X \to X$ above (going down then right, right then down, and across the diagonal) are the unique extensions of the corresponding maps on generic fibers, which are all the same map by assumption of X_K being a group scheme. Thus this diagrom above commutes as well, and the other group axioms follow by identical logic.

Finally, for assertion (e), we know (arbitrary) base change preserves morphisms that are smooth, separated and of finite type, so it only remains to verify the Néron mapping property. Take a smooth S'-scheme Y and a K'-morphism $u : Y_{K'} \to X_{K'}$, we then wish to show a unique extension $Y \to X_{S'}$. We know we want to use the Néron mapping property of X, but a priori we cannot – there is no reason for u to be a K-morphism. The key observation here is that *étaleness* of $S' \to S$ ensures that $\mathrm{Spec}(K')$ is the fiber product of $\mathrm{Spec}(K)$ and S. Indeed, any étale map is open, and thus generic points are mapped to generic points; taking the fiber pullback of $\mathrm{Spec}(K)$ would thus yield precisely the scheme of generic points of S'. We thus get the following commutative diagram:

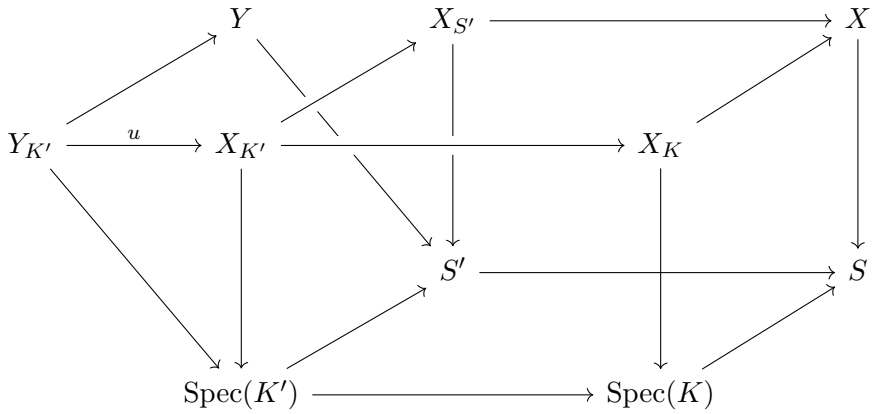

From this diagram above we see that composing the morphism u with the map $X_{K'} \to X_K$ yields a K-morphism. Moreover, since $S' \to S$ is étale, it is smooth, and thus the composition $Y \to S' \to S$ (which is also the structure morphism of Y as an S-scheme) is smooth. Thus we can use the Néron mapping property of X to extend this to an S-morphism $Y \to X$. Then by the universal property of $X_{S'}$ as a fiber product we get a unique S'-morphism $Y \to X_{S'}$, as desired. □

Assertion (c) of the above Proposition states that the condition of being a Néron model can be checked on an open cover, i.e. it is a *local* property. We claim that we can go even *more* local, by examining "stalks" i.e. the schemes $X \times_S \mathcal{O}_{S,s}$ for closed points $s \in S$. The key result is as follows.

> PROPOSITION 8.1.2.2. Let X be an S-scheme of finite type. Then X is a Néron model of its generic fiber if and only if the same is true for $X \times_S \mathrm{Spec}(\mathcal{O}_{S,s})$ for all closed points $s \in S$.

Note that it is in fact not sufficient to check an arbitrary S-scheme X being a Néron model on base change to stalks; one must *a priori* assume the finite-type condition. Moreover, we also remark that to show the forward direction – that stalk base changes of a Néron model are also Néron models – it is insufficient to use the base change property (assertion (e) in Proposition 8.1.2.1), since in general $\mathrm{Spec}(\mathcal{O}_{S,s})$ need not

be étale over S. It is however, constructed as a colimit of open neighborhoods of s, so we will need some technical lemmas on the behavior of such limits to prove the Proposition.

LEMMA 8.1.2.3. Let S be an arbitrary base scheme with closed point s, and let X, Y be finitely presented S-schemes.

(a) The canonical map

$$\varinjlim_{\text{open } U \ni s} \mathrm{Hom}_U(X_U, Y_U) \to \mathrm{Hom}_{\mathcal{O}_{S,s}}(X_{\mathcal{O}_{S,s}}, Y_{\mathcal{O}_{S,s}})$$

is bijective.

(b) Let $X^{(s)}$ be a finitely presented $\mathcal{O}_{S,s}$ scheme. Then there is an open neighborhood $U \ni S$ and a U-scheme X of finite presentation such that $X_{\mathcal{O}_{S,s}} \cong X^{(s)}$.

(c) Let $f : X \to Y$ be an S-morphism. Then the morphism

$$f_{\mathcal{O}_{S,s}} : X_{\mathcal{O}_{S,s}} \to Y_{\mathcal{O}_{S,s}}$$

has property P if and only if there exists some open $U \ni s$ such that the base change $f_U : X_U \to Y_U$ has property P for all of the following properties: isomorphism, immersion (open or closed), separated, surjective, (quasi)-affine, (quasi)-finite, proper.

Proof. See [EGA, §IV, 8.8.2] for (a) and (b), and [EGA, §IV, 8.10.5] for (c). The proof is essentially that those theorem statements applied to the inverse system

$$\{(X_U, Y_U, U)\}_{\text{open } U \ni s},$$

whose inverse limit is precisely $(X_{\mathcal{O}_{S,s}}, Y_{\mathcal{O}_{S,s}}, \mathrm{Spec}(\mathcal{O}_{S,s}))$. □

Before proceeding to the proof of the Proposition we remark that

Proof of 8.1.2.2. For one direction, suppose X is a Néron model of its generic fiber. Pick an arbitrary closed point $s \in S$, and denote

$$X^{(s)} := X \otimes_S \mathrm{Spec}(\mathcal{O}_{S,s}).$$

We know smoothness, separatedness and being of finite type are preserved under base change, so it remains to verify the Néron mapping property. Let $Y^{(s)}$ be a smooth $\mathcal{O}_{S,s}$-scheme, and let $L := Frac(\mathcal{O}_{S,s})$, and take an L-morphism

$$f_L : Y^{(s)}_L \to X^{(s)}_L = X_L$$

Note being smooth implies $Y^{(s)}$ is of finite presentation over $\mathrm{Spec}(\mathcal{O}_{S,s})$. By (b) of Lemma 8.1.2.3 and locality of smoothness, there exists some open neighborhood $U \ni s$ and a smooth finitely presented U-scheme Y whose fiber over s is $Y^{(s)}$. WLOG we may assume U is integral, so that $\mathscr{K}(U) = L$.
By (c) of Proposition 8.1.2.1, we know that X_U is a Néron model of its generic fiber $X_L = X^{(s)}_L$, so we have a unique map

$$f_U : Y \to X_U$$

extending f_L. But then note that

$$f^{(s)} := f_U \otimes_U \mathrm{Spec}(\mathcal{O}_{S,s}) : Y^{(s)} \to X^{(s)}$$

is a $\mathcal{O}_{S,s}$-morphism extending f_L. We claim this morphism must be unique. If not, then by bijectivity of the sets in part (a) of Lemma 8.1.2.3, there must be some open set $U' \ni s$ such that f_L extends to $f_{U'} : Y_{U'} \to X_{U'}$ in two distinct ways; a contradiction. Thus the Néron mapping property holds.

For the other direction, suppose the base change of X to the stalk of any closed point of S is a Néron model, we wish to show that X itself is a Néron model. For any closed point s, we have that the structure morphism

$$X_{\mathcal{O}_{S,s}} \to \mathrm{Spec}(\mathcal{O}_{S,s})$$

is smooth, separated and of finite type. By part (c) of Lemma 8.1.2.3, for each such s there must exist an open neighborhood $U_s \ni s$ so that the base changed structure morphism

$$X_{U_s} \to U_s$$

is smooth, separated and of finite type (technically the lemma only gives us the latter two, the smoothness follows simply by definition). Since the $\{U_s\}_{\text{closed } s \in S}$ form an open cover of S, it follows that $X \to S$ is smooth, separated and *locally* of finite type (hence why we needed to assume X was a priori finite type). Thus it only remains to verify the Néron mapping property. But since we've already shown in Proposition 8.1.2.1 that this property is also local on the base we may WLOG assume that S is integral. Note in this case then that X_K, Y_K are also the generic fibers of the base changes of X, Y to $\mathrm{Spec}(\mathcal{O}_{S,s})$. Thus any morphism

$$f_K : X_K \to Y_K$$

extends uniquely to $f^{(s)} : X_{\mathcal{O}_{S,s}} \to Y_{\mathcal{O}_{S,s}}$, and hence (by bijectivity of the map of Homsets in part (a) in Lemma 8.1.2.3) uniquely to a morphism on an open neighborhood $U_s \ni S$. Gluing all such morphisms together – note they must agree on intersections by uniqueness – we get a unique morphism $f : X \to Y$ extending f_K, as desired. This completes the other direction. □

The above result motivates the following definitions.

> DEFINITION 8.1.2.4. Given an S-scheme X which is a Néron model of its generic fiber X_K, we call the schemes $X \times_S \mathrm{Spec}(\mathcal{O}_{S,s})$ (for any closed point $s \in S$) **local Néron models** for X_K. Thus we may also view X itself as a **global Néron model**.

Proposition 8.1.2.2 says that if a K-scheme X_K admits a global Néron model, then it must also admit all its local ones. However, the converse need not be true, as we will see [cite when].

§8.1.3 Local Existence Prelude: Boundedness

We want to characterize the conditions under which *local* Néron models of a K-scheme X_K exist. Since we are working over Dedekind schemes, this fundamentally amounts to asking when Néron models exist for the spectrum of a DVR. A key element of that theory is the notion of **boundedness**, so we expound upon this here. Throughout this section we let R denote a DVR with field of fractions K.

Spoiler. It will turn out (Theorem 8.1.4.1) that X_K admits a Néron model exactly when the rational points over K^{sh} (the field of fractions of the strict henselization of R) is bounded.

If we have a faithfully flat extension of DVRs $R \subseteq R'$, we get coinciding absolute values on their corresponding field of fractions $K \subseteq K'$, which we simply denote by $|\cdot|$. Let X_K be any K-scheme, and $x \in X_K(K')$ be any K'-rational point. For any global section $s \in \mathcal{O}_{X_K}$, recall $s(x)$ (the value of s at x) is an element of $\kappa(x) = K'$. We may thus speak of the absolute value $|s(x)|$. Our goal is to give a rigorous definition of what we intuitively expect to be a bounded subset of the vector space K^n, for example.

Definition 8.1.3.1. With the notation as above, let X_K be finite type K-scheme. We define the notion of a subset $E \subseteq X_K(K')$ being **bounded in** X_K (relative to R', K') as follows:

(a) If $X_K = \mathbb{A}^n_K$, then E is bounded if and only if the set

$$\left\{ |g(x)| : g \in \mathcal{O}_{\mathbb{A}^n_K}, x \in E \right\} \subseteq \mathbb{R}$$

is bounded.

(b) If X_K is some affine scheme, E is bounded if and only if there exists some closed immersion $X_K \hookrightarrow \mathbb{A}^n_K$ mapping E onto a bounded subset, in the sense of (a).

(c) In general, E is bounded if and only if there exists a finite covering of X_K by affine open subschemes $U_1, \dots U_m$ and a decomposotion $E = \bigcup_{i \in [m]} E_i$ so that $E_i \subseteq U_i(K')$ is bounded, in the sense of (b).

Sanity Check. Verify that if $E \subseteq X_K(K')$ is bounded then so is any $E' \subseteq E$.

Observation. For affine $X_K = \mathrm{Spec}(A)$, boundedness amounts to saying there exist elements $g_1, \dots g_n \in A$ generating A as a K-algebra which, when viewed as maps $X(K') \to K'$, are bounded on E. In particular the condition in (b) is independent of choice of closed immersion.

Slightly less obvious is that the condition in (c) is independent of choice of open cover $\{U_i\}$. We will discuss this right after giving some key examples. In particular we want to emphasize that the notion of boundedness is always *relative* to the ambient scheme.

Example 8.1.3.2 ((Un)Bounded subsets). Suppose $K' = K$. Firstly consider $\mathbb{G}_a = \mathbb{A}^1_K = \mathrm{Spec}(K[t])$. Let $\pi \in R$ be a non-unit (i.e. a power of a uniformizer), and consider the subset

$$E := \{\langle t - \pi^m \rangle : m \in \mathbb{N}\} = \{\pi^m\}_{m \geq 0} \subseteq \mathbb{G}_a(K) = K$$

Writing any $s \in K[t]$ as $s_n t^n + \dots + s_1 t + s_0$ for $s_i \in K$, we have

$$\begin{aligned} |s(\langle t - \pi^m \rangle)| &= |s_n \pi^{nm} + \dots + s_1 \pi^m + s_0| \\ &\leq |s_n| + \dots + |s_1| + |s_0| \ \ \forall m \geq 0. \end{aligned}$$

Thus it follows E is bounded. This aligns with our classical intution: all the points in E may be viewed as elements of the unit ball in K.

Now suppose we define the "same subset" E, but of $\mathbb{G}_m(K)$ instead now. In particular we write

$$E := \{\langle t - \pi^m \rangle : m \in \mathbb{N}\} = \{\pi^m\}_{m \geq 0} \subseteq \mathbb{G}_m(K) = K^\times.$$

Intuitively one might expect this to be bounded as well, but it is not! Considering the closed immersion

$$\begin{aligned} \mathbb{G}_m = \operatorname{Spec}(K[t, t^{-1}]) &\xrightarrow{\sim} Z(xy - 1) \subseteq \mathbb{A}^2_K \\ (t, t^{-1}) &\leftarrow (x, y) \end{aligned}$$

We note the global section $y \in \mathcal{O}_{\mathbb{A}^2_K}$ does not take bounded values on E, since $\left|(\pi^m)^{-1}\right| = |\pi^m|^{-1}$ grows arbitrarily large for large m. Since boundedness is independent of choice of closed immersion into affine space, it follows E is unbounded.

The takeaway here is that *boundedness is highly dependent on the ambient scheme*, and cannot be judged on some "intrinsic" level of the subset E.

We now deal with the well-definedness of part (c) of Definition 8.1.3.1.

LEMMA 8.1.3.3. Take a subset $E \subseteq X_K(K')$, and suppose there exists an affine covering $\{U_i\}_{i \in [n]}$ of X_K satisfying condition (c) of 8.1.3.1. Let $\{V_j\}_{j \in [m]}$ be another affine open cover, then there is a partition $E = \bigcup_{j \in [m]} E'_j$ into subsets $E'_j \subseteq V_j(K')$ such that each E'_j is bounded in V_j.

Proof. Since conditions (b) and (c) of Definition 8.1.3.1 are visibly equivalent for any affine scheme, it follows that we may (by replacing V_j with affine opens covering each intersection $V_j \cap U_i$) that $\{V_j\}_{j \in [m]}$ is a refinement of the cover $\{U_i\}_{i \in [n]}$.

By reordering if necessary, we may suppose U_1 is covered by (and equal to the union of) $V_1, \dots V_k$. WLOG we may also assume each of the V_i is a distinguished affine open; i.e. if $U_i = \operatorname{Spec}(A)$, then $V_j = \operatorname{Spec}(A_{f_j})$, for some $f_j \in A$ such that $f_1, \dots f_k$ generate A. In particular there must exist $a_1, \dots a_k \in A$ so that

$$\sum_{j \in [k]} a_j f_j = 1 \in A. \tag{8.1.3-I}$$

Now take $E_1 := E \cap U_1(K')$. Note that this above equality implies that

$$\epsilon := \inf \{\max \{|f_j(x)| : j \in [k]\} : x \in E_1\} > 0.$$

If this weren't the case, then we would have

$$\left| \left(\sum_{j \in [k]} a_j f_j \right)(x) \right| \leq \sum_{j \in [k]} |a_j(x)|\,|f_j(x)| < 1 = 1(x)$$

for some $x \in E_i$, contradicting Eqn. 8.1.3-I. Thus we may define

$$F_j := \{x \in E_1 : |f_j(x)| \geq \epsilon\}.$$

For any element $a/(f_j)^\ell \in \mathcal{O}_{V_j} = \operatorname{Spec}(A_{f_j})$ and any $x \in F_j$, note we have

$$\left|(a/f_j^\ell)(x)\right| = |a(x)| \cdot |f_j(x)|^{-\ell} \leq |a(x)|\,\epsilon^{-\ell}.$$

By assumption of E_1 being bounded in $U_1 = \operatorname{Spec}(A)$, it follows that $|a(x)|$ is bounded as x ranges over $E_i \cap V_j(K') = F_j$ for all $a \in A$, and thus the same is true for any $a/(f_j)^\ell \in \mathcal{O}_{V_j}$. Thus each F_j defined is a

bounded subset of V_j, and moreover by construction we have

$$E_1 = \bigcup_{j \in [k]} F_j.$$

Repeating this construction for all $E_i \subseteq U_i(K')$, we get the desired result. □

Now that we have settled the definition of boundedness, we log some immediate consequences of the definition.

PROPOSITION 8.1.3.4. Let R'/R be a faithfully flat extension of DVRs with field of fractions K'/K, and let X_K be a finite-type K-scheme. Let $E \subseteq X_K(K')$ be a bounded subset.

(a) Boundedness of E is equivalent to boundedness of its image (via precomposing by the morphism $\mathrm{Spec}(K') \to \mathrm{Spec}(K)$) $E' \subseteq X_{K'}(K')$.

(b) For any finite-type K-scheme Y_K and a K-morphism $X_K \to Y_K$, the image of a bounded subset $E \subseteq X_K(K')$ under the canonical map $X_K(K') \to Y_K(K')$ is also bounded.

Proof. We verify these assertions when X_K, Y_K are affine, and the general case follows by using Lemma 8.1.3.3 to reduce to the affine case.

We start with assertion (b). Let $X_K = \mathrm{Spec}(A), Y_K = \mathrm{Spec}(B)$, and suppose the morphism $\phi : X_K \to Y_K$ is determined by a K-algebra map $\phi^\# : B \to A$. Take $x \in E \subseteq X_K(K')$, determined by a map $A \to K'$. Its image in $Y_K(K')$ is defined by precomposing the map $A \to K'$ with $\phi^\#$. Thus for any $b \in B$ and $y := \phi(x)$, we have

$$b(y) = \phi^\#(b)(x) \in K'$$

It then immediately follows (from the characterization of boundedness in the Observation following definition 8.1.3.1) that boundedness of E implies boundedness of its image in $Y_K(K')$, proving assertion (b).

The reverse direction of assertion (a) is a consequence of (b), so it remains to show that if $E \subseteq X_K(K')$ is bounded then so is the corresponding $E' \subseteq X_{K'}(K')$. But again in the affine case if $X_K = \mathrm{Spec}(A)$ then $X_{K'} = \mathrm{Spec}(A \otimes_K K')$, and thus boundedness on X_K must imply boundedness on $X_{K'}$. □

As one might intuitively expect, boundedness is implied by properness, in the sense that any subset of the rational points of a proper scheme is bounded.

PROPOSITION 8.1.3.5. If X_K is any proper K-scheme, then any $E \subseteq X_K(K')$ is bounded in X_K.

Proof. By Chow's Lemma ([Stacks, §30.18]), X_K admits a proper surjective K-morphism onto it by a projective scheme, so by Proposition 8.1.3.4 we may reduce to the projective case. In fact by applying the Proposition again we may assume $X = \mathbb{P}^n_K$.

Covering $\mathbb{P}^n_K$ with the standard affine charts (i.e by making some choice of projective coordinates $[x_0 : x_1, \ldots : x_n]$ and setting $U_i \cong \mathbb{A}^n_K$ to be the set of points where the i-th coordinate is non-vanishing). We may write K'-rational points as $x = [x_0 : \ldots : x_n]$, with each $x_i \in K'$. Then consider

$$E_i := \{[x_0 : \ldots : x_n] : |x_i| = \max(|x_0|, \ldots |x_n|)\} \subseteq \mathbb{P}^n_K(K').$$

Note that $E_i \subseteq U_i(K')$, and moreover $\bigcup_{i=0}^n E_i = \mathbb{P}^n_K(K')$, so it suffices to show each E_i is bounded in U_i. But, as a subset of $U_i(K')$, note that

$$E_i = \{(x_1, \dots x_n) \in \mathbb{A}^n_K(K') : |x_j| \leq 1 \ \forall j\},$$

which is visibly bounded. Thus $\bigcup_i E_i = \mathbb{P}^n_K(K')$ is bounded, as desired. □

We end this subsection by describing a more non-trivial criterion for boundedness (which we won't end up using for our purposes but is interesting nonetheless). Before we get to the Proposition statement we will need a definition. Recall that any separated finite-type K-scheme admits a (non-canonical!) compactification $\overline{X_K}$ (see §38.33 in [Stacks]), i.e. a proper K-scheme that has X_K as a dense open subscheme.

DEFINITION 8.1.3.6. Let X_K be a separated finite-type K-scheme. We say X_K has **no rational point at infinity** if there exists a compactification $\overline{X_K}$ with $(\overline{X_K} - X_K)(K) = \emptyset$.

EXAMPLE 8.1.3.7. Consider the scheme

$$Z := \mathrm{Proj}(\mathbb{Q}[x, y]/\left\langle x^2 + y^2 \right\rangle) \subseteq \mathbb{P}^1_{\mathbb{Q}},$$

i.e. the closed subscheme cut out by x^2+y^2. Note this is a non-empty scheme which has no points over $\mathbb{Q}$. We may then set $X := \mathbb{P}^1_{\mathbb{Q}} - Z$, a dense open subscheme, therefore separated and of finite-type. We thus have

$$(\overline{X} - X)(\mathbb{Q}) = (\mathbb{P}^1_{\mathbb{Q}} - X)(\mathbb{Q}) = Z(\mathbb{Q}) = \emptyset,$$

and thus we see that X is a scheme with no rational point at infinity.

PROPOSITION 8.1.3.8. Let R be an excellent henselian DVR with field of fractions K, and suppose X_K is a finite-type K-scheme with no rational point at infinity, and let $\overline{X_K}$ be some compactification for which this is the case. Then if $E \subseteq X_K(K)$ is bounded in $\overline{X_K}$, it is also bounded in X_K.

This result can be seen as a converse to assertion (b) of Proposition 8.1.3.4, granting certain conditions on R and X_K. Note again that boundedness in $\overline{X_K}$ does not trivially imply boundedness in X_K simply because it is a subscheme; we saw in Example 8.1.3.2 that this need not be the case.

Proof. [later] □

COROLLARY 8.1.3.9. Let R be an excellent DVR with fraction field K, and let X_K be a separated K-scheme of finite-type with no rational point at infinity. Then $X_K(K)$ is bounded (in X_K).

Proof. Combine Propositions 8.1.3.8 and 8.1.3.5. □

§8.1.4 (An Outline of) Local Existence

We've seen in the previous section that the existence of a global Néron model over a Dedekind scheme S implies the existence of local ones at all closed points of S, which are all spectra of fields or DVRs. In this

section we will discuss the criteria for the existence of Néron models over such schemes. Throughout this section we let R denote a DVR with field of fractions K. We denote by R^{sh} and K^{sh} the strict henselization of R and its field of fractions, respectively.

Spoiler. We will see later (Theorem 8.1.5.4) that we construct global Néron models – without too much effort – by gluing local ones together!

Theorem 8.1.4.1 (Local Existence of Néron Models). Let R be a DVR with fraction field K. Let X_K be a finite type smooth K-group scheme. Then X_K admits a Néron model X over R if and only if $X_K(K^{sh})$ is bounded in X_K.

Corollary 8.1.4.2. Any abelian variety A_K over K (fraction field of a DVR R) admits a Néron model over R.

Proof. By definition, A_K is a smooth and finite type K-group scheme. Moreover, it is proper, hence $A_K(K^{sh})$ is bounded, by Proposition 8.1.3.5. Thus it satisfies the conditions of the above theorem, yielding the desired result. □

The proof of Theorem 8.1.4.1 will be constructive, and will take up multiple future chapters. In this section, we give a brief outline of the proof strategy. Firstly, we want to see why boundedness of the set of K^{sh}-rational points is a *necessary* condition for a Néron model of *any* finite-type K-scheme to exist; this is rather easy to show with our current tools. Recall if X_K is the generic fiber of the R-scheme X, then for any (faithfully flat) extension R'/R with corresponding K'/K, we have a canonical map

$$X(R') \to X_K(K').$$

Proposition 8.1.4.3. Let X_K be a finite-type K-scheme. Then $E \subseteq X_K(K')$ is bounded if and only if there exists a finite-type R-model X of X_K such that the image of the canonical map $X(R') \to X_K(K')$ contains E.

Proof. For one direction, let E be a bounded subset. Covering X_K with affines $\{U_{i,K}\}$, if each affine admits an R-model U_i such that the image of $U_i(R') \to U_{i,K}(K')$ contains $E_i := E \cap U_{i,K}(K')$, then gluing U_i together along the generic fiber will yield an R-model X of X_K whose image of R' points contains E. Conversely, if such an R-model exists, then we get affine-local R-models U_i. Thus we reduce to the case where X_K is affine.

For a closed subscheme $X_K \subseteq \mathbb{A}^n_K$, we may take X to be its schematic closure in $\mathbb{A}^n_R$; then it is clear that X is an R-model of X_K. Moreover then the image of the canonical map

$$X(R') \to X_K(K') \subseteq \mathbb{A}^n_K(K')$$

would consists of points whose coordinates lie in R', i.e whose absolute values our bounded by 1 (geometrically this can be thought of as a cube in affine space). Thus any set of K'-rational points E that lies within this image is visibly bounded.

Conversely, given any such bounded E, we can always scale the coordinate functions on $\mathbb{A}^n_K$ by some suitable constants so that the value they take on any point of E is constrained by the values taken on in the image of $X(R')$. □

COROLLARY 8.1.4.4. For X_K a finite-type K-scheme, there exists a finite-type R-model X of X_K satisfying the extension property of étale points if and only if $X_K(K^{sh})$ is bounded. In particular, boundedness is *necessary* for existence of a Néron model of X_K.

Proof. Recall Lemma 8.1.1.8 tells us that the extension property for étale points can be checked on the strict henselization, i.e. it is equivalent to surjectivity of $X(R^{sh}) \to X_K(K^{sh})$. The result then follows from the above Proposition, setting $E := X_K(K^{sh})$. □

WARNING. Although the construction in the proof of 8.1.4.3 yields a finite-type R-model satisfying the integral property for étale points, in general this need not be smooth nor separated.

EXAMPLE 8.1.4.5. Suppose $X_K = \mathbb{P}^1_K$, and we can take our affine cover to be the standard charts; i.e covered by two copies of $\mathbb{A}^1_K$. Our procedure in the proof of Proposition 8.1.4.3 results in an R-model that is two copies of $\mathbb{A}^1_R$ glued along the generic fiber. Though this is still smooth, we claim it is not separated. We'll verify this via the valuative criterion for separatedness using R and K themselves.

Picking any section of the structure morphism $\mathbb{A}^1_R \to \mathrm{Spec}(R)$, we note the generic point is sent to the generic fiber and the special point to the special fiber. Thus any K-valued point in the generic fiber $\mathbb{A}^1_K \subseteq X$ can be extended to an R-valued point two different ways; we can just switch the copy of $\mathbb{A}^1_R$ we are looking at! This violates the valuative criterion, as desired.

Pictorially, the image is the following.

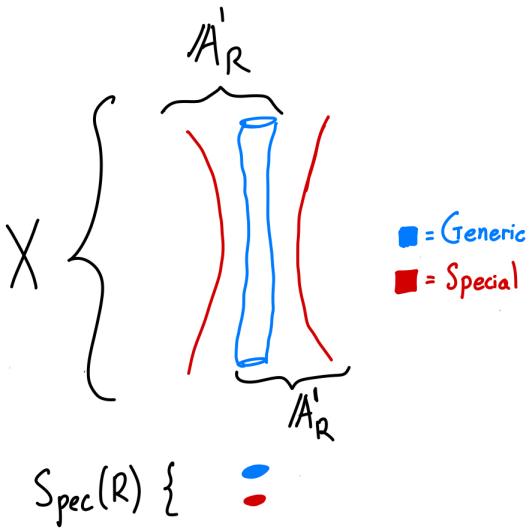

Figure **8.1.4-A**: Two affine lines glued along the generic fiber fail to be separated, much akin to the classic "line with doubled origin" example.

The above example shows a failure of separatedness after gluing; in general there is no reason why the gluing should be smooth, as the generic fibers need not be an open subschemes of the various U_i.

Construction Blueprint of Theorem 8.1.4.1

Our current situation may be summarized as follows: we have a smooth finite-type group scheme X_K over the field of fractions of a dvr R such that $X_K(K^{sh})$ is bounded. Corollary 8.1.4.4 gives us a finite-type R-model satisfying the extension property for étale points, let this be X_0. From here, we wish to arrive at a Néron model for X_K. This will require considerable more work.

Our strategy can be summarized in four main steps.

(I) We break down the construction of X_0 (recounted in the proof of Proposition 8.1.4.3 and build it up in a new way to obtain a *separated* R-model X_1 satisfying the extension property for étale points, by the end of §8.2.1. This will require some highly technnical lemmas on "flattening by blowups" procedure, the background of which we expound upon in Appendix §C.2. However in general this X_1 may not be smooth.

(II) To our X_1 we apply a **smoothening process** – a series of blowups done in such a manner so as to preserve the extension property for étale points – to obtain, by the end of §8.2.3, a smooth separated X_2, which will in fact be a *weak* Néron model (recall Definition 8.1.1.10). This will satisfy something known as the **weak Néron mapping property**, which we discuss in §8.2.4. Such models are not unique, nor do they admit an extension of the group structure of X_K, so further work will be required.

(III) To modify X_2 to allow an extension of the group structure of X_K, we throw out some so-called *minimal components* of the special fiber of X_3. Doing so in a precise enough way will yield (in the middle of §8.3.2) X_3. With this smaller model, the multiplication map

$$m_K : X_K \times X_K \to X_K$$

will yield an extension to an R-birational map

$$m : X_3 \times X_3 \to X_3,$$

which will define an **R-birational group law**. This X_3 will already be a dense open subscheme of our final Néron model. However, in throwing out those minimal components we may have lost our extension property for étale points, so X_3 will be neither an abelian scheme nor a weak Néron model.

(IV) To turn X_3 into an honest Néron model, we take its "saturation under the R-birational group law". In short we use a theorem of Weil to enlarge X_3 into a smooth, separated, finite-type R-scheme X_4 with a legitimate R-group law extending the group structure of X_K. Finally, we check that X_4 actually satisfies the Néron mapping property (spoilers: it will). We will complete the verification of X_4 as a Néron model by the end of §8.3.2, and in the following subsection §8.3.3 we provide the proof of the key theorem of Weil used in our construction.

After these steps, our task will be complete, and we can finally rest and watch the sun rise on a grateful universe.

§8.1.5 Local Existence Postlude: Global Existence

If we are to construct a global Néron model, we have to find a way to glue local Néron models. The issue, as demonstrated by the necessity of imposing the finite-type condition in Proposition 8.1.2.2, is that the resulting global model may not be of finite type, and in general this can be difficult to control. As it turns out though, things work nicely when dealing with abelian varieties, and we will be able to relatively easily prove existence of global Néron models for them, assuming the existence of local ones. Throughout this section, S denotes a Dedekind scheme with ring of rational functions K.

To begin we describe a relationship between global and local Néron modesl that will be useful for us.

PROPOSITION 8.1.5.1. Let X_K be a smooth, separated finite-type K-scheme. Then X_K admits a Néron model X (over S) if and only if the following condition is satisfied:

(*) There exists a dense open $S' \subseteq S$ such that X_K admits a Néron model X' over S' and local Néron models at the (finitely many!) closed points of $S - S'$.

Proof. If X_K has a global Néron model X over S, then condition (*) is an immediate consequence of Proposition 8.1.2.2 and assertion (c) of Proposition 8.1.2.1 (locality on the base).

For the reverse direction, suppose (*) holds true. We may again use locality on the base to reduce to assuming that S is integral. Let $s_1, \dots s_n$ be the finitely many closed points in the complement of the dense $S' \subseteq S$, and let X_{s_i} be a local Néron model of X_K over each $\mathcal{O}_{S,s_i}$. By applying assertion (b) of Lemma 8.1.2.3, each of these spread out to some Néron model X_i over a neighborhood $S_i \ni s_i$. Moreover, since X_i and X' coincide over the generic point of S, they must coincide over a nonempty open part of S'. Thus by shrinking further if necessary we may assume

$$S_i \cap (S - S') = \{s_i\}, \quad X_i|_{S_i \cap S'} = X'|_{S_i \cap S}.$$

Thus we may glue each of the X_i to X' along $S \cap S'$ to obtain a a smooth, separated, finite-type S-model X of X_K such that the base change to all stalks is a Néron model of X_K. Thus X is a Néron model of its generic fiber, by Proposition 8.1.2.2. □

DEFINITION 8.1.5.2. Let S be a connected Dedekind scheme with function field K, and let A_K be an abelian variety over K. We say A_K has **good reduction** at a closed point $s \in S$ if A_K extends to a smooth and proper scheme $A_{(s)}$ over $\mathcal{O}_{S,s}$. Otherwise we say s has **bad reduction**.

PROPOSITION 8.1.5.3. Let S be a connected Dedekind scheme with function field K, and let A_K be an abelian variety over K. Suppose A_K extends to a smooth and proper S-scheme A. Then one can endow A with a group structure extending the one of A_K to make it an abelian scheme. In particular A is a Néron model of A_K.

Proof. By Proposition 8.1.2.2 (which we can apply since A is proper!), we may check this on stalks, and thus we may reduce to the case where S is a DVR. Moreover the valuative criterion of properness implies that A is already a weak Néron model of A_K.

Assuming the local existence theorem 8.1.4.1, we know A_K admits a Néron model X. By the Néron mapping property, there exists a unique S-morphism

$$A \to X$$

extending the identity map $A_K \xrightarrow{\sim} A_K$. This morphism must be an open immersion by [cite], and has closed image by properness of A. Moreover since X, S are finite type Noetherian schemes and $X \to S$ is flat (since it is smooth), it follows that the structure morphism is open ([Har77, Ex. 9.1]). This, combined with connectedness of the geometric fiber $X_K = A_K$, implies X is connected. Thus the map $A \to X$ is in fact an isomorphism, so A is the Néron model of A_K. Then assertion (d) of Proposition 8.1.2.1 tells us the group structure of A_K extends to one on A, and thus A is an abelian scheme. □

We can now give a proof of a more general case of our Main Theorem 8.0.0.1.

THEOREM 8.1.5.4. Let S be a connected Dedekind scheme with function field K, and let A_K be an abelian variety over K. Then A_K admits a Néron model A over S. Morever, the set of all closed points of S over which A_K admits good reduction, along with the generic point, form a dense open subscheme $S' \subseteq S$, and $A_{S'}$ is an abelian scheme over S'.

Proof. By definition, A_K extends to a smooth proper scheme $A_{(s)}$ over $\mathcal{O}_{S,s}$ for every $s \in S$ that is the generic point or a point of good reduction. By applying parts (b) and (c) of Lemma 8.1.2.3, each of these spread out to a scheme A_s that is smooth and proper over some neighborhood U_s of each s. By Proposition 8.1.5.3, each of these can be given the structure of an abelian scheme and a Néron model of A_K over U_s. Then we can glue them together to obtain an abelian scheme and Néron model A' over S', where

$$S' = \bigcup_{s \text{ generic or good reduction}} U_s \subseteq_{\text{dense, open}} S.$$

By the local existence result (Theorem 8.1.4.1) we also know A_K admits local Néron models over each of the closed points not of good reduction. Thus all the conditions of Proposition 8.1.5.1 are ssatisfied, and thus A_K admits a Néron model over S. □

Section 8.2

Weak Néron Models (Construction Steps I-II)

Throughout this section, R denotes a dvr with fraction field K, uniformizer π and residue field $R/\pi = k$. In the previous section, we defined the notion of a Néron model of a smooth, separated, finite-type K-scheme X_K and reduced the existence theorem for global Néron models to that of existence of local ones (i.e over the spectrum of a DVR with field of fractions K). In §8.1.4, Corollary 8.1.4.4 told us that boundedness of $X_K(K^{sh})$ was equivalent to the existence of a finite-type R-model X satisfying the extension property for étale points. Following that, we gave an outline for the proof of local existence in a series of four main steps. This section covers the first two steps in that outline, which essentially takes our R-model and "smoothens" and "separates" it to yield a weak Néron model for X_K (recall Definition 8.1.1.10).

The main tool we will use for both the separating and smoothening processes will be blowups, and lots of them. For the readers convenience we include a a review of some basic definitions, computations and properties in Appendix §A.3. We remark here that in our exposition we first make the model separated and then make it smooth, but the specific order doesn't matter, since taking blowups (along with the other operations we will perform) don't ruin smoothness nor separatedness.

Section Outline

- §8.2.1 will produce a separated finite-type R-model satisfying the extension property for étale points.
- §8.2.2 and 8.2.3 give an outline of the "smoothening" technique that will produce make our separated finite-type R-model above smooth, thereby yielding a weak Néron model.
- §8.2.4 reviews the notion of relative rational morphisms in order to provide a characterization of weak Néron models in terms of a so-called *weak Néron mapping property*, which is like an relatively rational version of the Néron mapping property. It is this key property that will be used in the subsequent section on Birational Group Laws (§[cite]) to show that, in the case that X_K is a group scheme, we can extend it to a Néron model.

§8.2.1 Separatedness

We saw in Example 8.1.4.5 that the R-model of X_K obtained in the proof of Proposition 8.1.4.3 need not be separated at all. In this subsection we rectify that and give a procedure to obtain a separated R-model X.

The first step is to see that w can still extract something useful from the proof: a finite set of separated R-models of X_K that, *collectively*, satisfy the extension property for étale points. Before that, however, we recall here a useful criterion for determining separatedness of glued schemes.

> Lemma 8.2.1.1. Let S be an arbitrary base scheme, and let X, Y be separated S-schemes with a common open subscheme $U \hookrightarrow X, Y$. Then the gluing $Z := X \sqcup_U Y$ is separated if and only if the

diagonal morphism

$$\Delta : U \to X \times_S Y$$

is a closed immersion.

Proof. Take the diagonal

$$\Delta : Z \to Z \times_S Z.$$

Note the product can be covered by $X \times_S X, Y \times_S Y, Y \times_S X$ and $X \times_S Y$, call these subsets $V_{XX}, V_{YY}, V_{YX}, V_{XY}$ respectively. Then to show separatedness of Z it is necessary and sufficient to check that each of the maps

$$\Delta^{-1}(V_{IJ}) \to V_{IJ}, \quad I, J \in \{X, Y\}$$

is a closed immersion. For V_{XX} and V_{YY} this follows from separatedness of X and Y respectively, and for V_{XY}, V_{YX}, checking this is precisely the condition stimulated in the lemma statement. □

PROPOSITION 8.2.1.2. For X_K a separated finite-type K-scheme, $X_K(K^{sh})$ is bounded if and only if there exists a finite set of separated finite-type R-models $\{X_i\}_{i\in[n]}$ such that the map

$$\bigsqcup_{i\in[n]} X_i(R^{sh}) \to X_K(K^{sh})$$

is a surjection.

Proof. When X_K is affine, recall in the proof of Proposition 8.1.4.3 that the schematic closure of X_K in $\mathbb{A}^n_R$ is an R-model X. Moreover, as a closed subscheme of affine space, X is R-separated. This model satisfying the extension property for étale points is then equivalent to boundedness of $X_K(K^{sh})$, by the arguments in the aforementioned Proposition and the subsequent Corollary.

In the general case, we covered X_K with affines $U_{i,K}$, each of which have a separated R-model U_i acquired via the above procedure. We saw that boundedness of $X_K(K^{sh})$ was equivalent to any K^{sh}-valued point of X_K extending to an R^{sh}-valued point of (at least) one of the U_i. So we're pretty much there, the remaining issue is that U_i is an R-model of $U_{i,K}$, not the entire generic fiber. The naive thing to try would be to glue U_i and X_K along $U_{i,K}$, which yields a finite type R-model X_i of X_K – it turns out this naive guess is precisely what gives us separatedness!

Note that X_K, being separated over K, is also separated over R. Thus we are gluing two separated schemes along a common open subscheme $U_{i,K}$. By Lemma 8.2.1.1, this is separated if and only if

$$U_{i,K} \to U_i \times_R X_K \tag{8.2.1-I}$$

is a closed immersion. But note since X_K is a K-scheme, the special fiber of the fiber product $U_i \times_R X_K$ is empty! In particular the fiber product is a K-scheme, so we have

$$\begin{aligned} U_i \times_R X_K &\cong \mathrm{Spec}(K) \times_K (U_i \times_R X_K) \\ &\cong U_{i,K} \times_R X_K. \end{aligned}$$

But then the diagonal embedding in Eqn. 8.2.1-I is the graph of a morphism into a separated target, and hence must be a closed immersion, as can be seen via the following commutative diagram (using preservation of closed immersions under base change)

$$\begin{array}{ccc} U_{i,K} & \xrightarrow{\Delta} & U_{i,K} \times_R X_K \\ \downarrow & & \downarrow \\ X_K & \xrightarrow{\Delta} & X_K \times_R X_K \end{array}$$

It thus follows this X_i we obtain via gluing is separated, as desired. □

EXAMPLE 8.2.1.3. Working with the same setup as Example 8.1.4.5, suppose we have $X_K = \mathbb{P}^1_K$, with the same affine charts $U_{1,K} \cong U_{2,K} \cong \mathbb{A}^1_K$. Each of them admits a separated R-model $\mathbb{A}^1_R$. Gluing this to the entire $\mathbb{P}^1_K$ along $U_{i,K}$ yields our two separated models X_1, X_2, each one isomorphic to the projective R-line minus the special point at infinity, which we know to be separated.

Pictorially we have the following image of X_1, which visually affirms our separatedness intuition.

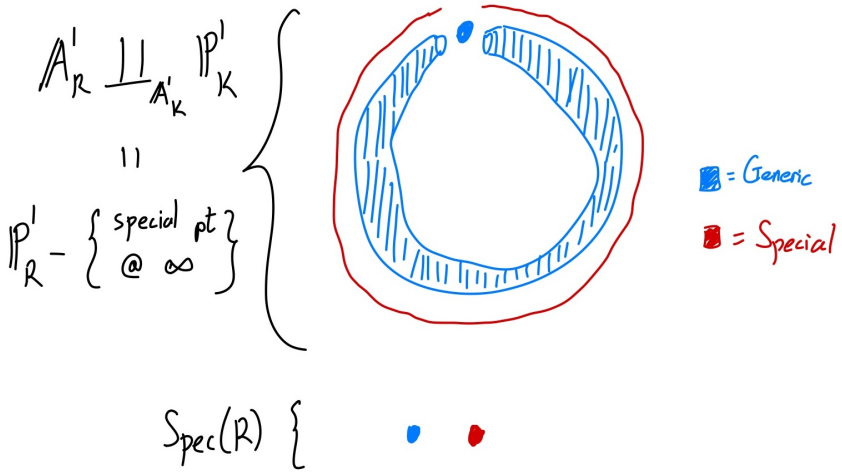

Figure **8.2.1-A**: The separated R-model of $\mathbb{P}^1_K$ obtained by gluing the affine R-line (the separated R-model of $\mathbb{A}^1_K$) via the common open subscheme $\mathbb{A}^1_K$

Note of course that X_1 not satisfy the extension property for étale points alone, which we can see visually from the image above; the generic point at infinity doesn't extend to an R-valued point (hence the same must hold true for strict henselizations). However, it would extend in the other chart X_2. Thus X_1, X_2 collectively satisfy the extension property for étale points.

Proposition 8.2.1.2 gives us the input to the following main theorem of this subsection, which will yield our desired separated model.

THEOREM 8.2.1.4. Let X_K be a separated finite-type K-scheme, and let $X_1, \dots X_n$ be separated finite-type R-models of X_K. Then there exists a separated finite-type R-model X of X_K and proper X_i-schemes X_i' such that the given isomorphisms $X'_{i,K} \xrightarrow{\sim} X_K$ extend to open immersions $X_i' \hookrightarrow X$.

COROLLARY 8.2.1.5. $X_K(K^{sh})$ is bounded if and only if X_K admits a separated finite-type R-model X satisfying the extension property for étale points.

Proof. By Proposition 8.1.4.3, boundedness is equivalent to the existence of a finite set of separated R models $\{X_i\}_{i\in[n]}$ such each K^{sh}-point of X_K extends to an R^{sh}-point of *at least* one of the X_i. By the

valuative criterion of properness (using the fact that $X_{i,K} \cong X'_{i,K} \cong X_K$), we get bijections

$$X'_i(R^{sh}) \xrightarrow{\sim} X_i(R^{sh}),$$

and thus every K^{sh}-point of X_K also extends to an R^{sh} point of one of the X'_i, all of which are embedded in some separated R-model X, completing our proof. □

The main strategy we will employ is to construct these X'_i via a finite sequence of blowups along the special fiber. To this end we make use of a convenient (albeit non-standard) terminology.

DEFINITION 8.2.1.6. For an R-scheme U, a **special blowup tower** is an R-morphism $U' \to U$ constructed via a finite sequence of blowups along closed subschemes of special fibers.

The most difficult stepping stone towards this is the following result on flattening via blowups. For our purposes we will choose to stand on giants' shoulders and assume this as a big technical hammer.

PROPOSITION 8.2.1.7 (Raynaud-Gruson's Flattening by blowups). Let $f : U \to V$ be a morphism of separated flat finite-type R-schemes such that the restriction f_K to generic fibers is flat. Then there exists a special blowup tower $V' \to V$ such that strict transform $f' : U' \to V'$ of f is flat.

SANITY CHECK. Naturally by "strict transform" of a special blowup tower we mean the sequence of strict transforms of each blowup composed together.

Proof. See §5.2 in [GR71]; this is the original source of and classical (sane) proof of this result. Alternatively, for a proof from "first principles" using a modernized (insane) approach via the theory of valuative spaces see Appendix §C.2. This result is then simply a special case of Theorem C.2.0.1. □

Finally, we prove a useful technical lemma for flat morphisms with nice generic base changes.

LEMMA 8.2.1.8. Let $f : X \to Y$ be a flat morphism of flat finite-type separated R-schemes such that $f_K : X_K \to Y_K$ is an open immersion. Then f is an open immersion.

Proof. Since X, Y are finite-type and Noetherian, f is an open morphism [cite maybe?], and for our purposes we may assume it is also surjective. Note it is enough to show that f is an open immersion after faithfully flat base-change. In particular we may base change via $U \to V$ itself, so WLOG we may assume f has a section

$$\epsilon : X \times_Y Y \hookrightarrow X \times_Y X.$$

We want to verify ϵ is an isomorphism. Note since f is open, it is separated, and thus its section ϵ is a closed immersion. In particular the induced map

$$\alpha : \mathcal{O}_X \to \epsilon_* \mathcal{O}_Y$$

is surjective, lets say with kernel I. Note that I, as a sub-R-module of $\mathcal{O}_X$, must also be a flat R-module (since we're working over a DVR, flatness is equivalent to having no torsion). But note since f_K is an isomorphism, the map $\alpha \otimes_R K$ has trivial kernel $I \otimes_R K$. But this implies I is entirely R-torsion, and thus must be trivial. It follows α is an isomorphism, and thus so is ϵ, as desired. □

We are now equipped to prove the main theorem of this subsection.

Proof of 8.2.1.4. Note it suffices to verify the theorem statement for $n = 2$ (i.e when there are just two separated R-schemes X_1, X_2), and the general case follows inductively.

To this end, denote by Γ_K the graph of the isomorphism

$$X_{1,K} \xrightarrow{\sim} X_{2,K},$$

and let Γ be its schematic closure in $X_1 \times_R X_2$. Applying the "flattening by blowups" procedure of Proposition 8.2.1.7 to the projection morphisms $p_i : \Gamma \to X_i$ (which we may do since on the generic fiber $p_{i,K}$ is an isomorphism, hence flat), there exists a special blowup tower $X_i' \to X_i$ such that both strict transforms

$$p_i' : \Gamma_i' \to X_i'$$

are flat. Note each $\Gamma_i' \to \Gamma$ is also a special blowup tower. Let's say the first step of each tower is a blowup

$$\Gamma_{i,1}' \to \Gamma$$

defined by some ideal sheaf $\mathscr{I}_i \trianglelefteq \mathcal{O}_\Gamma$, for $i = 1, 2$. We may then take the blowup of Γ along the ideal $\mathscr{I}_1 \cdot \mathscr{I}_2$, which is then the blowup of each $\Gamma_{i,1}'$ along the pullback of $\mathscr{I}_{j\neq i}$ along the blowup morphism. Continuing this way for each successive blowup, we get the following commutative diagram of special blowup towers:

$$\begin{array}{ccc} \Gamma' & \longrightarrow & \Gamma_1' \\ \downarrow & & \downarrow \\ \Gamma_2' & \longrightarrow & \Gamma \end{array}$$

Moreover, flatness of both the $p_i' : \Gamma_i' \to X_i'$ allow us to apply Lemma 8.2.1.8 to deduce each p_i' is an open immersion. Thus by taking the sequence of blowups of X_i' along the closed subschemes defining the blowup tower $\Gamma'' \to \Gamma_i$, we get special blowup towers

$$X_i'' \to X_i'$$

such that the restriction to $\Gamma_i' \subseteq X_i'$ yields the blowup tower $\Gamma' \to \Gamma_i'$. In particular $\Gamma' \hookrightarrow X_i''$ is an open immersion for both i. Thus we may glue X_1'' and X_2'' along Γ' to obtain X, i.e

$$X := X_1'' \sqcup_{\Gamma'} X_2''.$$

Note that this X itself is an R-model of X_K, since any and all blowups done were within special fibers. We claim that this X is what we want; i.e that X is separated.

To see this, consider Γ^*, the schematic closure of the graph of the isomorphism $X_{1,K}'' \xrightarrow{\sim} X_{2,K}''$ inside $X_1'' \times_R X_2''$. Since Γ' was flat over R, the canonical isomorphism $\Gamma_K' \xrightarrow{\sim} \Gamma_K^*$ extends to a morphism $\Gamma' \to \Gamma^*$. Similarly, the isomorphism $\Gamma_K^* \xrightarrow{\sim} \Gamma_K$ extends to a morphism $\Gamma^* \to \Gamma$. The composition

$$\Gamma' \to \Gamma^* \to \Gamma$$

is proper [say why], and by cancellation it follows $\Gamma' \to \Gamma^*$ is proper. Thus Γ' is closed in Γ^*, thus closed in $X_1'' \times_R X_2''$. It follows then by Lemma 8.2.1.1 that X is separated, as desired. □

Thus we have reduced to the case of a separated R-model X of X_K satisfying the extension property for étale points. However, this X need not be R-smooth at all, which we will address in the next few

sections of this chapter.

§8.2.2 Smoothening: Theorem Statement

Our goal in this section (and the following few sections of this chapter) will be to take a separated R-model of a smooth, separated, finite-type X_K and make it smooth.

In all that follows, R is a DVR with field of fractions K, residue field k and uniformizer π. We denote by R^{h} (resp. R^{sh}) the henselization (resp. strict henselization) of R. Recall [Link reference here maybe?] that R^h and R^{sh} are also dvrs with π as a uniformizer, and the residue field of R^{sh} is the separable closure of k. For any R-scheme X we denote by X_K its generic fiber and X_k its special fiber.

DEFINITION 8.2.2.1. Let X be a finite-type R scheme with smooth generic fiber X_K. A **smoothening** of X is an R-morphism $f : X' \to X$ satisfying the following conditions:

(i) f is proper and an isomorphism on generic fibers.

(ii) The canonical map
$$\mathrm{sm}(X'/R)(R^{sh}) \to X(R^{sh})$$
induced by f is bijective.

REMARK 8.2.2.2. In [BLR90a], condition (ii) of the above definition was phrased as follows:

(*) For each étale R-algebra R', each R'-valued point of X lifts uniquely to an R'-valued point X' which factors through the R-smooth locus $\mathrm{sm}(X'/R)$.

We explain why this is equivalent to our condition (ii), assuming condition (i) holds.

Any étale R-algebra R' is integral over R, hence is semi-local (i.e has finitely many maximal ideals, see [Mat70, §5.E]). Thus note that
$$\mathrm{Spec}(R') = \bigcup_{\mathfrak{m}_i \trianglelefteq R' \text{maximal}} \mathrm{Spec}(R'_{\mathfrak{m}_i}),$$
where here naturally we identify all the points contained in $\mathfrak{m}_i$ with the points of $\mathrm{Spec}(R'_{\mathfrak{m}_i})$. Note each localization at the maximal ideal is a local ring. Thus in order to test condition (ii) of the above, it suffices to restrict to local étale extensions R'/R; all such rings are flat dvrs over R. But the valuative criterion of properness, which we can apply by condition (i), implies that the map $X'(R') \to X(R')$ induced by f is automatically bijective for any such R'. Thus condition (ii) is tantamount to saying this bijection comes solely from the R-smooth points of X'. Since the strict henselatization R^{sh} is built as a direct limit of étale local extensions of R, checking this condition on all R-algebras is equivalent to checking bijectivity of
$$\mathrm{sm}(X'/R)(R^{sh}) \to X(R^{sh}),$$
which is precisely our formulation of condition (ii).

KEY IDEA. Given a non-smooth R-model X of a K-scheme X_K that satisfies the extension property of étale points, by taking the smooth locus $X'' := \mathrm{sm}(X'/R)$ of a smoothening of X, we get a smooth R-model of X_K that still satisfies the extension property for étale points, by condition (ii) of the definition.

In general, a smoothening $X' \to X$ is not a **resolution of singularities** (i.e a projective morphism $X'' \to X$, with X' regular, that is an isomorphism over $\mathrm{reg}(X)$), because the points in the complement of the smooth locus of X' do not need to be regular.

EXAMPLE 8.2.2.3. Take $X = \mathbb{A}^2_R$ the affine plane. Blowing up at any closed point p yields a smoothening (trivially), but not a desingularization, since the map $\mathrm{Bl_p(X)} \to \mathrm{X}$ is not an isomorphism over $\mathrm{reg}(X) = X$.

WARNING (Convention). What we are calling a "resolution of singularities" is often referred to as a **strong** resolution of singularities; other authors may not require an isomorphism over the regular locus, for example, but just general birationality. Other authors still may change the "projective" condition to "proper", and so on and so forth. In general there is no consenus on a list of properties that determines something to be a resolution of singularities, only a lot of similar ideas.

However, a resolution of singularities of X (in the sense of our definition) is always a smoothening. We expound a little on this result, for which we first need a technical lemma from commutative algebra

LEMMA 8.2.2.4. Let $\iota : R \to A$ and $\varepsilon : A \to R$ be morphisms of regular local rings such that $\varepsilon \circ \iota = \mathrm{id}_R$. Then:

(a) The image of each regular system of parameters of R under ι is part of a regular system of parameters of A.

(b) $\ker(\varepsilon)$ is generated by part of a regular system of parameters.

(c) If $t_1, \dots t_n$ is a minimal set of generators for $\ker(\varepsilon)$, then the completion of A with respect to $\ker(\varepsilon)$ is (canonically) isomorphic to $R[\![t_1, \dots t_n]\!]$.

Proof. Let $\mathfrak{m}_R \trianglelefteq R, \mathfrak{m}_A \trianglelefteq A$ denote the respective maximal ideals, and let $r_1, \dots r_m$ be a minimal set of generators of $\mathfrak{m}_R$ (i.e a regular system of parameters).

We first deal with assertion (a). Since $\varepsilon \circ \iota = \mathrm{id}_R$, it follows they are both local maps, and so $\iota(r_i) \in \mathfrak{m}_A$ for all i. We can also see they are A-linearly independent; indeed, if not, we'd have

$$\varepsilon\left(\sum_{i=1}^{m-1} a_i \iota(r_i)\right) = \varepsilon(\iota(r_m)) \implies \sum_{i=1}^{m-1} \varepsilon(a_i) r_i = r_m,$$

contradicting the assumed R-linear independence of the r_i. Thus $\iota(r_i)$ is part of a regular system of parameters for A.

For assertion (b), by assertion (a) we may write a minimal set of generators

$$\mathfrak{m}_A = \langle \iota(r_1), \dots, \iota(r_m), t_1, \dots t_n \rangle .$$

By replacing each t_i with $t_i - \iota(\varepsilon(t_i))$ if necessary (which preserves linear independence by the same type of argument given for (a)!), we may assume $t_i \in \ker(\varepsilon)$. Thus we visibly have $\langle t_1, \dots t_n \rangle \subseteq \ker(\varepsilon)$, and we claim the reverse containment also holds. Indeed, consider the ring

$$A' := A/\langle t_1, \dots t_n \rangle,$$

which is a regular local ring of dimension m (since we're quotienting out by elements of the regular system). Note ϵ must factor through this quotient, since it was a sub-ideal of the kernel, so the induced

$$\varepsilon' : A' \to R$$

is a surjective map of regular local rings of equal dimension, hence an isomorphism. It follows $\ker(\varepsilon) = \langle t_1, \dots t_n \rangle$, as desired.

Finally, for the last assertion, recall that for any noetherian ring A with ideal $I = \langle t_1, \dots t_n \rangle$, the I-adic completion is given by the following quotient of the formal power series ring:

$$A[\![x_1, \dots x_n]\!]/\langle x_1 - t_1, \dots, x_n - t_n \rangle \cong R[\![x_1, \dots x_n]\!],$$

as desired. □

> PROPOSITION 8.2.2.5. Any resolution of singularities $X' \to X$ is a smoothening.

Proof. By definition a resolution of singularities is projective, hence proper, and moreover we know the non-R-smooth locus lies in the special fiber, so we have an isomorphism on generic fibers. Thus it remains to verify the map

$$sm(X'/R)(R^{sh}) \to X(R^{sh})$$

is bijective. By our discussion in Remark 8.2.2.2, this is equivalent to requiring that, for any étale R-algebra R', every $a \in X'(R')$ factors through the R-smooth locus of X'. Note a resolution of singularities is preserved under base change, so the map

$$X' \otimes_R R' \to X \otimes_R R'$$

is also a resolution of singularities. Moreover if

$$(a, \mathrm{id}) : \mathrm{Spec}(R') \to X_{R'}$$

factors through the smooth locus, the same is true for $a \in X'(R')$ [Cite maybe]. Thus it suffices to assume $R = R'$. Then [cite some EGA]. □

The main result of the rest of this chapter we will work towards is the following theorem.

> THEOREM 8.2.2.6. Let X be a finite-type R-scheme with smooth generic fiber X_K. Then X admits a smoothening $f : X' \to X$, constructed as a special blowup tower.

> COROLLARY 8.2.2.7. Given a smooth, separated, finite-type K-scheme X_K, boundedness of $X_K(K^{sh})$ is equivalent to the existence of a weak Néron model X of X_K.

Proof. We have seen (several times now) that boundedness of $X_K(K^{sh})$ is implied by the extension property for étale points. For the other direction, Corollary 8.2.1.5 tells us boundedness is equivalent to existence of a separated finite-type R-model X satisfying the extension property for étale points. Applying the above Theorem and taking

$$X'' := \mathrm{sm}(X'/R),$$

we get an R-model of X_K (since $X'_K \cong X_K$, which itself is R-smooth so lies within X'') that remains separated and of finite type (both conditions are unaffected by blowups) and is also smooth. Moreover by definition of a smoothening we have a bijection on R^{sh}-valued points of X'' and X', so the extension property for étale points is preserved. Thus X'' is a weak Néron model of X_K. □

We remark that such X'' in the above statement need not be unique nor proper over R (even if X is proper over R!). [give example maybe?]

The key tool for blowup we will use is *dilatations*.

PROPOSITION 8.2.2.8 (Dilatations Exist). Let X be a finite-type R-scheme, and let Y_k be a closed subscheme of the special fiber (given the reduced structure), defined by some sheaf of ideals $\mathcal{I} \trianglelefteq \mathcal{O}_X$. Let $Dil_{Y_k}(X) \to X$ denote the restriction of $Bl_{Y_k}(X)$ to the open subscheme where $\mathcal{I}\mathcal{O}_{Bl_{Y_k}(X)}$ is generated by π. Then $Dil_{Y_k}(X)$ is a flat R-scheme, and the map on special fibers factors through Y_k. Moreover it is universal among all such flat R-schemes, i.e for any R-flat Z, there is a unique dotted arrow making the following diagram commute:

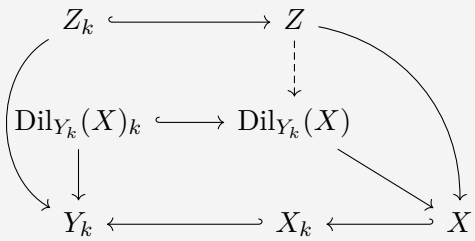

We call $Dil_{Y_k}(X)$ the **dilatation** of Y_k in X.

Naturally if $\mathrm{Dil}_{Y_k}(X)$ in the above Proposition exists, it must be unique up to canonical isomorphism, so calling it "the" dilatation makes sense.

ASIDE (Lexical Pedantry). The words "dilation" and "dilatation" originate from different grammatical tenses of the original Latin root *dilatare* and are often used interchangeably in colloquial parlance. However, in medical terminology, "dilation" is used for naturally occurring widenings (e.g. that of pupils in a dark room), whereas "dilatation" is used for forceful or more extreme unnatural widenings, e.g. one caused by an illness. For this reason we opt for the latter variation of the word, since we are forcibly enlarging a scheme via blowups.

Proof. We will first construct the dilatation affine-locally, and the global case will follow easily. Recall the blowup of X is defined locally as follows: if $X = \mathrm{Spec}(A)$, then Y_k is cut out by some ideal $I \trianglelefteq A$. Since A is finitely generated over a dvr, it is noetherian, so I is generated by finitely many elements $g_0 = \pi, g_1, \dots g_n$. Note that $\pi \in I$ since I cuts out a closed subscheme of the special fiber of X. Then the blowup $Bl_I(A)$ is

the closed subscheme of $\mathbb{P}^n_A$ cut out by the homogeneous ideal

$$I' := \ker\left(A[t_0, \dots t_n] \to \oplus_{n\geq 0} I^n,\ t_i \mapsto g_i\right).$$

If U_i is the affine open subscheme of $\mathbb{P}^n_A$ where t_i doesn't vanish, then

$$Bl_I(A) \cap U_i = \text{Spec}\left(A\left[\frac{g_0}{g_1}, \dots, \frac{g_n}{g_i}\right] \Big/ (g_i\text{-torsion})\right).$$

We now define

$$\text{Dil}_I(A) := Bl_I(A) \cap U_0 = \text{Spec}\left(A\left[\frac{g_0}{\pi}, \dots, \frac{g_n}{\pi}\right] \Big/ (\pi\text{-torsion})\right)$$

We claim this is exactly what we want, in the local case. Note that since by construction the ring of global sections of $\text{Dil}_I(A)$ is π-torsion free, it is flat over R (even if A is not!). We will then check the universal property for the case of affine $Z = \text{Spec}(B)$, where B is a flat R-algebra. Suppose

$$\phi : Z \to X$$

is a morphism given by the map of R-algebras $\phi^\# : A \to B$. The fact that $\phi_k : Z_k \to X_k$ factors through $Y_k = V(I) \subseteq X_k$ implies the ideal $\phi^\# \cdot B$ is contained inside πB. Thus, for every generator g_i of I, there exist $h_i \in B$ so that

$$\phi^\#(g_i) = \pi h_i.$$

Moreover these h_i are unique, since B has no π-torsion. This, combined with R-flatness of B, tells us the map

$$\begin{aligned} A[t_0, \dots t_n] &\to B \\ t_i &\mapsto h_i \\ a &\mapsto \phi^\#(a) \end{aligned}$$

must factor uniquely through the quotient

$$A[t_0, \dots, t_n] \Big/ \langle \pi\text{-torsion},\ g_i - \pi T_i \rangle_{j\neq i \in [n]} \cong A\left[\frac{g_0}{\pi}, \dots, \frac{g_n}{\pi}\right] \Big/ (\pi\text{-torsion}) = \mathcal{O}_{\text{Dil}_I(A)}$$

In other words, the morphism $Z \to X$ factors uniquely through $\text{Dil}_I(A) \to X$. This completes our proof in the local case.

The global case, for X some arbitrary separated finite-type R-scheme and Y_k some closed subscheme of the special fiber defined by some sheaf of ideals $\mathcal{I} \trianglelefteq \mathcal{O}_X$, we define

$$\text{Dil}_{Y_k}(X) := \left\{x \in \text{Bl}_{Y_k}(X) : \mathcal{I}'_x \text{ is generated by } \pi\right\}.$$

Here the $\mathcal{I}'$ is the sheaf of ideals generated locally by what we called I' in our above argument. This construction is an open subscheme of $\text{Bl}_{Y_K}(X)$ [why?], whose intersection with any affine $\text{Spec}(A)$ is precisely the subscheme $\text{Dil}_I(A)$ we described above. Since all the properties of the dilatation we want to verify are local, they follow from our arguments in the affine case! □

EXAMPLE 8.2.2.9. Suppose $X = \text{Spec}(R[t_1, \dots t_n])$, and suppose Y_k is the origin of the special fiber,

i.e.

$$Y_k = \{\langle \pi, t_1, \dots t_n \rangle\} \subseteq X_k.$$

Then using the construction given in the above proof we see that

$$\mathrm{Dil}_{Y_k}(X) = \mathrm{Spec}\left(R\left[\frac{t_1}{\pi}, \dots, \frac{t_n}{\pi}\right]\right).$$

The map $\mathrm{Dil}_{Y_K}(X)$ to X is naturally the one induced by sending $t_i \to t_i$ on coordinate rings.

PROPOSITION 8.2.2.10 (Elementary Properties of Dilatations). Let X be an R-scheme, with Y_k a closed subscheme of the special fiber. Then

(a) All dilatations along *any* closed subscheme of the special fiber factor through the ideal of π-torsion in $\mathcal{O}_X$ (i.e. the largest R-flat subscheme of X).

(b) Dilatations commute with flat base-change $R \to R'$ that send uniformizers to uniformizers.

(c) Suppose X is a closed subscheme of an R-scheme Z. Then $\mathrm{Dil}_{Y_k}(X)$ is a closed subscheme of $\mathrm{Dil}_{Y_k}(Z)$.

(d) Dilatations commute with products, in the sense that if $Y_{1,k}, Y_{2,k}$ are closed subschemes of the special fibers of X_1, X_2, then

$$\mathrm{Dil}_{Y_{1,k} \times_k Y_{2,k}}(X_1 \times_R X_2) = \mathrm{Dil}_{Y_{1,k}}(X_1) \times_R \mathrm{Dil}_{Y_{2,k}}(X_2).$$

Proof. All these properties follow directly by our construction given in the proof of 8.2.2.8; we leave it to the reader to explicitly work out the details if they wish. □

§8.2.3 Defect of Smoothness

As in previous subsections, we denote by R a DVR with field of fractions K, and X_K a smooth separated K-scheme of finite-type. In this subsection we construct a measure of how "far away" an R-model X (of finite type!) of X_K is from being smooth. Let

$$a : \mathrm{Spec}(R^{sh}) \to X$$

be an R^{sh}-valued point of X. We consider the pullback $a^*\Omega^1_{X/R}$ of the $\mathcal{O}_X$-module of relative Kähler differentials from X to $\mathrm{Spec}(R^{sh})$. Note since X was finite type over R, the module of global sections of the pullback is a finite-rank R^{sh}-module. Since R^{sh} is also a DVR, the structure theorem for modules over a PID tells us this module splits into a direct sum of free parts and torsion parts.

DEFINITION 8.2.3.1. With the notation above, we call

$$\delta(a) := \text{ Length of } R^{sh}\text{−torsion}(\Gamma(a^*\Omega^1_{X/R}))$$

the **defect of smoothness** at a.

Of course, we need to justify that $\delta(a)$ indeed deserves its name, in the sense that it actually measures "distance from smoothness".

PROPOSITION 8.2.3.2. Any $a \in X(R^{sh})$ factors through $\mathrm{sm}(X/R)$ if and only if $\delta(a) = 0$.

Proof. Let a_K, a_k denote the generic and special fibers of a, respectively. One direction of the proposition statement is clear: if $a \in \mathrm{sm}(X/R)$, then $\Omega^1_{X/R}$ is locally free at a_k, and hence $a^*\Omega^1_{X/R}$ is free, so $\delta(a) = 0$.

For the other direction, if $\delta(a) = 0$, then $a^*\Omega^1_{X/R}$ is freely generated by d elements. Note that since X_K is smooth at a_K, the rank d of $a^*\Omega^1_{X/R}$ is simply the dimension of the generic fiber X_K at a_K. By Chevalley's upper semi-continuity theorem (see [EGA, IV, §13.1.3]), we have

$$\dim_{a_k}(X) \geq d.$$

Moreover, $a^*\Omega^1_{X/R}$ being free of rank d tells us $\Omega^1_{X/R}$ (and hence $\Omega^1_{X_k/k}$ is generated by d elements at a_k. In other words, X is k-smooth at a_k. We wish to deduce R-smoothness.

Now pick some neighborhood $U \ni a_k$, and pick some closed immersion $\iota U \subseteq \mathbb{A}^n_R$, and by abuse of notation we identify the point and the open set with their images under the immersion. By applying the Jacobi criterion for k-smoothness of X_k at a_k, it follows that there exist sections $g_{d+1}, \ldots, g_n$ of $\mathcal{O}_{\mathbb{A}^n_R} \cong R[t_1, \ldots t_n]$ that cut out U_k along with π, i.e.

$$\mathcal{O}_{U_k} \cong R[t_1, \ldots, t_n] / \langle \pi, g_{d+1}, \ldots, g_n \rangle .$$

such that the differentials $\mathrm{d}g_i\ (a_k)$ are k-linearly independent in $\Omega^1_{\mathbb{A}^n_k/k} = \Omega^1_{\mathbb{A}^n_R/R} \otimes \kappa(a_k)$. But then we can apply the reverse direction of the Jacobi criterion to deduce that the closed subscheme $Z \subseteq \mathbb{A}^n_R$ cut out by $g_{d+1}, \ldots, g_n$ is R-smooth of relative dimension d at a_k. In other words, X is R-smooth at a_k, and thus a factors through the R-smooth locus of X. □

We can generalize our use of the Jacobi criterion in the above proof to give an explicit method to compute $\delta(a)$. Take a neighborhood $U \ni a$, and pick a closed immersion $U \hookrightarrow \mathbb{A}^n_R$. Suppose $g_1, \ldots g_m \in \mathcal{O}_{\mathbb{A}^n_R}$ generate

$$\mathcal{I}_U \trianglelefteq \mathcal{O}_{\mathbb{A}^n_R} \cong R[z_1, \ldots z_n]$$

defining U as a closed subscheme. Note $\mathrm{d}z_1, \ldots, \mathrm{d}z_n$ constitute a basis of $\Omega^1_{\mathbb{A}^n_R/R}$. Then we can represent the relative differentials $\mathrm{d}g_i$ in terms of the $\mathrm{d}z_i$, which we write as

$$\mathrm{d}g_i = \sum_{j=1}^{n} \frac{\partial g_i}{\partial z_j} \mathrm{d}z_j ,$$

where here $\frac{\partial g_i}{\partial z_j}$ is simply our notation for the uniquely determined coefficients. Recall we then define the Jacobian matrix as

$$J := \left[\frac{\partial g_i}{\partial z_j} \right]_{i \in [m], j \in [n]} \in \mathrm{Mat}_{m \times n}(R).$$

LEMMA 8.2.3.3. Take some $a \in X(R^{sh})$, and suppose X has relative dimension d at a_K. With the notation as above, let Λ denote the set of all $(n-d)$ minors of the Jacobian J. Then

$$\delta(a) = \min \{ \nu_\pi(a^*\Delta) : \Delta \in \Lambda \},$$

where here $\nu_\pi(a^*\Delta)$ denotes the π-order of generating ideal of the matrix $a^*\Delta$ (once put into Smith

normal form).

Proof. By (K-)smoothness of X_K, the Jacobi criterion tells us there must exist some $\Delta \in \Lambda$ such that $a_K^*\Delta \neq 0$. But since $a_K^*\Delta$ is simply the image of $a^*\Delta$ under the natural injection $R^{sh} \hookrightarrow K^{sh}$, it follows $a^*\Delta \neq 0$ as well. [finish] □

By a similar technique we can show that $\delta(a)$ is in fact bounded as a varies over $X(R^{sh})$; in other words, there is some finite measure of the global "defect of smoothness" of X.

PROPOSITION 8.2.3.4. For X a finite-type R-model of a finite-type smooth K-scheme X_K,

$$\max\left\{\delta(a) : a \in X(R^{sh})\right\} < \infty.$$

Proof. Since any finite-type scheme is quasi-compact, we may reduce to the case of X being affine, i.e

$$X = \mathrm{Spec}(A) = \mathrm{Spec}\left(R[t_1, \dots t_n]/\left\langle g_1, \dots g_m\right\rangle\right).$$

If we let X_K^d denote the union of all d-dimensional irreducible components of X_K, we know by smoothness that X_K is the disjoint union of all the (finitely many!) X_K^d. Let X^d denote the schematic closure of X_K^d in X (i.e subscheme defined by the preimage of the ideal of definition of X_K^d under the map $A \to A \otimes_R K$), and define A^d to be its ring of global sections.

Using the same notation as in the preamble to and statement of Lemma 8.2.3.3, the Jacobi criterion for any $x \in X_K^d$ tells us that there exists an $(n-d)$-minor Δ_x of the Jacobian at x such that $\Delta_x(x) \neq 0$. Altogether then, this implies that the family $\{\Delta_x\}_{x \in X_K^d}$ generate the unit ideal in $A^d \otimes_R K$.

By chasing denominators then, we can find elements $f_1, \dots f_r \in A^d$ and minors $\Delta_{x_1}, \dots \Delta_{x_r}$ such that

$$\sum_{i=1}^{r} f_1 \Delta_{x_1}|_{X^d} = \pi^c,$$

for some $c \geq 0$. Lemma 8.2.3.3 then tells us that $\delta(a) \leq c$ for all $a \in X(R^{sh})$ such that $a_K \in X_K^d$. Since there's only finitely many non-empty X_K^d, it follows $\delta(a)$ achieves a global maximum. □

Thus we now have a global measure of an R-models defect of smoothness. Our goal will be to find suitable closed subschemes along the special fiber along which blowups reduce this value. For this we introduce the following definition.

DEFINITION 8.2.3.5. Let X be a finite-type R-model of X_K. We say a closed subscheme of the special fiber $Y_k \subseteq X_k$ is **abundant** (in X) if the set of k^{sep}-valued points that lift to R^{sh}-valued points of X is schematically dense in Y_k.

LEMMA 8.2.3.6 (Properties of abundance). If Y_k is abundant in X, then

(a) The set of closed points of Y_k that lift to R'-valued points of X, for R'/R a local étale extension, is schematically dense in Y_k.

(b) The sheaf of ideals defining Y_k is equal to the intersection of all kernels of morphisms

$$a^* : \mathcal{O}_X \to a_*\mathcal{O}_{\mathrm{Spec}(k^{\mathrm{sep}})},$$

where a varies over the R^{sh}-valued points of X factoring through Y_k.

(c) Y_k is geometrically reduced, and its k-smooth locus is open and dense in Y_k.

Proof. Assertion (a) is a simply a consequence of the definition of the strict henselization R^{sh} as a colimit of étale-local extensions of R, thus allowing us to deescend the lifting property to any such extension.

Assertion (b) [look in EGA maybe]. □

Now we are ready for the key process that we will use in our proof of Theorem 8.2.2.6: we will construct a dilatation of X along an abundant subscheme of the special fiber, and doing so repeatedly will decrease the defect of smoothness inside it.

PROPOSITION 8.2.3.7. Let Y_k be abundant in X, and let U_k be a k-smooth open subscheme of Y_k such that the pullback $\Omega^1_{X/R}|_{U_k}$ is locally free. For each $a \in X(R^{sh})$, denote by a' its lift to $(\mathrm{Dil}_{Y_k}(X))(R^{sh})$. If $a_k \in U_k$, then

$$\delta(a') \leq \max\{0, \delta(a) - 1\}.$$

In particular, $\delta(a) < \delta(a)$ for all non-smooth R^{sh}-valued points of X which specialize into points of U_k.

Recall that by a_k we mean the specialization of $a \in X(R^{sh})$ into the special fiber. Before delving into the proof we illustrate with an example.

EXAMPLE 8.2.3.8. Suppose we have

$$X := \left\{st = \pi^2\right\} \subseteq \mathbb{A}^2_R = \mathrm{Spec}(R[s,t])$$

Let Y_k be the closed subscheme defined by $\langle \pi, s, t\rangle$; this consists of a single k-valued point. This is relatively easily seen to be abundant, as the point lifts an R-valued point of X, in particular the the map defined by

$$\begin{aligned} R[s,t]/\left\langle st - \pi^2 \right\rangle &\to R \\ s, t &\mapsto \pi \end{aligned}$$

The same map also defines an R^{sh}-valued point. We now compute the defect of smoothness at this point (which we denote by $a \in X(R^{sh})$.

Taking the derivative of the defining equation of X with respect to s and t, we see the Jacobian is the 1×2 matrix given by

$$[t, s].$$

We see immediately that $a^*\Delta = \pi$ for either of the two minors Δ of this Jacobian (just plug in $s = t = \pi$), and subsequently we deduce from Lemma 8.2.3.3 that $\delta(a) = 1$.

Computating the dilatation of X with respect to Y_k, we get

$$\mathrm{Dil}_{Y_k}(X) = \mathrm{Spec}(A') = \mathrm{Spec}\left(A[x,y]\Big/\langle x - \pi s, y - \pi t\rangle\right) = \mathrm{Spec}\left(R[x,y]\Big/\langle xy - 1\rangle\right).$$

This is immediately seen to be R-smooth, and thus we know the lift of a to $a' \in (\mathrm{Dil}_{Y_k}(X))(R^{sh})$ must satisfy $\delta(a') = 0$, by Proposition 8.2.3.2.

Proof of 8.2.3.7. The claim is local on X, so we may reduce to the case of X being affine and Y_k being irreducible of pure dimension r. Then the sheaves $\Omega^1_{Y_k/k}$ and $\Omega^1_{X/R}|_{Y_k}$ are locally free, and the first is obtained from the second by quotienting out the submodule generated by $\{\, \mathrm{d}g \ : g|_{Y_k} = 0\}$. By shrinking further, we can in fact assume that both sheaves are free, and moreover that there exists a basis

$$\mathrm{d}y_1 \ , \ldots , \ \mathrm{d}y_r \ , \ \mathrm{d}z_1 \ , \ldots , \ \mathrm{d}z_{n-r}$$

of $\Omega^1_{X/R}|_{Y_k}$ such that $\mathrm{d}y_1 \ , \ldots , \ \mathrm{d}y_r$ give rise to a basis of $\Omega^1_{Y_k/k}$, and $z_1, \ldots z_{n-r}$ vanish on Y_k. [complete] □

To formulate Proposition 8.2.3.7 in a way that will help us construct the smoothening, we introduce the following definition.

DEFINITION 8.2.3.9. Let E be a subset of $X(R^{sh})$. A closed subscheme Y_k of the special fiber X_k is called **E-permissible** if the following conditions hold:

1. Y_k is abundant in X (in the sense of Definition 8.2.3.5).
2. If U_k denotes the largest k-smooth open subscheme of Y_k where $\Omega^1_{X/R}|_{Y_K}$ is locally free, then there is no k^{sep}-valued point in $Y_k - U_k$ which lifts to a point in E.

OBSERVATION. The associated subscheme U_k of an E-permissible Y_k is always Zariski dense in Y_k by Lemma 8.2.3.6.

We now have a more succint way of formulating Proposition 8.2.3.7.

COROLLARY 8.2.3.10. Let Y_k be an E-permissible closed subscheme of the special fiber of X, and take some $a \in E$. Let $a' \in (\mathrm{Bl}_{Y_k}(X'))(R^{sh})$ denote note the unique lift of a to the blowup. Then

(a) If a doesn't specialize into a point of Y_k, then $\delta(a) = \delta(a')$.

(b) If a specializes into a point of Y_k, then $\delta(a') \leq \max\{0, \delta(a) - 1\}$.

Proof. Let a_k denote the specialization of a. If $a_k \notin Y_k$, there exists an open neighborhood $U \ni a$ over which the blowup is an isomorphism, hence $\delta(a) = \delta(a')$.

Otherwise if $a_k \in Y_k$, then by the universal property of the dilatation $\mathrm{Dil}_{Y_k}(X)$ (as written in Proposition 8.2.2.8) tells us that in fact $a' \in \mathrm{Dil}_{Y_k}(X)$. Since the dilatation is constructed as an open subscheme of the blowup $\mathrm{Bl}_{Y_k}(X)$ and Y_k is E-permissible, we see the conditions of Proposition 8.2.3.7 are met, thus giving us the desired bound on $\delta(a')$. □

We introduce one more convenient piece of notation before proceeding to the proof of the main result on smoothening (Theorem 8.2.2.6).

DEFINITION 8.2.3.11. Given a subset $E \subseteq X(R^{sh})$, the **defect of smoothness of X along E** is defined to be

$$\delta(X, E) := \max\{\delta(a) : a \in E\}.$$

SANITY CHECK. Recall Proposition 8.2.3.4 tells us $\delta(X, E)$ is always a well-defined natural number.

Proof of 8.2.2.6. It suffices to prove that, for any $E \subseteq X(R^{sh})$, that there exists a finite sequence of E-permissible blowups such that each $a \in E$ factors through $\mathrm{sm}(X'/R)$.

We proceed via induction on $\delta(X, E)$. For the base case, if $\delta(X, E) = 0$ then Proposition 8.2.3.2 tells us a factors through the smooth locus of X itself, so the assertion is trivial.

For the inductive step, we will rearrange things so that Corollary 8.2.3.10 can be applied. First of all, we see that we only care about blowing up along subschemes of the non-smooth locus, so WLOG we may assume $\delta(a) > 0$ for all $a \in E$. Let $E_k \subseteq X(k^{\mathrm{sep}})$ denote the subset induced by specialization. Identifying points in E_k with their associateed closed points in X_k, we denote by $\overline{E_k}$ their Zariski closure in the special fiber (given the reduced structure). We observe that that $\overline{E_k}$ is abundant in X. We now proceed to partition E in a convenient way for us.

Set $F^1 := E$ and $Y_k^1 := \overline{F_k^1} = \overline{E_k}$. Let U_k^1 be the largest open k-smooth subscheme of Y_k^1 where $\Omega^1_{X/R}|_{Y_k^1}$ is locally free. Define

$$E^1 := \left\{a \in F^1 : a_k \in U_k^1\right\}.$$

Now define $F^2 := F^1 - E^1$, and proceed the same way. This yields:

(i) A filtration $X(R^{sh}) \supseteq F^1 \supseteq F^2 \supseteq F^3 \cdots$.

(ii) A disjoint partition

$$E = E^1 \sqcup E^2 \sqcup \cdots E^i \sqcup F^{i+1},$$

for all i.

Since each Y_k^i is abundant, part (c) of Lemma 8.2.3.6 (combined with reducedness of Y_k^i, which tells us the restriction of any coherent $\mathcal{O}_X$-module must be locally free on a dense subset) tells us that each U_k^i is a dense open in Y_k^i. In particular we deduce that

$$Y_k^{i+1} \subseteq Y_k^i - U_k^i \implies \dim(Y_k^{i+1}) < \dim(Y_k^i).$$

It therefore follows that eventually we must have $Y_k^{t+1} = \emptyset$ for some t, and consequently that $F^{t+1} = \emptyset$. Thus we arrive at a partition

$$E = \bigsqcup_{i=1}^{t} E^i.$$

By construction, we have that each U_k^i is k-smooth and $\Omega^1_{X/R}|_{Y_k}$ is locally free on it, and thus each Y_k^i is E^i-permissible. Moreover, recalling that we assumed that E_k was disjoint from the R-smooth locus of X, we also observe that each $Y_k^i = \overline{F_k^i}$ is disjoint from $\mathrm{sm}(X/R)$ as well (since the non-smooth locus is Zariski-closed, taking closure preserves disjointness).

We can now carry out the inductive step. Note Corollary 8.2.3.10 tells us that

$$\delta(\mathrm{Bl}_{Y_k^t}(X), E^t) < \delta(X, E^t).$$

By the inductive hypothesis, then, there exists a map

$$X' \to \mathrm{Bl}_{Y_k^t}(X)$$

that is a sequence of E^t-permissible blowups such that each $a \in E^t$ factors through the smooth locus of X'. Note the composition

$$X' \to \mathrm{Bl}_{Y_k^t}(X) \to X$$

leaves $E^t - E$ undisturbed. Thus this is a sequence of E-permissible blowups.

If we write F' for the (canonical!) lifting of a subset $F \subseteq X(R^{sh})$ to $X'(R^{sh})$, then note we have

$$(E - E^t)' = (E^1)' \sqcup \cdots \sqcup (E^{t-1})'.$$

Note that

$$\delta(X', (E - E^t)') = \delta(X, E - E^t) \leq \delta(X, E).$$

Now we may apply the whole process once over and induct on the number of elements in the partition of E. Running this game again we eventually arrive at a map

$$X'' \to X'$$

that is a sequence of $(E - E^t)'$-permissible (and hence E'-permissible) blowups such that all points of $(E - E^t)'$ factor through $\mathrm{sm}(X''/R)$. Thus

$$X'' \to X' \to \mathrm{Bl}_{Y_k^t}(X) \to X$$

gives us our sequence of E-permissible blowups satisfying the desired property stated at the beginning of this proof, thereby completing it. □

§8.2.4 *S*-Rational Maps and the Weak Néron Mapping Property

We have thusly proved Theorem 8.2.2.6, and thus Corollary 8.2.2.7, and so we have the existence of a Weak Néron model. We emphasize here the key property of such models.

DEFINITION 8.2.4.1. We say an R-model X of X_K satisfies the **weak Néron mapping property** (sometimes abbreviated to **WNMP**). if, for any R-smooth scheme Z with irreducible special fiber Z_k, and any K-rational map

$$f_K : Z_K \dashrightarrow X_K,$$

there exists an extension of f to an R-rational map $f : Z \dashrightarrow X$.

OBSERVATION. Setting $Z = \mathrm{Spec}(R^{sh})$, we see the weak Néron mapping property is a *strengthening* of the defining characteristic of a weak Néron model (8.1.1.10). In particular, any smooth, separated, finite-type R-model satisfying the WNMP is a weak Néron model. The goal of this subsection will be to show the reverse direction.

Though the notion of a rational morphism $X \dashrightarrow Y$ between two schemes is a rather elementary concept in algebraic geometry, the notion of a relative (to a given base S) rational morphism is not considered nearly as generally known. To that end we recall and expound upon it here. The motivation here is that when

working over a base scheme S, the notion of a rational morphism does not play well with respect to base-change $S' \to S$. Thus we want a relative notion of rational morphism that does.

DEFINITION 8.2.4.2. Let X be an S-smooth scheme. We say an open $U \subseteq X$ is **S-dense** if, for each $s \in S$, the fiber $U_s = U \times_S \mathrm{Spec}(\kappa(s))$ is dense in the fiber $X_s = X \times_S \kappa(s)$.

An **S-rational map** $X \dashrightarrow Y$ is an equivalence class of S-morphisms $\phi : U \to Y$, where U is an S-dense subscheme, where two such morphisms $(U, \phi), (U', \phi')$ are equivalent if they coincide on an S-dense open subset of $U \cap U'$.

SANITY CHECK. The above definition assumes the intersection of two S-dense subsets is S-dense, which the reader should check if not immediately obvious. Moreover, we leave the reader to verify mentally that the notions of domain of definition, birationality, and base change, defined as in the case of rationl morphisms, also make sense here.

EXAMPLE 8.2.4.3. Set $S = \mathrm{Spec}(Z)$, and consider the S-rational map

$$\phi : \mathbb{A}^1_{\mathbb{Z}} \dashrightarrow \mathbb{A}^1_{\mathbb{Z}}$$

given by the rational function $\frac{t+1}{t-1}$.

Our first result will be to show that this relative notion of rationality, in particular the domain of definition of an S-rational morphism, behaves nicely with respect to flat base change. Before this we will need a lemma.

LEMMA 8.2.4.4. Let U be an S-dense open subscheme of an S-smooth and quasicompact X. Then U contains a quasicompact S-dense open subscheme. In particular any S-rational $\phi : X \dashrightarrow Y$ can be represented by an S-morphism on a quasicompact S-dense open.

Proof. Let $\{U_i\}_{i \in I}$ be an affine open covering of U, and consider

$$\tau_i := X \times_S U_i \to U_i$$

given by projection onto the second coordinate, for each i. This map admits the tautological section $\delta_i : U_i \to X \times_S U_i$. Let V_i denote the union of all connected components of fibers of τ_i that intersect $\delta_i(U_i)$. Then smoothness of τ_i tells us that V_i is an open subscheme of the fiber product $X \times_S U_i$ (this is an immediate consequence of [EGA, §IV, 15.6.5]).

Define

$$\mathrm{Sat}(U_i) := \mathrm{Im}\left(V_i \hookrightarrow X \times_S U_i \to X\right)$$

to be the image under the projection to the first coordinate. If we view $X \times_S U_i$ as the base change of X under the structure morphism $U_i \to S$, then we note that S-flatness of U_i tells us the projection morphism must be flat as well, and locally of finite presenation (via smoothness of X). Thus it is an open map, and thus the image of the open V_i is open. In particular we may view $\mathrm{Sat}(U_i)$ as a "saturation" of U_i with respect to the structure morphism $X \to S$.

We claim that $\{\mathrm{Sat}(U_i)\}_{i \in i}$ is in fact an open covering of X. Take any $x \in X$, and let $s \in S$ be its image under the stucture morphism. Let $C \subseteq X_s$ be the connected component containing x. S-density of

U tells us that U_s is dense in X_s, so there must exist some $u \in U_s \cap C$. If we take some $i \in I$ such that $u \in U_i$, it follows from our construction earlier that

$$(x, u) \in V_i \implies x \in \mathrm{Sat}(U_i).$$

Thus we see that the saturations of the U_i indeed form an open cover of X. By quasicompactness of X they admit a finite subcover $\left\{\mathrm{Sat}(U_{i_j}\right\}_{j\in[n]}$. Consider then the set

$$U' := \bigcap_{j=1}^{n} U_{i_j} \subseteq U.$$

This is visibly S-dense in U by construction, and moreover quasicompact (as a finite union of affines). □

Definition 8.2.4.5. Let $\phi : X \dashrightarrow Y$ be an S-rational morphism represented by a morphism $U \to Y$ on a quasicompact open. We define the **graph of** ϕ to be the schematic closure of the graph Γ_U of $U \to Y$ inside $X \times_S Y$.

Sanity Check. The definition given above is well-defined i.e. independent of choice of representative U, since any quasicompact S-dense open $V \subseteq U$ is schematically dense in U (this is an immediate consequence of [EGA, §IV, 11.10.10]), thus $V \cong \Gamma_V$ and $U \cong \Gamma_V$ would have the same schematic closure in $X \times_S Y$.

For an S-rational $\phi : X \dashrightarrow Y$, suppose Ω is the largest open subscheme of X such that the projection $p_X : X \times_S Y \to X$ onto the first coordinate induces an isomorphism

$$\Gamma \cap p^{-1}(\Omega) \xrightarrow{\sim} \Omega.$$

Then note we must have $\Omega \subseteq \mathrm{dom}(\phi)$. Moreover, if Y is separated over S, then any graph Γ_U, for $U \to Y$ any representative of ϕ, must be closed inside $U \times_S Y$. In particular we have $\Gamma \cap U \times_S Y = \Gamma_U$ We would thus have an isomorphism

$$\Gamma \cap p^{-1}(U) \xrightarrow{\sim} U$$

and thus $U \subseteq \Omega$. Thus $\mathrm{dom}(\phi) \subseteq \Omega$ and we have an equality

$$\Omega = \mathrm{dom}(\phi).$$

In particular there is a unique S-morphism

$$\Omega = \mathrm{dom}(\phi) \to Y$$

representing to the S-rational ϕ.

Proposition 8.2.4.6. Suppose X, Y, X' are smooth S-schemes, all locally of finite presentation. Suppose Y is separated, and that we have an S-rational map $\phi : X \dashrightarrow Y$ and a flat S-morphism $X' \to X$. Then $f^{-1}(\mathrm{dom}(\phi))$ is an S-dense open subscheme of X' and $\phi \circ f$ is an S-rational map satisfying

$$\mathrm{dom}(\phi \circ f) = f^{-1}(\mathrm{dom}(\phi)).$$

Proof. Since these are all smooth finite-type schemes, all the morphisms are locally of finite presentation.

This combined with flatness of f tell us it must be universally open [Stacks, p. 29.26.10]. It follows then that $f^{-1}(\mathrm{dom}(\phi))$ is actually S-dense in X', since the base change $X'_s \to X_s$ is also an open map for any fiber.

Let $\Gamma \subseteq X \times_S Y$ denote the graph of ϕ and $\Gamma' \subseteq X' \times_S Y$ the graph of $\phi \circ f$. Then since schematic images commute with flat base change [cite maybe], it follows that

$$\Gamma' = X' \times_X \Gamma.$$

Let $p : \Gamma \to X$ and $p' : \Gamma' \to X'$ denote the projections onto the first factors. If we set

$$U' := \mathrm{dom}(\phi \circ f), \quad , U := f(U'),$$

then faithful flatness of $U' \to U$ tells us that p is an isomorphism over U if and only if p' is an isomorphism over U'. Therefore $U \subseteq \mathrm{dom}(\phi)$, and the desired equality follows. □

Finally we show that the domain of an S-rational map is also compatible with flat base change.

PROPOSITION 8.2.4.7. Let $\phi : X \dashrightarrow Y$ be an S-rational map between finite-type smooth S-schemes, and suppose Y is separated over S. Let $S' \to S$ be a flat morphism, and let $\phi' : X' \dashrightarrow Y'$ denote the S'-rational base change. Then this is well-defined and

$$\mathrm{dom}(\phi') = \mathrm{dom}(\phi) \times_S S'.$$

Proof. We leave the reader to verify the well-definedness of the base change, and once we have that we visibly see that

$$\mathrm{dom}(\phi') \supseteq \mathrm{dom}(\phi) \times_S S'.$$

For the reverse inclusion, let Γ, Γ' denote the graphs of ϕ, ϕ' respectively. Since schematic closure commutes with flat base change we have

$$\Gamma' = \Gamma \times_S S'.$$

If $p : \Gamma \to X$ and $p' : \Gamma' \to X'$ denote the projections onto the X coordinates, and we pick out a point $x' \in \mathrm{dom}(\phi')$, and denote $x := f(x')$, we get a commutative diagram

$$\begin{array}{ccc} \mathrm{Spec}(\mathcal{O}_{X',x'}) & \longrightarrow & \mathrm{Spec}(\mathcal{O}_{X,x}) \\ \downarrow & & \downarrow \\ X' & \longrightarrow & X \end{array}$$

Note the horizontal top map is faithfully flat. Therefore p' being an isomorphism over $\mathrm{Spec}(\mathcal{O}_{X',x'})$ implies that it is an isomorphism over $\mathrm{Spec}(\mathcal{O}_{X,x})$. Finite-typeness of Y tells us that Γ must be finite-type over X. Thus we have an affine open neighborhood $W \ni x$ such that $p^{-1}(W) \to W$ is a closed immersion. Let $Z \subseteq X$ denotes the schematic image of this morphism, and U a quasicompact S-dense open of X where ϕ is defined (which we know exists by Lemma 8.2.4.4). Then

$$U \cap W \subseteq Z,$$

and moreover $U \cap W$ is S-dense in W. Therefore $Z = W$, and so $p^{-1}(W) \to W$ is an isomorphism, and so $x \in \mathrm{dom}(\phi)$. Since the x' chosen was arbitrary, the proposition follows. □

Now that we've built up the requisite background on S-rational morphisms, we assume the reader is fully happy with the statement of Definition 8.2.4.1. We're now ready to prove the key result.

PROPOSITION 8.2.4.8. Any weak Néron model X of X_K satisfies the WNMP (8.2.4.1).

Proof. Fix a smooth R-scheme Z with irreducible special fiber and a K-rational map

$$f_K : Z_K \dashrightarrow X_K,$$

we wish to construct an extension of f_K to Z.

Let $V_K \subseteq Z_K$ be an open dense subscheme on which f_K is defined, let $Y_K := Z_K - V_K$, and let Y denote the schematic closure of Y_K in Z. Since we are working over a spectrum of a DVR, we note that taking the schematic closure of any affine patch is torsion-free [worth citing a comm.alg fact here or no?], and thus R-flat. Thus Y itself is R-flat, giving us

$$\dim(Y_k) = \dim(Y) < \dim(Z_K) = \dim(Z_k).$$

Thus, $Z_k - Y_k$ is Zariski-dense inside the special fiber. It follows then that

$$V := Z - Y$$

is R-dense in Z. If we can find an R-rational map $V \dashrightarrow X$ extending the morphism $f_K : V_K \to X$, we will be done, since the domain of definition will be R-dense in Z as well. Thus WLOG we may replace Z with V and rename to reduce to the case where $f_K : Z_K \to X_K$ is a morphism.

Moreover, if we solve the problem locally on Z, then we can glue the resulting R-rational maps together, since separatedness of the weak Néron model X gives us unique R-rational extensions (see the discussion following Definition 8.2.4.5). Thus we reduce to the case where Z is of finite presentation.

Let $\Gamma \subseteq Z \times_R X$ denote the schematic closure of the graph of f, with projections p_Z, p_X onto the Z and X coordinates respectively. Our strategy will be to show that p_Z is R-birational (i.e. invertible on an R-dense open of Z), and then the map

$$\phi := p_X \circ p_Z^{-1} : Z \dashrightarrow X$$

will yield the desired result.

Define

$$T := p_Z(\Gamma) \subseteq Z.$$

Suppose we knew the special fiber T_k contained a dense open of Z_k. If η denotes the generic point of Z_k (recall we suppose Z has irreducible special fiber!), and we pick some $\xi \in p^{-1}(\eta) \subseteq \Gamma$, it would follow that $\mathcal{O}_{Z,\eta}$ is a dvr dominated by $\mathcal{O}_{\Gamma,\xi}$ (i.e the inclusion is a local map). But the induced fraction field extension is trivial, since p is a projection from (the schematic closure of) a graph, and so p_K is an isomorphism. Thus we have an isomorphism

$$\mathcal{O}_{Z,\eta} \xrightarrow{\sim} \mathcal{O}_{\Gamma,\xi}.$$

Part (b) of Lemma 8.1.2.3 then tells us that p_Z must be an isomorphism in an open neighborhood of η, and thus is R-birational.

So the problem boils down to showing that T_k does in fact contain a nonempty (hence dense) open of Z_k. First off recall by Chevalley's theorem that, since T is the image of a (closed!) graph under a

quasi-compact, finitely-presented morphism, we know T is constructible, and thus so is T_k. Hence, if we were to show T_k is dense in Z_k, it would imply it contains a dense open, so this is what we will show.

We claim that it suffices to check density after a faithfully flat base-change. Indeed, if $R \to R'$ is a faithfully flat ring extension and we denote by $Z', X', k', f', \Gamma', T'$ all the respective base changes, then we note that we have a natural inclusion

$$\Gamma_{f'} \subseteq \Gamma' \subseteq Z' \times_{R'} X',$$

since the base change of Γ remains closed and contains the graph of f', and thus contains the schematic closure of its graph. Denoting by $T_{f'}$ the image of $\Gamma_{f'}$ in Z' under the projection, we see that if

$$(T_{f'})_{k'} \subseteq Z'_{k'} = (Z_k)_{k'},$$

is dense, then $(T_k)_{k'}$ must also be dense, since it contains the former. Then, faithful flatness of $k \to k'$ allows us to descend density to deduce T_k is dense in Z_k.

Thus WLOG $R = R^{sh}$ and thus $k = k^{\text{sep}}$. This then allows us to use the extension property for étale points of the weak Néron model X. Recall the set of k-rational points is Zariski dense in a k-smooth scheme when k is separably closed [cite maybe], and so we have $Z_k(k) \subseteq Z_k$ is dense. Since R is henselian, we also know every $z_k \in Z_k(k)$ lifts to a point $z \in Z(R)$. Define

$$x_K := f_K(z_K) \in X_K(K).$$

Since X is a weak Néron model, x_K extends to some $x \in X(R)$. Thus (z, x) is in the R-closure of f_K in $Z \times_R X$, in other words

$$(z, x) \in \Gamma(R) \implies z \in T(R) \implies z_k \in T_k(k).$$

Thus we get

$$Z_k(k) \subseteq T_k(k) \subseteq T_k,$$

and so T_k is dense, completing our proof. □

> COROLLARY 8.2.4.9. Let S be any smooth R-scheme, and let η be a generic point of an irreducible component of the special fiber S_k of S. Set $R' := \mathcal{O}_{S,\eta}$ and $K' := \text{Frac}(R')$. Then if X is a weak Néron model of X_K, the base change $X_{R'}$ is a weak Néron model of $X_{K'}$.

Proof. Since R' is a finite R-algebra, the base change preserves smoothness, separatedness and finite-typeness. Thus it remains to show the canonical map

$$X_{R'}((R')^{sh}) \to X_{K'}((K')^{sh})$$

is a bijection. Since $(R')^{sh}$ (resp. $(K')^{sh}$) is built as filtered limits of étale extensions of R' (resp. their fraction fields), any $p \in X_{K'}((K')^{sh})$ descends down to some finite level K'', which the fraction field of some pointed étale neighborhood $(V, v) \to (U, \eta)$, with U being some shrunken down open of $\text{Spec}(R')$, and $R'' := \mathcal{O}_{V,v}$ being some finite étale extension of R'.

Note étaleness of $V \to U$ and R-smoothness of U implies V is R-smooth. Moreover, we know $v \in V_k$, since it lies over $\eta \in U_k \subseteq S_k$, and WLOG we can suppose V_k is irreducible. Finally, by shrinking the étale neighborhood of the descended point p down further, we can realize it as a K'-rational morphism $V_{K'} \dashrightarrow X_{K'}$, and composing it with the canonical $X_{K'} \to X_K$ we get a K-rational map $f_K : V_K \dashrightarrow X_K$.

Thus all the conditions of Proposition 8.2.4.8 are satisfied, so it follows we have an R-rational extension

$$f : V \dashrightarrow X,$$

defined in a neighborhood of $v \in V$. Restricting to $\mathcal{O}_{V,v} = \mathrm{Spec}(R'')$, we get an R-morphism

$$g : \mathrm{Spec}(R'') \to X,$$

If $s : \mathrm{Spec}(R'') \to \mathrm{Spec}(R') \to \mathrm{Spec}(R)$ denotes its structure morphism, then consider the map

$$(g, s) : \mathrm{Spec}(R'') \to X \times_R R' = X_{R'}.$$

tensoring with K', we get

$$\mathrm{Spec}(K'') \to X_K \times_K K' = X_{K'}$$

But it is not too hard to see this is precisely the descended point we begin with! In other words, we have construced a (unique, by separatedness!) extension of the K''-point to R''. Doing this for all étale extension yields the weak Néron model condition on strict henselizations. □

We end this section by remarking that a weak Néron model need not be unique at all.

EXAMPLE 8.2.4.10. [later].

Section 8.3

Birational Groups (Construction Steps III-IV)

§8.3.1 Invariant Differential Forms

Throughout this section, we denote by G a group scheme with multiplication, inversion and unit maps (denoted m, ι, ϵ respectively) over a base scheme S.

For any S-scheme T and any T-valued point $g : T \to G$, we get a T-valued point of the T-scheme G_T via

$$g_T := (g, \mathrm{id}_T) : T \to G \times_S T = G_T.$$

Definition 8.3.1.1. For any S-scheme T and $g \in G(T)$, we define **left translation by** g to be the map

$$L_g : G_T \xrightarrow{\sim} T \times_T G_T \xrightarrow{g_T \times \mathrm{id}_{G_T}} G_T \times_T G_T \xrightarrow{m_T} G_T,$$

and **right translation by** g to be

$$R_g : G_T \xrightarrow{\sim} G_T \times_T T \xrightarrow{\mathrm{id}_{G_T} \times g_T} G_T \times_T G_T \xrightarrow{m_T} G_T.$$

Sanity Check. Verify that both L_g and R_g are isomorphisms for any $g \in G(T)$.

For the sake of notational convenience, often we will drop the subscript T and just write

$$\begin{aligned} L_g : G &\to G \\ x &\mapsto gx, \end{aligned}$$

for example. Moreover in the special case that $T = G$ and $g = \mathrm{id}_G$, we call L_g, R_g the **universal** left and right translations, which we denote via

$$\begin{aligned} \mathcal{L} : G \times_S G &\to G \times_S G, \\ (x, y) &\mapsto (x, xy), \end{aligned} \qquad \begin{aligned} \mathcal{R} : G \times_S G &\to G \times_S G \\ (x, y) &\mapsto (xy, y). \end{aligned}$$

Each left and right translation for any T-valued point $g : T \to G$ is obtained from the universal left and right translations $\mathcal{L}$ and $\mathcal{R}$ via the base change $g : T \to G$.

Now take the module $\Omega^k_{G/S}$ of relative Kähler differentials of degree k. Since pullback maps are functorial, for any S-scheme T and $g \in G(T)$ we get an induced isomorphism by left translation

$$L_G^* \Omega^k_{G_T/T} \xrightarrow{\sim} \Omega^k_{G_T/T}.$$

Right translation induces a similar isomorphism R_G^*. Moreover we recall that the base-change projection morphism $p : G_T \to G$ induces an isomorphism

$$p^* \Omega^k_{G/S} \xrightarrow{\sim} \Omega^k_{G_T/T}.$$

DEFINITION 8.3.1.2. With the above notation, A relative differential k-form ω on G (i.e a global section of $\Omega^k_{G/S}$) is **left-invariant** if $L_g^*\omega_T = \omega_T$ for all $g \in G(T)$ and for all T (here $\omega_T := p^*(\omega) \in \Omega^k_{G_T/T}$). We define **right-invariant** differential forms similarly.

Since each translation on G_T is obtained via the universal translations $\mathcal{L}$ and $\mathcal{R}$, it suffies to check invariance of differential forms only for the universal translations. Thus, again for notational conveience, we often drop the T subscript to write ω instead of ω_T.

In the ensuing discussion we will frequently make use of the following key fact.

PROPOSITION 8.3.1.3. For G a group scheme over S, two global sections ω, ω' of any sheaf $\mathcal{F}$ on G are equal if

$$g_T^*\omega_T = g_T^*\omega'_T \in g_T^*\mathcal{F}_T$$

for every $g \in G(T)$. Moreover two sheaves $\mathcal{F}, \mathcal{G}$ are isomorphic if their restrictions to each $g \in G(T)$ are isomorphic.

Proof. Take the universal point $g = \mathrm{id}_G \in G(G)$, then we have the commutative diagram

$$\begin{array}{ccc} G & \xrightarrow{g_G = \mathrm{id}_G \times \mathrm{id}_G} & G \times_S G \\ & \searrow^{g} & \downarrow \\ & & G \end{array}$$

Thus the statements are trivially true for $T = G$; and they thus hold for any pullback along any $T \to G$. □

PROPOSITION 8.3.1.4. Let G be an S-group scheme with unit section $\epsilon : S \to G$. Then for each $\omega_0 \in \Gamma(S, \epsilon^*\Omega^k_{G/S})$ there exists a unique left-invariant differential form $\omega \in \Gamma(G, \Omega^k_{G/S})$ such that $\epsilon^*\omega = \omega_0$. The same assertion holds for right-invariance.

Proof. Naturally it suffices to prove the Proposition solely in the left-invariant case, the right invariant case follows by symmetry. Moreover, since pullbacks commute with exterior powers, we only have to show the proposition statement for one-forms, i.e. when $k = 1$. Moreover, the question is local on the base S, so we may assume ω_0 lifts to a section

$$\omega' \in \Omega^1_{G/S}(U),$$

where U is some neighborhood of the unit section $\epsilon : S \hookrightarrow G$. Recall we have a decomposition of the Kähler differential modules given by

$$\Omega^1_{G\times_S G/S} \cong p_1^*\Omega^1_{G/S} \oplus p_2^*\Omega^1_{G/S}.$$

Over $V := m^{-1}(U)$, then (recall m is the multiplication map), we have

$$m^*\omega' = \omega_1 \oplus \omega_2 \in p_1^*\Omega^1_{G/S}(V) \oplus p_2^*\Omega^1_{G/S}(V).$$

If

$$\delta = \iota \times_S \mathrm{id}_G : G \to G \times_S G$$

$$x \mapsto (x^{-1}, x)$$

denotes the *twisted diagonal* morphism, note that $\delta(G) \subseteq V := m^{-1}(U)$. Thus we may define the global section

$$\omega := \delta^* p_2^* \omega_2 \in \Omega^1_{G/S}(G).$$

We claim this is our desired left-invariant global one-form satisfying $\epsilon^*\omega = \omega_0$.

Firstly, to see left invariance, take some S-scheme T and an arbitrary $g \in G(T)$. We have the commutative diagram

$$\begin{array}{ccc} G & \xrightarrow{L_g} & G \\ {\scriptstyle\delta}\downarrow & & \downarrow{\scriptstyle\delta} \\ G \times G & \xrightarrow{R_{g^{-1}} \times L_g} & G \times G \end{array}$$

From this diagram we deduce that

$$L_g^* \omega = L_g^* \delta^* p_2^* \omega_2 = \delta^* (R_{g^{-1}} \times L_g)^* p_2^* \omega_2 \in \Omega^1_{G_T/T}.$$

Thus to show left invariance it will suffice to show

$$(R_{g^{-1}} \times L_g)^* p_2^* \omega_2 = p_2^* \omega_2.$$

Note that $m \circ (R_{g^{-1}} \times L_g) = m$, and moreover pulling back by $R_{g^{-1}} \times L_g$ and m both preserve the direct sum decomposition of the one-forms over $V := m^{-1}(U)$, so we have

$$\begin{aligned} \omega_1 \oplus \omega_2 &= m^* \omega' \\ &= (R_{g^{-1}} \times L_g)^* \circ m^* \omega' \\ &= (R_{g^{-1}} \times L_g)^* (\omega_1 \oplus \omega_2) \\ &= (R_{g^{-1}} \times L_g)^* \omega_1 \oplus (R_{g^{-1}} \times L_g)^* \omega_2. \end{aligned}$$

Since the decomposition is unique it follows $\omega_2 = (R_{g^{-1}} \times L_g)^* \omega_2$, as desired. This shows left-invariance of ω.

Next, we have to show that $\epsilon^*\omega = \omega_0$. Consider the morphism

$$\epsilon_G : G \to G \times_S G,$$

obtained by base changing the unit section $\epsilon : S \to G$. Here we think of the G we are base changing along to be in the second factor of $G \times_S G$. Thus, $\epsilon_G^* p_1^* \Omega^1_{G/S}$ is the zero map. Moreover, since $m \circ \epsilon_G = \mathrm{id}_G$, we have:

$$\epsilon_G^* \omega_2 = \epsilon_G^*(\omega_1 + \omega_2) = \epsilon_G^*(m^* \omega') = \mathrm{id}_G^* \, \omega' = \omega' \in \Omega^1_{G/S}.$$

Combining this deduction with the equality $\delta \circ \epsilon = \epsilon_G \circ \epsilon$ tells us

$$\begin{aligned} \epsilon^* \omega &= \epsilon^* \delta^* \omega_2 \\ &= \epsilon^* \epsilon_G^* \omega_2 \\ &= \epsilon^* \omega' = \omega_0 \in \Omega^1_{G/S}, \end{aligned}$$

as desired.

Finally, to show uniqueness, suppose we have global left-invariant one-forms ω, ω'such that

$$\epsilon^*\omega = \epsilon^*\omega' = \omega_0 \in \Gamma(S, \epsilon^*\Omega^1_{G/S}).$$

Then for any $g \in G(S)$, using the fact that $g = L_f \circ \epsilon$, we have

$$\begin{aligned} g^*\omega &= \epsilon^* \circ L_g^*\omega \\ &= \epsilon^*\omega = \epsilon^*\omega' \\ &= \epsilon^* \circ L_g^*\omega' \\ &= g^*\omega'. \end{aligned}$$

Thus ω and ω' take on the same value on all points of $G(S)$. This remains true after base-change, so ω and ω' coincide on the universal point $\mathrm{id}_G \in G(G)$. Thus by the same reasoning as in the proof of Proposition 8.3.1.3, they are the same section. This settles uniqueness. □

Writing $p : G \to S$ for the structure morphism of G, we can state the result of the above Proposition more elegantly.

Corollary 8.3.1.5. There exists canonical isomorphisms

$$p^*\epsilon^*\Omega^k_{G/S} \xrightarrow{\sim} \Omega^k_{G/S},\ k \in \mathbb{N},$$

obtained by extending sections in $\epsilon^*\Omega^k_{G/S}$ to left (or right) invariant sections in $\Omega^k_{G/S}$.

Observation. When G is a smooth group scheme of relative dimension d over a local scheme S, we know each $\Omega^k_{G/S}$ is locally free, and moreover that $\epsilon^*\Omega^k_{G/S}$ is free over $\mathcal{O}_S$. The above Corollary then tells us that each $\Omega^k_{G/S}$ is generated by $\binom{d}{k}$ left-invariant (or right-invariant) k-forms. In particular, the top-degree d-forms are generated by a single left-invariant $\omega \in \Omega^d_{G/S}$.

Proposition 8.3.1.6. Let G be a smooth S-group scheme of relative dimension d, and pick some left-invariant generator $\omega \in \Omega^d_{G/S}$. Then there exists a unique character $\chi : G \to \mathbb{G}_m$ such that

$$\mathrm{int}_g^*\,\omega_T = \chi(g)\omega_T$$

for every $g \in G(T)$, for every S-scheme T. (Here by int_g we mean the inner automorphism $G \to G$ induced by conjugation by g, i.e. $\mathrm{int}_g := L_g \circ R_g^{-1}$).

Proof. Since left and right translations commute, and moreover ω was assumed to be left-invariant, we have for any $h, g \in G(T)$ that

$$L_h^*(\mathrm{int}_g^*\,\omega_T) = L_h^* \circ R_{\in g}^* \circ L_g^*(\omega_T) = R_{g^{-1}}^* \circ L_h^* \circ L_g^*(\omega_T) = R_{g^{-1}}^*(\omega_T).$$

Since R_g is an isomorphism it follows $\omega'_T := R_{g^{-1}}^*(\omega_T)$ must also generate $\Omega^d_{G_T/T}$. Thus the pullbacks via the unit section

$$\epsilon_T^*\omega'_T, \epsilon_T^*\omega_T \in \epsilon^*\Omega^d_{G_T/T}$$

must also be generators. Thus they must differ by a well-defined unit $\chi(g) \in \mathcal{O}_T^*$. Pulling back via the

projection morphism $p_T : G_T \to T$ to get back ω'_T and ω_T (by Corollary 8.3.1.5), it follows

$$\mathrm{int}_g^* \omega_T = \chi(g)\omega_T$$

Since left/right translations and pullbacks are functorial it follows the map

$$\chi : g \mapsto \chi(g)$$

defines a group homomorphism $\chi : G \to \mathbb{G}_m$, as desired. □

Remark 8.3.1.7. In the case where G is a smooth group scheme over K the field of fractions of a DVR, and $G(K)$ is bounded, we can in fact say something stronger about the character discussed in the above Proposition. Note $\chi(G(K))$ would a bounded subgroup of $K^\times$, but such a subgroup can only consists of elements of valuation exactly 1!

Moreover if G is connected, it turns out χ is actually trivial! Take a maximal K-torus T; any non-trivial χ would then induce a surjection $T \twoheadrightarrow \mathbb{G}_m$, and thus T must contain a torus isogenous to $\mathbb{G}_m$. But then $G(K)$ cannot be bounded, since we saw in Example 8.1.3.2 that $\mathbb{G}_m$ is unbounded (and thus anything isogenous to it must be as well).

§8.3.2 Local Néron Models Exist (assuming Weil's Theorem)

Throughout this subsection, we denote by R a DVR with uniformizer π, field of fractions K, and residue field k. As before we also let R^{sh} denote the strict henselization of R and K^{sh} its field of fractions. Let X_K be a smooth finite-type d-dimensional K-group scheme (hence it is also separated by 1.1.2.5). Finally suppose $X_K(K^{sh})$ is bounded in X_K.

Then we know by Corollary 8.2.2.7 that X_K admits a weak Néron model X. Our goal will be to show the groupl law on X_K extends to an "R-birational group law" (to be defined later) on X. Our strategy will be to throw out certain irreducible components of the special fiber of X.

To define these, consider the unique-up-to-scaling left-invariant generating d-form $\omega \in \Omega^d_{X_K/K}$ (which we know exists by the observation following Corollary 8.3.1.5). Recall that when the local ring at a point $x \in X$ is a DVR with uniformizer π_x – for example when x is the generic point of an irreducible component of the special fiber – then we may define the **order** of a differential form to be the unique integer n such that $\pi_x^{-n}\omega$ generates $\Omega^d_{X_K/K,x}$. We denote this as

$$\mathrm{ord}_x(\omega).$$

Definition 8.3.2.1. Let $\{C_i\}_{i\in I}$ denote the irreducible components of the special fiber X_k of X, each with generic point η_i. For a global generating left-invariant $\omega \in \Omega^d_{X_K/K}$, we say C_i is **ω-minimal** if $\mathrm{ord}_{\eta_i}(\omega)$ is minimal amongst the orders of ω at all the generic points.

Sanity Check. Since the left-invariant form ω is unique up to scalar invariance, we see that the above definition is independent of choice of ω.

EXAMPLE 8.3.2.2. Suppose $X_K = \mathbb{G}_{m,K}$, and consider the following two R-models of it:

$$X_1 := \text{Spec}\left(R[z, z^{-1}]\right), \quad X_2 := \text{Spec}\left(R[z, z^{-1}, (z-1)/\pi]\right)$$

[complete]

LEMMA 8.3.2.3. Let X_i denote the open subscheme of the weak Néron model X where all irreducible components of the special fiber except C_i are thrown out, and let η_i denote the generic point of $(X_i)_k$. If there exists an R-rational map

$$\phi : X_i \dashrightarrow X_j$$

which is an isomorphism on generic fibers, then we know there is a unit $a \in \Gamma(X_K, \mathcal{O}_{X_K})^*$ satisfying $\phi_K^* \omega = a\omega$. Suppose there exists some $x_K \in X_K(K^{sh})$ that extends to a point in $x \in X_i(R^{sh})$, and suppose $|a(x_K)| = 1$. Then:

(a) $\text{ord}_{\eta_i}(\omega) \geq \text{ord}_{\eta_j}(\omega)$.

(b) The morphism $\text{dom}(\phi) \to X_j$ is an open immersion if and only if the above inequalit is an equality.

Proof. For ease of notation let

$$n_i, n_j := \text{ord}_{\eta_i}(\omega), \text{ord}_{\eta_j}(\omega).$$

We know $\pi^{-n_i}\omega$ is nonvanishing and regular over the generic fiber X_K, and by construction of n_i we know it is the same at the generic point of the irreducible fiber η_i. In other words, we know $\pi^{-n_i}\omega$ is defined everywhere on X except perhaps some points of codimension ≥ 2. Note that since $\Omega^d_{X_i/R}$ is a finite locally-free module, it is *reflexive*, and thus normality of X_i tells us (via [Stacks, § 31.13.13]) that this is equivalent to the existence of an open subscheme $\iota : U \hookrightarrow X_i$, whose complement has codimension ≥ 2, such that

$$\Omega^d_{X_i/R} \cong \iota_* \left(\Omega^d_{X_i/R}|_U\right).$$

But this is simply equivalent to saying that any global section defined on U extends all the way to a global section on X_i. Therefore, (applying identical reasoning for X_j as well) we know $\pi^{-n_i}\omega$ and $\pi^{-n_j}\omega$ generate $\Omega^d_{X_i/R}$ and $\Omega^d_{X_j/R}$ respectively, and consequently there must be a section $b \in \Gamma(X_i, \mathcal{O}_{X_i})$ such that

$$\phi^*(\pi^{-n_j}\omega) = b\pi^{-n_i}\omega$$

over X_i (again, technically b is only defined over $\text{dom}(\phi)$ since its complement is codimension ≥ 2 by the same reflexivity reasoning as before it extends to all of X_i). Recalling that we defined a such that $\phi_K^*(\omega) = a\omega$, substituting yields

$$a = \pi^{n_j - n_i} b$$

over X_K. By matching valuations (letting η denote the generic point of X_K) it follows that

$$n_j - n_i \geq \text{ord}_\eta(b) \geq 0,$$

proving assertion (a).

For assertion (b), recall that, since ϕ_K is an isomorphism, Zariski's Main Theorem tells us that ϕ is

etale on $\mathrm{dom}(\phi)$ if and only if it is an open immersion. Note ϕ is étale over its domain if and only if

$$\phi^*\Omega^d_{X_j/R} \to \Omega^d_{\mathrm{dom}(\phi)/R}$$

is bijection, i.e. if and only if the section we have entitled "b" in this proof is invertible over $\mathrm{dom}(\phi)$ (and hence X_i), which, as we have seen, is equivalent to $n_j - n_i = 0$. This completes our proof. □

DEFINITION 8.3.2.4. Using the notation of Lemma 8.3.2.3, we say two irreducible components of the special fiber C_i, C_j are **biequivalent** if the identity morphism on X_K extends to an R-rational map between the corresponding submodels X_i, X_j. If C_i (or equivalently C_j is ω-minimal we say X_i, X_j are **ω-minimal submodels** of X.

SANITY CHECK. Verify that this is indeed an equivalence relation on the set of ω-minimal irreducible components of our weak Néron model X, and moreover that any ω-minimal submodel is still a weak Néron model.

The key fact will be that if we throw out all irreducible components except for one from each of those equivalence classes, we will be left with an "R-birational group" X.

PROPOSITION 8.3.2.5. Let X^{weak} be a weak Néron model of a smooth finite-type d-dimensional group scheme X_K. Let $X \subseteq X^{weak}$ denote the subscheme obtained by removing all components C_i of the special fiber X_k, save for one representative from each biequivalence class of the ω-minimal components[a]. Then the group law m_K on X_K extends to an R-birational group law on X.

[a] Alternatively one may think of X as being constructed by gluing the corresponding submodels X_i along their common generic fiber.

Specifically by an R-birational group law, we meen an R-rational map

$$m : X \times_R X \dashrightarrow X$$

such that the universal translations

$$(x, y) \mapsto (x, m(x, y)), \quad (x, y) \mapsto (m(x, y), y)$$

are R-birational. Furthermore m is associative, in the sense that the usual diagram in the associativity axiom for group schemes commutes wherever the rational maps are defined.

LEMMA 8.3.2.6. Let X be a submodel of a weak Néron model X^{weak} of X_K, constructed by retaining only the irreducible components of X_k are ω-minimal. Let η be one of the generic points of its special fiber, and define $R' := \mathcal{O}_{X,\eta}$, and $K' := \mathrm{Frac}(R')$. Then every translation on $X_{K'}$ by a K'-valued point of $X_{K'}$ extends to an R'-birational endomorphism of $X_{R'}$, which is an open immersion on its domain of definition.

Proof. Firstly note that Corollary 8.2.4.9 tells us that the formation of weak Néron models is compatible with such a base change, i.e. $X^{weak}_{R'}$ will be a weak Néron model of $X_{K'}$. Moreover, any generic point ξ' of the irreducible fiber of $X_{R'}$ will lie over a generic point ξ of an ω-minimal component of X. By smoothness

we also know R' has the same uniformizer as R, and thus the valuation of any element pulled back along $\mathrm{Spec}(R') \to \mathrm{Spec}(R)$ remains unchanged. Thus,

$$\mathrm{ord}_\xi(\omega) = \mathrm{ord}_{\xi'}(\omega'),$$

where naturally ω' is the R'-pullback of ω. Thus, all irreducible components of $X_{R'}$ are ω'-minimal. All of this is to say it suffices to consider the case $R = R'$ and $K = K'$.

Now take

$$\tau_K : X_K \to X_K$$

to be translation by a K-valued point. Let $U \subseteq X$ be any ω-minimal submodel. Then the weak Néron mapping property of X^{weak} implies τ_K extends to an R-rational

$$\tau : U \dashrightarrow X^{weak}.$$

Since U is ω-minimal, we can use Lemma 8.3.2.3 to deduce that the image must land in some ω-minimal submodel $X_i \subseteq X^{weak}$ (note the fact that τ_K satisfies the necessary conditions for the lemma to apply is a consequence of Remark 8.3.1.7). In particular τ lands in X. Applying this over all such $U \subseteq X$ (using the fact they they must agree on interesections), we see that we get an R-rational extension

$$\tau : X \dashrightarrow X$$

that is open on its domain of definition (by part (b) of Lemma 8.3.2.3). The exact same logic applies for the inverse translaton τ_K^{-1}, giving us an inverse

$$\tau^{-1} : X \dashrightarrow X,$$

thus shows τ is R-birational, as desired. □

> EXERCISE 8.3.2.A. Show the result of Lemma 8.3.2.6 holds if we replace R' with the local ring $\mathcal{O}_{Z,\eta}$ of any smooth R-scheme Z at a generic point η of the special fiber Z_k.

Now we can apply the above lemma to prove the R-birational groupl law extension.

Proof of 8.3.2.5. Consider the universal left and right translations (recall from the discussion below Definition 8.3.1.1)

$$\mathcal{L}_K : (x, y) \mapsto (x, m_K(x, y)).$$

Let η be a generic point of an irreducible component of the special fiber X_k of X, and set $R_\eta := \mathcal{O}_{X,\eta}$ and $S_\eta := \mathrm{Spec}(R_\eta)$ for notational convenience. Note that S' has an X-scheme structure, let $p : S_\eta \to X$ be the structure morphism. Lemma 8.3.2.6 then tells us the translation τ_K obtained from $\mathcal{L}_K$ via the base change $S_{\eta,K} \to X_K$ extends to an S_η-birational map

$$\tau_\eta : S_\eta \times_R X \dashrightarrow S_\eta \times_R X.$$

If $\mathcal{L}$ denote the partially-defined universal translation on $X \times_R X$ induced by m_K, note we actually have the commutative diagram

$$\begin{array}{ccc} \tau_\eta : S_\eta \times_R X & \dashrightarrow & S_\eta \times_R X \\ \downarrow & & \downarrow \\ \mathcal{L} : X \times_R X & \dashrightarrow & X \times_R X \end{array}$$

Here the downward maps are given by (p, id_X). By Proposition 8.2.4.7, we see that the domain of Φ actually contains any generic points of the special fiber of $X \times_R X$ that project to η under the first projection. Applying this over all generic points of all irreducible components of X_K, it follows that $\mathcal{L}$ is in fact R-rational. Since each τ_η is S_η-birational, that $\mathcal{L}$ is R-birational. Identicall logic shows that the universal right translation $\mathcal{R}_K$ extends to an R-birational $\mathcal{R}$ on $X \times_R X$.

Let $W \subseteq X \times_R X$ be an R-dense open contained in $\mathrm{dom}(\mathcal{L}) \cap \mathrm{dom}(\mathcal{R})$ and also containing the generic fiber. Composing $\mathcal{L}$ with projection onto the second factor and $\mathcal{R}$ with projection onto the first factor, we obtain two extensions of m_K to morphisms $W \to X$ which must coincide. Thus m_K extends to an R-rational map on $X \times_R X$, and one easily verifies the relevant axioms from uniqueness of the extension of m_K. This completes our proof. □

There is a classical well-known result of Weil that say any R-birational group as constructed in Proposition 8.3.2.5 uniquely defines an R-group scheme.

THEOREM 8.3.2.7 (Weil). Let X_K be a smooth, finite-type K-group scheme with group law m_K. Suppose X is a smooth, separated R-model of X_K such that m_K extends to an R-birational group law m on X. Then there exists a smooth, separated, finite-type R-group scheme $\overline{X}$, containing X as an R-dense (see 8.2.4.2) open, extending m and having generic fiber X_K. Moreover $\overline{X}$ is determined uniquely up to isomorphism by these properties.

We will prove this theorem in the following subsection, it is a special case of Theorem 8.3.3.4. For now, granting the theorem, we will show that this $\overline{X}$ is in fact our desired Néron model of X_K; all that is really left to show is that it satisfies the Néron mapping property. We will just need one technicall lemma.

EXERCISE 8.3.2.B. Let Z be a normal noetherian scheme and $f : Z \dashrightarrow \mathrm{Spec}(A)$ a rational map into *any* affine scheme. Then $Z - \mathrm{dom}(f)$ has pure codimension 1, i.e. if f is defined in ≤ 1, it is defined everywhere. *[Hint: show this is a consequence of the standard commutative algebra fact that any noetherian normal domain is the intersection of all localizations at primes of height one.]*

From this we can deduce the following generalization.

PROPOSITION 8.3.2.8. Let S be a normal noetherian base scheme, and suppose $u : Z \dashrightarrow G$ be an S-rational map from smooth S-scheme Z to a smooth, separated S-group scheme G. Then if u is defined in codimension ≤ 1, it is defined everywhere.

Proof. Consider the rational map

$$\begin{aligned} v : Z \times_S Z &\to G \\ (z_1, z_2) &\mapsto u(z_1)u(z_2)^{-1}, \end{aligned}$$

and let U, V denote the domains of definition of u, v, respectively (note $U \times_S U \subseteq V$). Note that $v|_{V \cap \Delta}$ factors through the unit section $\epsilon : S \to G$ (here Δ denotes the diagonal of $Z \times_S Z$).

First we will show that V contain the diagonal Δ. In other words, setting $F := Z \times_S Z - V$, we want to show $F \cap \Delta = \emptyset$. Suppose forthe sake of contradiction we have some $x \in F \cap \Delta$, and let $s \in S$ be its image under the structure morphism. Let $H \ni \epsilon(s)$ be an affine neighborhood, and let $W \ni x$ be a neighborhood such that we have a rational map

$$v|_W : W \dashrightarrow H.$$

If V' denotes the domain of definition of the above map, we know $V' \subseteq V$. By normality of $Z \times_S Z$, hence of W, and affineness of H, we can apply Exercise **??** to deduce that

$$F' := W - V'$$

is of pure codimension 1 in W. Note that, if we identify Z with the diagonal Δ, we can wright

$$F' \cap \Delta \subseteq F \cap \Delta \subseteq Z - U,$$

and thus normality of Z tells us that

$$\operatorname{codim}_\Delta(F' \cap \Delta) \geq 2.$$

If we let d be the relative dimension of Z over S at x, it follows by smoothness that $\Delta \subseteq Z \times_S Z$ is defined locally by d equiations, and so we get (recalling we deduced F' was of pure codimension 1 in W)

$$\dim_x(F' \cap \Delta) \geq \dim_x(F') - d = \dim_x(Z \times_S Z) - 1 - d = \dim_x(\Delta) - 1,$$

i.e. $F' \cap D$ has codimension at most 1 in Δ. But this contradicts what we deduced earlier. Thus it follows such an x cannot exist, and so V must in fact contain the diagonal.

We will now show this fact then implies that $U = Z$. It suffices. By Proposition 8.2.4.6, it suffices to show there exists a faithfully flat morphism of S-schemes $f : Z' \to Z$ (with Z' smooth over S and locally of finite presentation) such that $u \circ f$ is defined everywhere.We let

$$Z' := V \cap (Z \times_S U) \subseteq Z \times_S Z,$$

then the projection onto the first coordinate yields a flat $f : Z' \to Z$. We wish to show f is surjective. Suppose $z : T \to Z$ be a geometric point, then we note the fiber product $T \times_S V$ is non-empty (here we view V as a Z-scheme via projection onto the first coordinate, and we deduce non-emptyness since we showed $V \supseteq \Delta$). Moreover, by S-density of $U = \operatorname{dom}(u) \subseteq Z$, we know

$$T \times_Z Z' = T \times_Z V \cap T \times_S U \neq \emptyset.$$

Thus every geometric fiber of f is nonempty, so f is surjective. Note that the morphism

$$\begin{aligned} V \cap (Z \times_S U) &\to G \\ (z_1, z_2) &\mapsto v(z_1, z_2)u(z_2) = u(z_1). \end{aligned}$$

This is visibly the map $u \circ f$. Therefore the S-rational map u is defined everywhere, completing the proof. □

We are now ready to show the R-group scheme $\overline{X}$ given by Theorem 8.3.2.7 satisfies the Néron mapping property; in particular give the proof of the main local existence theorem, which we restate here for convenience

THEOREM (§8.1.4.1). Let R be a DVR with fraction field K. Let X_K be a finite type smooth K-group scheme. Then X_K admits a Néron model X over R if and only if $X_K(K^{sh})$ is bounded in X_K.

Proof. We begin by recapping the key steps through the multiple subsections. By Corollary 8.1.4.4, we know boundedness is a necessary condition, so we wish to show it is sufficient. By Corollary 8.2.2.7, we know that boundedness implies the existence of a weak Néron model X' of X_K, and Weil's result on birational group scheme extnesions (8.3.2.7) tells us we can (uniquely!) enlarge this obtain a weak Néron model $X := \overline{X'}$ that is an R-group scheme extending the group law on m_K on X_K.

Thus it remains to show this X satisfies the Néron mapping property, i.e. for any R-smooth Z with a morphism of generic fibers $u_K : Z_K \to X_K$, there exists a unique extension $u : Z \to X$.

It is enough to treat the case where the special fiber Z_k is irreducible. Let η be its generic point and let $R' := \mathcal{O}_{Z,\eta}$ be its local ring. Consider the rational map

$$\begin{aligned} \phi : Z \times_R X &\dashrightarrow Z \times_R X \\ (z, x) &\mapsto (z, u_K(z)x) \end{aligned}$$

defined on the generic fiber. Exercise 8.3.2.A tells us that the base change along $\mathrm{Spec}(R') \to Z$ turns this into an R'-rational map $\phi_{R'}$. Then flatness of $X \to X'$ allows us to apply Proposition 8.2.4.6 to deduce that the map

$$\begin{aligned} \tau : Z \times_R X &\dashrightarrow X \\ (z, x) &\mapsto u_K(z)x \end{aligned}$$

is defined at all generic points of the special fiber sitting over $\eta \in Z$. Thus, τ is an R-rational morphism defined on the generic fiber, which is of codimension 1. Moreover, R-smoothness of Z tells us that $Z \times_R X$ is also normal. Thus Proposition 8.3.2.8 applies to tell us that τ is in fact defined everywhere.

Thus if $p : Z \to \mathrm{Spec}(R)$ denotes the structure morphism and $\epsilon : \mathrm{Spec}(R) \to X$ denotes the unit section, we see the map

$$u := \tau \circ (\mathrm{id}_Z, \epsilon \circ p) : Z \to Z \times_R X \to X$$

is an R-morphism $u : Z \to X$ extending u_K. For uniqueness, suppose $u' : Z \to X$ is another extension of u_K. Then the equalizer $E \to Z$ is a closed subscheme of Z (by separatedness of X) that contains Z_K. But R-flatness of Z tells us the injection $Z_K \hookrightarrow Z$ is schematically dense (apply the equivalence of flatness over a domain to torsion-freeness of modules with [Stacks, p. 29.7.5]). But any closed schematically dense subscheme must be all of Z, and thus $E = Z$ so $u = u'$, proving uniqueness. So X is indeed our desired Néron model of X_K. □

§8.3.3 Solving Birational Group Laws (proving Weil's Theorem)

The goal of this subsection is to prove Theorem 8.3.2.7. Throughout this subsection, let S be an arbitrary base scheme, and suppose that $X \to S$ is a smooth, separated, finite-type S-scheme that is faithfully flat over S.

DEFINITION 8.3.3.1. An **S-birational group law** on X is an S-rational map

$$\begin{aligned} m : X \times_S X &\dashrightarrow X \\ (x, y) &\mapsto xy \end{aligned}$$

such that

(i) The S-rational "universal left and right translations"

$$\begin{aligned} \mathcal{L} : X \times_S X &\dashrightarrow X \times_S X, \\ (x, y) &\mapsto (x, xy), \end{aligned} \qquad \begin{aligned} \mathcal{R} : X \times_S X &\dashrightarrow X \times_S X \\ (x, y) &\mapsto (xy, y), \end{aligned}$$

and

(ii) m is associative, i.e. $(xy)z = x(yz)$ (whenever both sides are defined, of course).

By our discussion in §8.2.4, we know the S-birational group law is compatible with flat base change. In particular, if $X \hookrightarrow \overline{X}$ is an inclusion of a dense open subscheme, and $\overline{X}$ is an S-group scheme, we get an S-birational group law on X. We want to work in the other direction.

DEFINITION 8.3.3.2. Let m be an S-birational group law on a smooth, separated, faithfully flat, finite-type S-scheme X. A **solution** of m is a smooth, separated, finite-type S-group scheme $\overline{X}$ with group law $\overline{m}$, together with an S-dense open $X' \subseteq X$ and an open immersion $X' \hookrightarrow \overline{X}$ such that

(i) The image of X' is S-dense in $\overline{X}$.

(ii) $\overline{m}$ restricts to m on X'.

The first claim is that solutions to such birational group laws are unique. As an intermediary step we leave the reader to prove (or recall) the following well-known lemma.

EXERCISE 8.3.3.A. Let G be a smooth S-group scheme, and let $U \subseteq G$ be an S-dense open. Then the morphism

$$\begin{aligned} U \times_S U &\to G \\ (x, y) &\mapsto xy \end{aligned}$$

is smooth and surjective.

PROPOSITION 8.3.3.3. Let m be an S-birational group law on a smooth separated finite-type faithfully flat S-scheme X. Then any solution of m is uniquely determined up to canonical isomorphism.

Proof. Let

$$\sigma_1 : X_1' \hookrightarrow \overline{X_1}, \quad \sigma_2 : X_2' \hookrightarrow \overline{X_2}$$

be solutions of the S-birational group law m on X. Set Y to be the S-dense open $X_1' \cap X_2' \subseteq X$; note that each $\sigma_i(Y)$ is S-dense in $\overline{X_i}$. Then the extended group laws $\overline{m_i}$ of $\overline{X_i}$ give rise to morphisms

$$\overline{m_i} \circ (\sigma_i \times \sigma_i) : Y \times_S Y \to \overline{X_i},$$

which are faithfully flat by Exercise 8.3.3.A. Moreover consider the following map:

$$\alpha := \sigma_2 \circ \sigma_1^{-1} : \overline{X_i} \dashrightarrow \overline{X_2},$$

which we visibly see to be compatible with the group laws, i.e

$$\overline{m_2} \circ (\sigma_2 \times \sigma_2) = \alpha \circ \overline{m_1} \circ (\sigma_1 \times \sigma_1).$$

Thus Proposition 8.2.4.6 tells us α is in fact defined everywhere. This, combined with compatibility with the group laws tells us α must be in fact be a homomorphism (of group schemes). Identical reasoning tells us $\beta := \sigma_1 \circ \sigma_2^{-1}$ is a group homomorphism as well, which is by construction inverse to α, thus they are both isomorphisms between $\overline{X_1}$ and $\overline{X_2}$, as desired. □

The main existence theorem we wanted to prove in this subsection (Theorem 8.3.2.7) hass the following convenient rephrasing in our new language.

THEOREM 8.3.3.4. Suppose $S = \mathrm{Spec}(R)$ is the spectrum of a field or DVR. Then for any m an S-birational group law on a smooth, separated, finite-type, faithfully flat S-scheme X, a solution to m exists.

REMARK 8.3.3.5. Actually the above theorem is true over an arbitrary base-scheme S; [BLR90a, §6.6] gives an argument (courtesy of Artin) using algebraic spaces. But of course for our purposes, the case of the spectrum of a field or dvr suffices.

EXERCISE 8.3.3.B. Suppose the theorem holds over a strictly henselian base. Setting $R' := R^{sh}$, we know that a solution $\overline{X'}$ to the base changed m' on $X' := X_{R'}$ exists, with respect to an R'-dense open $U' \subseteq X'$. Seting and $R'' := R' \times_R R'$ with projections p_1, p_2 onto the first and second coordinates respectively. Show that the canonical descent datum

$$p_1^* U' \to p_2^* U'$$

extends uniquely to a canonial descent datum on $\overline{X'}$ that satisfies the cocycle condition and is compatible with the group structure. Conclude using Theorem C.1.3.1 that a solution to m on X exists.

If X is an S-dense open of an S-group scheme $\overline{X}$, and T is an arbitrary S-scheme, the for every $x \in X(T)$, the set

$$\left\{ y \in X_T : xy, x^{-1}y, xy^{-1} \in X_T \right\}$$

is T-dense and open in X_T (as the intersection of 3 T-dense and open preimages of translations in $\overline{X_T}$). Thus we get an induced S-birational group law on X with the property in the following definition.

DEFINITION 8.3.3.6. An S-birational group law on X is **strict** if there exists an X-dense (hence also S-dense) open $U \subseteq X \times_S X$ on which m is defined, such that the universal translations are isomorphisms onto images which are X-dense with respect to *both* projections.

PROPOSITION 8.3.3.7. Let S be the spectrum of a field or a DVR, and let X be a smooth separated finite-type S-scheme with a birational group law m. Then there exists an S-dense open $X' \subseteq X$ such that m restricts to a strict group law on X'.

Proof. Let $U \subseteq \mathrm{dom}(m) \subseteq X \times_S X$ be such that the universal translations $\mathcal{L}_X, \mathcal{R}_X$ are open immersions on U. Set $V := \mathcal{L}_X(U), W := \mathcal{R}_X(U)$. Since U, V, W are all S-dense, it follows their intersection

$$Z := U \cap V \cap W \subseteq X \times_S X$$

is also S-dense. By [Stacks, §37.25.5], we know the set

$$T_1 := \{x \in X : Z \cap X_x \text{ is not dense in } X_x\}$$

is constructible in X (here of course we identify X_x with $x \times_S X \subseteq X \times_S X$). By S-density of Z, it follows the generic points of fibers of $X \to S$ are not in T_1. This then implies the closure $\overline{T_1} \subseteq X$ cannot be dense in any fiber of $X \to S$ either; if it were, then T_1 would be dense in an irreducible component of the fiber and must contain its generic point by constructibility, which we have shown cannot happen). Thus the open subscheme

$$\Omega_1 := X - \overline{T_1}$$

is S-dense in X, and moreover has the property that $Z \cap (\Omega_1 \times_S X)$ is Ω_1-dense in $\Omega_1 \times_S X$ (treating projection to the first coordinate as the structure morphism). By an analogous construction with respect to projection to the second coordinate, we construct a scheme Ω_2. Then setting

$$X' := \Omega_1 \cap \Omega_2 \subseteq X,$$

we get that X' is an S-dense open such that $Z \cap (X' \times_S X')$ is X'-dense with respect to either projection. We claim this X' has the desired properties, i.e. if we set

$$\begin{aligned} U' &:= U \cap (X' \times_S X') \cap m^{-1}(X'), \\ V' &:= \mathcal{L}_{X'}(U'), \\ W' &:= \mathcal{R}_{X'}(U'), \end{aligned}$$

it remains to show that these are all X'-dense, with respect to both projections $p_1, p_2 : X' \times_S X' \to X'$. First we show U' is X'-dense with respect to p_1. Take any $a \in X'$, note for the purposes of demonstrating X'-density we may just treat the base scheme S as a field and a as an S-valued point of X', so the fiber product of X and a over S makes sense. Then base changing along $a \to X$ turns $\mathcal{L}_X$ into

$$\mathcal{L}_X(a, \cdot) : U \cap (a \times_S X) \xrightarrow{\sim} V \cap (a \times_S X),$$

which is an isomorphism with dense image in $a \times_S X$. Thus the open subscheme

$$\mathcal{L}_X(a, \cdot)^{-1}(V \cap (a \times_S X)) = m^{-1}(X') \cap (a \times_S X)$$

is also dense in $a \times_S X$. This proves X'-density of U' with respect to p_1. A symmetric argument using W and $\mathcal{R}_X$ shows X'-density with respect to p_2. Thus we move on to V' and W'.

For any $a \in X'$, we note that $V' \cap (a \times_S X')$ is the image of the dense open $U' \cap (a \times_S X') \subseteq a \times_S X'$ under the open immersion $\mathcal{L}_{X'}(a, \cdot)$, and thus thus dense. This shows X' of V' with respect to p_1. Symmetrically

we deduce X'-density of W' with respect to p_2.

For what remains, consider the isomorphisms

$$\mathcal{L}_X : m^{-1}(a) \xrightarrow{\sim} W \cap (a \times_S X) =: W_a$$
$$\mathcal{R}_X : m^{-1}(a) \xrightarrow{\sim} V \cap (X \times_S a) =: V_a.$$

Since $a \in X'$, we know by construction that $Z \cap (X \times_S a)$ is dense in $X \times_S a$, and thus it follows V_a must be dense too. Similarly it follows that W_a is dense in $a \times_S X$. Restricting to X' also preserves density, and thus it follows that

$$m^{-1}(a) \cap U' = \mathcal{L}_{X'}^{-1}(V_a \cap (X' \times_S a)) \cap \mathcal{R}_{X'}^{-1}(W_a \cap (a \times_S X))$$

is open and dense in $m^{-1}(a)$. Taking its image under $\mathcal{R}_{X'}$, we get

$$\mathcal{R}_{X'}(m^{-1}(a) \cap U') = \mathcal{R}_{X'}(m^{-1}(a)) \cap \mathcal{R}_{X'}(U') = (a \times_S X') \cap W'$$

must be dense in $a \times_S X'$. This shows X'-density of W' with respect to p_1. A symmetric argument shows X'-density of V' with respect to p_2. This completes our proof. □

The upshot of the above Proposition is that it reduces the proof of Theorem 8.3.3.4 to showing that we can enlarge a strict group law on X (over a strictly henselian base, which is no loss of generality due to Exercise 8.3.3.B) to a group scheme $\overline{X}$. The big-picture idea of the proof is relatively intuitive, though the minute details get rather technical. The gist is as follows.

INTUITION. Picking some $a \in X(S)$, we can consider gluing two copies of X together along the translation map $x \mapsto ax$. To obtain an enlargement X_1 to which the group law m also extends. The (S-rational!) translation map $\tau_a : x \mapsto ax$ will then be defined on the open subscheme $X \subseteq X_1$. Repeating this construction again we get $X_2 \subseteq X_3 \subseteq \ldots$ such that $\tau_a : X \dashrightarrow X_i$ makes sense, and applying a noetherian argument we show this process stabilizes at some X_n. We will then show this implies $\tau_a : X \to X_n$ is defined everywhere on X, and moreover that this will imply multiplication $m : X \times_S X \to X_n \times_S X_n$ will be defined everywhere. From this we will show that the multiplication law actually extends to all of X_n, and X_n will be our desired group scheme.

The reason for this technical difficulties, as we will see, essentially boils down to the problem that for any $a \in X$, the multiplication map $x \mapsto ax$ is only defined for "most" $x \in X$, and moreover precisely which x for which this is defined is dependent on a. So there will be a lot of massaging to ensure that all the steps described above can be done in a well-defined manner. In lieu of this we record here for convenience the notation and assumptions that will be used for the remainder of the subsection.

SETUP 8.3.3.8 (S, X, U). We set S to be the spectrum of a separably closed field or a strictly henselian DVR, and let X be a smooth, separated, faithfully flat finite-type S-scheme. Let m denote a strict group law on X, with S-rational left and right universal translations denoted by $\mathcal{L}_X$ and $\mathcal{R}_X$ respectively. Let U be an X-dense open on which m is defined, and let $V := \mathcal{L}_X(U), W := \mathcal{R}_X(U)$ be its X-dense images (with respect to both projections).

We set

$$\mathscr{R}_{X/S} : \underline{\mathrm{Sch/S}} \to \mathrm{Set}$$

to be the functor which associates to each S-scheme T the set of T-birational endomorphisms of X_T. For any $a \in X(T)$, let

$$\tau_a : T \times_S X \dashrightarrow T \times_S X$$

be the base change of the left translation $\mathcal{L}_X$ via the base change $a : T \to X$.

Furthermore, we denote by X^n the n-fold fiber product of X over S, and for $r \leq n$ and a set of integers $1 \leq i_1, \ldots, i_r \leq n$, we let

$$p_{i_1,\ldots,i_r} : X^n \to X^r$$

denote projection onto the product of the factors with indices $i_1, \ldots, i_r$. By $p_{i_1,\ldots,i_r}$-density of $U \subseteq X^n$ we will mean X^r-density of U with respect to the morphism $p_{i_1,\ldots,i_r}$.

We now present a series of technical lemmas assuming the above hypothesis, concluding the subsection with the construction of $\overline{X}$ in Theorem 8.3.3.4 making use of all the lemmas.

LEMMA 8.3.3.9. The morphism

$$X(T) \to \mathscr{R}_{X/S}(T)$$
$$a \mapsto \tau_a$$

for every S-scheme T defines a monomorphism of functors (here of course identifying X with its functor of points) that "resepcts the laws of composition on X", meaning for any $a, b, c \in X(T)$ satisfying $m(a, b) = c$ also satisfies $\tau_a \circ \tau_b = \tau_c$.

Proof. The fact that this map obeys the laws of composition is a direct consequence of associativity of the group multiplication law m. So we have to show the monomorphism assertion, i.e. that every $X(T) \to \mathscr{R}_{X/S}(T)$ is an injection. Suppose we have $a, b \in X(T)$ such that $\tau_a = \tau_b$, we want to show $a = b$.

By applying the base change $T \to S$, we may WLOG consider the case $T = S$. By X-density of U with respect to p_1, we see the compositions

$$\mathcal{R}_a : S \times_S X \xrightarrow{a \times \mathrm{id}_X} X \times_S X \overset{\mathcal{R}_X}{\dashrightarrow} X \times_S X$$
$$\mathcal{R}_b : S \times_S X \xrightarrow{b \times \mathrm{id}_X} X \times_S X \overset{\mathcal{R}_X}{\dashrightarrow} X \times_S X$$

are well-defined S-rational maps. Note, since $\mathcal{R}_a = (\tau_a, \mathrm{id}_X)$ and $\mathcal{R}_b = (\tau_b, \mathrm{id}_X)$, that $\tau_a = \tau_b$ implies that $\mathcal{R}_a = \mathcal{R}_b$ as S-rational morphisms. Since $\mathcal{R}_X$ is an open immersion on U, we must have that $a \times \mathrm{id}_X$ and $b \times \mathrm{id}_X$ coincide on the S-dense open subscheme

$$X' := (a \times \mathrm{id}_X)^{-1}(U) \cap (b \times \mathrm{id}_X)^{-1}(U) \subseteq S \times_S X,$$

hence on all of $S \times_S X$, as S-rational maps. In particular the first components must agree, so $a = b$, as desired. □

LEMMA 8.3.3.10. With the notation of 8.3.3.8, let Ω be the set

$$\Omega := \left\{(x, y, z, w) \in X^4 : (z, w), (y, w), (x, yw) \in U\right\} \subseteq X^4.$$

Then Ω is p_{123}-dense in X^4 (another way of saying this is the set $\{w \in X : (x, y, z, w) \in \Omega\}$ is dense

in X for all $(x, y, z) \in X^3$.

Proof. Note U is p_1-dense in X^2, so the conditions that $(z, w) \in U$ and $(y, w) \in U$ are both open p_{123}-dense conditions. Thus it suffices to check that the set

$$\{w \in X : (y, w) \in U, (x, yw) \in U\} \subseteq$$

is dense. We already know the condition $(y, w) \in U$ cuts out a dense subset of $U_y \subseteq X$. Note translation $\mathcal{L}_X : (y, w) \mapsto (y, yw)$ has a p_1-dense image (namely, V) in X^2. Then the restriction that $(x, yw) \in U$ means we intersect U_y with V_x, another dense open, so we get the desired density. Openness is clear, since this is the intersection of 3 open conditions. □

LEMMA 8.3.3.11. With the notation of 8.3.3.8, let $\Gamma \subseteq X^3$ be the schematic closure of the graph of $m : U \to X$. Then for any S-scheme T, then any $(a, b, c) \in \Gamma(T)$ satisfies $\tau_a \circ \tau_b = \tau_c$.

Proof. Let Ω be the p_{123}-dense open subscheme given by Lemma 8.3.3.10. Then the S-rational maps

$$\lambda : X^4 \dashrightarrow X^4, \qquad \mu : X^4 \dashrightarrow X^4$$
$$(x, y, z, w) \mapsto (x, y, x(yw), w), \qquad (x, y, z, w) \mapsto (x, y, zw, w)$$

are defined on Ω. Letting $\Omega' := \Omega \cap p_{12}^{-1}(U)$, we claim that $Z' := \Omega' \cap (\Gamma \times_S X)$ is schematically dense in $Z := \Omega \cap (\Gamma \times_S X)$. Indeed, this is visibly true for $p_{12}^{-1}(U) \cap (\Gamma \times_S X)$ (recall X is S-flat), and this remains dense when intersecting both sets with an open subscheme of X^4, in particular Ω.

By associativity of m, we know the morphism $\mu|_{Z'}$ factors through Λ, the schematic image of $\lambda|_\Omega$. By continuity, we also know $\mu|_Z$ factors through Λ. Thus we get a morphism

$$\mu : Z \to \Lambda.$$

Now set

$$\phi := (a, b, c) \times \mathrm{id}_X : T \times_S X \to X^4, \quad \Omega_{a,b,c} := \phi^{-1}(\Omega).$$

Then $\Omega_{a,b,c}$ is T-dense and open in X_T. Let ϕ_T be ϕ turned into a T-morphismm and let μ_T be the base change of μ along $T \to S$. Then we see that

$$(p_3 \circ \mu_T \circ \phi_T)|_{\Omega_{a,b,c}} = \tau_c|_{\Omega_{a,b,c}}.$$

But this also coincides with $\tau_a \circ \tau_b$ on $\Omega_{a,b,c}$, since $\mu \circ \phi$ factors through Λ. This yields the desired equality. □

An important consequence of the above Lemma is the following.

COROLLARY 8.3.3.12. With the notation of 8.3.3.8, let $\Gamma \subseteq X^3$ be the schematic closure of the graph of $m : U \to X$, and let $q_{ij} : \Gamma \to X^2$ be the morphisms induced by the projections $p_{ij} : X^3 \to X^2$. Then each q_{ij} is an open immersion and has an image which is p_1 and p_2-dense in X^2.

Proof. First we show injectivity of the q_{ij}. By Lemma 8.3.3.11, we know

$$\tau_a \circ \tau_b = \tau_c \ \forall (a, b, c) \in \Gamma(T) \ \forall T \in \underline{\mathrm{Sch}}\,/S.$$

Since this is an identity in the group $\mathscr{R}_{X/S}(T)$, any two of those maps determine the third. Moreover, we know from Lemma 8.3.3.9 that the natural transformation $X \to \mathscr{R}_{X/S}$ is a monomorphism. Thus, a point of Γ is determined by any two of its components, and in particular this means all the q_{ij} must be injective. In particular we see by construction that we have the following isomorphisms:

$$q_{12}^{-1}(U) \xrightarrow{\sim} U$$
$$q_{13}^{-1}(V) \xrightarrow{\sim} V$$
$$q_{23}^{-1}(W) \xrightarrow{\sim} W.$$

Thus Zariski's Main Theorem applies (recall we are working over a normal base S) to tell us that each q_{ij} is an open immersion. Then X-density of $U, V, W \subseteq X^2$, the image of each q_{ij} is X-dense (with respect to both projections!) in X^2. □

If we fix points $a, b, c \in X(T)$, the above Corollary tells us that there exists at most one point $x \in X(T)$ such that $ax = c$. Suggestively we could write $a^{-1}c := x$. Of course there are other variations of this with respect to other projections. But the most important takeaway is the following.

KEY IDEA. Corollary 8.3.3.12 implies that multiplication $m : X^2 \dashrightarrow X$ is defined at some $(x, y) \in X^2$ if the fiber $q_{12}^{-1}((x, y))$ is nonempty.

> LEMMA 8.3.3.13. With the notation of 8.3.3.8, for any $a, b \in X(S)$, then, $a \times_S b$ as a point of X^2, the multiplication law $m : X^2 \dashrightarrow X$ is defined at $a \times_S b$ if and only if the birational map $\tau_a : X \dashrightarrow X$ is defined at b.

Proof. One direction is clear, so for the other direction suppose $\tau_a : X \dashrightarrow X$ is defined at b. Let $\Gamma \subseteq X^3$ be the schematic closure of the graph of the multiplication map $m : U \dashrightarrow X$. Consider the S-dense open subscheme

$$U_a := U \cap (a \times_S X),$$

we know τ_a is defined on (at least) U_a. Let Γ_a be the schematic closure in X^2 of the graph of $\tau_a|_{U_a}$. Then we have

$$(a \times_S \Gamma_a) \cap (a \times_s U_a \times_S X) \subseteq \Gamma.$$

By continuity, we also know that $a \times_S \Gamma_a \subseteq \Gamma$. Note the image of the morphism

$$a \times_S \Gamma_a \hookrightarrow \Gamma \overset{q_{12}}{\to} X^2$$

contains the point $a \times_S b$, we know the fiber over it with respect to q_{12} is nonempty. By the Key Idea above the Lemma statement tells us this means that m is defined at $a \times_S b$, as desired. □

The upshot of the above Lemma is that if we want to expand the domain of definition of a birational group law $m : X^2 \dashrightarrow X$, then it suffices to enlarge the domain of definition of $\tau_a : X \dashrightarrow X$ for various points $a \in X(S)$. Armed with this insight, we introduce the following construction (intuitively, we make a "translated" copy of X by a and glue it to X).

Construction 8.3.3.14. With notation as in Setup 8.3.3.8, fix some $a \in X(S)$ and, as in the proof of Lemma 8.3.3.13, consider the schematic closure $\Gamma_a \subseteq X^2$ of τ_a. Then Corollary 8.3.3.12 tells us that the projections $p_i : \Gamma_a \to X$ are injective. Thus Zariski's Main Theorem tells us that p_1, p_2 are open immersions, moreover they have S-dense images in X. Thus, the projections define gluing data, and we obtain the S-scheme

$$X' := X \cup_{\Gamma_a} X.$$

This is of course smooth and finite-type over S, as it's covered by two S-dense subschemes isomorphic to X, and since Γ_a is closed in X^2 by construction, we see that X' is S-separated.

In order to distinguish the two copies of X sitting in X' we write the gluing data as

$$p_1 : \Gamma_a \to X(a), \quad p_2 : \Gamma_a \to X.$$

Here $X(a) \cong X$. The reason for this notation is that the birational map $\tau_a : X \dashrightarrow X$ extends to an isomorphism (namely, the canonical one) $\tau_a : X \xrightarrow{\sim} X(a)$. In particular, τ_a can be interpreted as an S-birational map $X' \dashrightarrow X'$.

Lemma 8.3.3.15. With the notation of Construction 8.3.3.14, the enlarged scheme X' contains the strict group law on m extends to a strict group law m' on X'.

Proof. By S-density of X in X', we know m extends to an S-birational group law m' on X, so it only remains to show it is strict, i.e. to find some $U' \subseteq X' \times_S X'$ which is X'-dense with respect to both projections and such that the universal translations are opne immersions with X'-dense images.

Take U as in Setup 8.3.3.8, and let U_1 be the image of U under the isomorphism

$$\tau_a \times \mathrm{id}_X : X \times_S X \xrightarrow{\sim} X(a) \times_S X.$$

Then m' is defined on U, and the isomorphism $\tau_a : X \xrightarrow{\sim} X(a)$ can be used to obtain the morphism

$$U_1 \to X(a)$$
$$(\tau_a(x), y) \mapsto \tau_a(xy).$$

Both this map and $m : U \to X$ coincide on an S-dense open of U, by associativity of m, so both of these maps make up the extension m'. Thus m' is defined on the open subscheme $U \cup U_1$, which we know by construction is p_1-dense in $X' \times_S X'$.

Now consider the open subscheme

$$\Omega := \left\{(x, y, z) \in X^3 : (x, y) \in U, (xy, z) \in U\right\} \subseteq X^3.$$

By similar reasoning to that of Lemma 8.3.3.10, we see that Ω is p_{23}-dense in X^3. Consider the image of Ω under the isomorphisms

$$(\mathrm{id}_X \times \tau_a) \circ p_{13} : X \times_S a \times_S X \xrightarrow{\sim} X^2 \xrightarrow{\sim} X \times_S X(a).$$

Let this image be U_2, it is an open subscheme which is $X(a)$-dense with respect to the second projection. Then the morphism

$$U_2 \to X$$

$$(x, \tau_a(y)) \mapsto (xa)y$$

is defined, and moreover coincides with multiplication $m : U \to X$ on an S-dense open by associativity of m. Thus writing

$$U' := U \cup U_1 \cup U_2,$$

we get an extension of m to m' on U', which is X'-dense with respect to both projections.

It remains to verify that the image of U' under the left and right universal translations are X'-dense themselves (with respect to both projections of course). First note that $\mathcal{L}_{X'}$ and $\mathcal{R}_{X'}$ are extensions of $\mathcal{L}_X$ and $\mathcal{R}_X$, since m' is an extension of m. Since $\mathcal{L}_X$ and $\mathcal{R}_X$ are open immersions on U by construction we must have that $\mathcal{L}_{X'}$ and $\mathcal{R}_{X'}$ are open immersions (in particular quasi-finite) on each of U, U_1, U_2 constructed above. Since these are all S-birational maps, Zariski's Main Theorem applies to show that they are open immersions on the union U'.

Let V_1 be the image of V under the isomorphism

$$\tau_a \times \tau_a : X \times_S X \xrightarrow{\sim} X(a) \times_S X(a).$$

Then $V' := \mathcal{L}_{X'}(U')$ contains $V \cup V_1$, and the latter is X'-dense in $X' \times_S X'$ with respect to both projections. A similar argument shows that $W' := \mathcal{R}_{X'}(U')$ is X'-dense with respect to the first projection. For X' density with respect to the other projection, we see that $W_1 := \mathcal{R}_{X'}(U_1)$ is X-dense with respect to the second projection. Moreover consider

$$W_2 := \mathcal{R}_{X'}(U_2) \subseteq X \times_S X(a),$$

which we see is $X(a)$-dense in $X' \times_S X(a)$ with respect to the second projection. Thus W' is X'-dense with respect to both projections, as desired. □

The above Lemma, combined with Construction 8.3.3.14, gives us a way to enlarge a strict birational group law to another strict birational group law. Thus we have the following construction.

> CONSTRUCTION 8.3.3.16. Let the notation be as in Setup 8.3.3.8. Let $a_0, a_1, \ldots,$ be a sequence of S-valued points of X. Defining $X_0 := X$ and, inductively,
>
> $$X_{n+1} := X_n \cup_{\Gamma_{a_n}} X_n,$$
>
> as in Construction 8.3.3.14. Note each X_i contains X as an X-dense open subscheme, and Lemma 8.3.3.15 tells us $m_0 := m$ extends to a strict group law m_i on each X_i. We remark that the notation of the X_i does not include the data of the choice of a_i; this should cause no confusion.

We want to use some sort of Noetherian argument to show the sequence $X = X_0 \subseteq X_1 \subseteq \ldots$ eventually stabilizes, and then, for a certain choice of sequence a_i, the end result yields the S-group scheme we're looking for.

> PROPOSITION 8.3.3.17. In the notation of 8.3.3.16, there exists a finite sequence $a_1, \ldots, a_n \in X(S)$ such that the S-rational group law m extends to an S-morphism $X \times_S X \to X_n$.

Proof. First. we show that we can find $a_1, \ldots a_n$ such that for any $a \in X(S)$ the (a priori only) S-rational translation $\tau_a : X \dashrightarrow X$ extends to a morphism $\tau_a : X \to X_n$. Assume FTSOC that this is not possible,

then this would imply the existence of an infinite sequence $a_1, a_2, \ldots$, such that

$$\tau_{a_{n+1}} : X \dashrightarrow X_n$$

is not defined everywhere for every n. Let $\Gamma_i \subseteq (X_i)^3$ be the schematic closure of the graph of the original multiplication law $m : U \to X$; this coincides with the schematic closure of m_i. Thus Corollary 8.3.3.12 tells us that the projection

$$p_{12} : \Gamma_i \to X_i$$

is an open immersion. Setting

$$\Omega_i := p_{12}(\Gamma_i) \cap (X \times_S X),$$

we see that the Ω_i form an increasing sequence of open subschemes. However, note $X \times_S X$ is *noetherian*, so this sequence stabilizes at some Ω_n. We claim that for $a := a_{n+1}$, the map

$$\tau_a : X \dashrightarrow X_n$$

is indeed defined everywhere, which would yields the desired contradicion. We know by definition the birational map

$$\tau_a : X \dashrightarrow X_{n+1}$$

is defined everywhere on its source (see the end of Construction 8.3.3.14). Thus Lemma 8.3.3.13 tells us that m_{n+1} on X_{n+1} is defined at $a \times_S b$ for any $b \in X$. Thus the fiber over $a \times_S b$ of

$$p_{12} : \Gamma_{n+1} \to X_{n+1} \times_S X_{n+1}$$

is nonempty, and thus $a \times_S b \in \Omega_{n+1} = \Omega_n$. But this then tells us the fiber over $a \times_S b$ of

$$p_{12} : \Gamma_n \to X_n \times_S X_n$$

must be nonempty, and thus (from the Key Idea following Corollary 8.3.3.12) it follows m_n is defined at $a \times_S b$. In particular $\tau_a : X \dashrightarrow X_n$ is defined at b. Since $b \in X$ was arbitary, it follows τ_a is in fact defined everywhere. Thus we obtain our contradiction, so it follows we can in fact pick some $a_1, \ldots a_n \in X(S)$ such that $\tau_a : X \to X_n$ is a morphism for every $a \in X(S)$.

It remains to show that the multiplication map $m : X \times_S X \dashrightarrow X_n$ is defined everywhere. By Lemma 8.3.3.13 we know it is defined on $a \times_S X$ for every $a \in X(S)$, but this is not enough to be a morphism everywhere in general. This is where we will use the assumption that our base is *strictly henselian* (or a field).

For any $t \in S$ (so t is either the generic or special point), let C_t denote the reduced subscheme of X_t whose underlying topological space is the closure of the set of points $\{a_t : a \in X(S)\}$. Then, as a consequence of Lemma 8.3.3.18 below, there exists an irreducible component X_t^0 of X_t contained C_t. moreover, if $k'/\kappa(t)$ is any field extension, and let $t' := \mathrm{Spec}(k')$. Then $C_{t'}$ coincides exactly with the reduced subscheme of $X_{t'}$ whose underlying topological space is the closure of the points $\{a_{t'} : a \in X(S)\}$; this is a consequence of [EGA, IV–3, §11.10.7]. In particular, if

$$Z_{t'} \subseteq X_{t'} = X \times_S t'$$

is a dense open subscheme, then there exists a point $a \in X(S)$ such that $a \times_S t'$ is a point in $Z_{t'}$.

We now resume with the proof. By Corollary 8.3.3.12, we know the map

$$\begin{aligned} q_{23} \circ q_{13}^{-1} : X \times_S X &\dashrightarrow X \times_S X \\ (w, x) &\mapsto (w^{-1}x, x) \end{aligned}$$

is an S-birational map (here of course by "$w^{-1}x$" we mean the unique element z such that $wz = x$). Let D be the domain of this map, then the Corollary also tells us D and its image are X-dense in X^2 with respect to both projections. Now take any $t \in X^2$. Again using methods similar to the proof of Lemma 8.3.3.10, we see that the set

$$Z := \left\{(w, x, y) \in X^3 : (w, x) \in D, (w^{-1}x, y) \in U\right\}$$

is open and p_{23}-dense in X^3, and thus open and dense in $X_t = X \times_S t$. Thus the above discussion about dense fibers over a strictly henselian base tells us that there must exist some $a \in X(S)$ such that $a \times_S t \in Z$. Then the S-rational map

$$\begin{aligned} X \times_S X &\dashrightarrow X \\ (x, y) &\mapsto (a^{-1}x)y \end{aligned}$$

is defined at t. Moreover, since we have already shown left translation $\tau_a : X \to X_n$ is defined everywhere, we see that

$$\begin{aligned} X \times_S X &\dashrightarrow X_n \\ (x, y) &\mapsto a((a^{-1}x)y) \end{aligned}$$

is defined at t. However, associativity of the strict group law tells us this map coincides exactly with m on its domain of definition. Thus, this is an extension of m to a morphism

$$X^2 \to X_n,$$

as desired. □

LEMMA 8.3.3.18. Let T be the spectrum of a field or a strictly henselian DVR, and let $Y \to T$ be a smooth surjective finite-type morphism. Then for both $t \in T$, there exists a connected component Y_t^0 of the fiber Y_t such that the set of points $\{a(t) : a \in Y(T)\}$ is dense in Y_t^0.

Proof. [later, or just cite somewhere] □

We are now (finally!) primed to prove Weil's theorem of finding a solution to a birational group.

Proof of 8.3.3.4. To recap, Exercise 8.3.3.B and Proposition 8.3.3.7 allows us to reduce our situation to the assumptions and notation as in Setup 8.3.3.8. Proposition 8.3.3.17 then tells us that we may construct some smooth, separated, finite-type S-scheme $\overline{X}$ that contains X as an S-dense open subscheme, and a construction of the strict group law m on X extends to a strict group law $\overline{m}$ on $\overline{X}$, whose domain of definition contains $X \times_S X$. We claim we can extend $\overline{m}$ to a group multiplication law on $\overline{X} \times_S \overline{X}$.

Recalling that Proposition 8.2.4.7 tells us that the domain of definition is compatible with flat base

change, we see it suffices to show that, for any $(b, c) \in \overline{X}^2$, the map

$$\overline{m}_X := \mathrm{id}_X \times \overline{m} : X \times_S \overline{X}^2 \dashrightarrow X \times_S \overline{X}$$

is defined at some (i.e. *any*) (a, b, c) in the domain lying above (b, c). For example, set (a, b, c) to be some generic point of the fiber above (b, c). Then $(a, b) \in X \times_S \overline{X}$. Note the map

$$\begin{aligned} X \times_S \overline{X} &\dashrightarrow X \\ (w, x) &\mapsto xw \end{aligned}$$

is defined on (a, b), since $\overline{m}$ is a strict group law on $\overline{X}$. Similarly, using Corollary 8.3.3.12, the map

$$\begin{aligned} X \times_S \overline{X} &\dashrightarrow X \\ (w, y) &\mapsto w^{-1}y \end{aligned}$$

is defined at (a, c), which is a generic point in the fiber of over x. Since $\overline{m}$ is defined on X^2, we see that the map

$$\begin{aligned} m' := X \times_S \overline{X} \times_S \overline{X} &\dashrightarrow X \times_S \overline{X} \\ (w, x, y) &\mapsto (w, (xw)(w^{-1}y)) \end{aligned}$$

is defined at (a, b, c), and the associativity of $\overline{m}$ shows that m' coincides with $\mathrm{id}_X \times \overline{m}$. Thus $\overline{m}$ is defined on all $\overline{X}^2$.

By essentially identical arguments, we see that

$$\begin{aligned} \overline{X} \times_S \overline{X} &\dashrightarrow \overline{X} \\ (x, y) &\mapsto x^{-1}y \end{aligned}$$

is also defined everywhere (again, here by $x^{-1}y$ we mean the unique element z satisfying $xz = y$). This, with our prior conclusion, shows that $\overline{m}$ indeed defines the structure of an S-group scheme on $\overline{X}$. This completes our proof. □

Section 8.4

Semistable Reduction of Abelian Varieties

Chapter 9

The Theory of Heights

Section 9.1

Reaching new Heights

The goal of height functions is to measure the "arithmetic complexity" of a point on a variety. This will allow us to generally study properties about rational and integral points of varieties, and in particular abelian varieties. Before we dive into the details of height functions, we note some important properties which our height functions will satisfy.

To begin, we note the basic structure of a height function. Let V be a projective variety over a field k. A height function h will then be a function

$$h : V(k) \to [0, \infty).$$

The first important property which our height functions will satisfy is the finiteness property which states that

$$\#\{P \in V(k) | h(P) \leq B\} < \infty$$

for each $B \in \mathbb{R}$. This property is at the heart of many of the important finiteness theorems throughout arithmetic geometry including the Mordell-Weil theorem (the group of rational points on an abelian variety over $\mathbb{Q}$ is finitely generated), Seigel's theorem (an affine curve of genus $g \geq 1$ over $\mathbb{Q}$ only has finitely many integral points), and the ultimate goal of this book: Falting's theorem (or the Mordell conjecture), which states that any non-singular curve of genus $g \geq 1$ over $\mathbb{Q}$ has only finitely many rational points.

The second feature of our height functions is that they will reflect the underlying geometry of the variety we will work with.

§9.1.1 Review of Absolute Values and Norms

Recall the standard general notion of absolute values

Definition 9.1.1.1. An **absolute value** on a field k is a function with the following properties:

(1) $|x| = 0$ if and only if $x = 0$;

(2) $|xy| = |x| \cdot |y|$;

(3) $|x + y| \leq |x| + |y|$.

Further, we call an absolute value **non-archimedean** if it satisfies

(3') $|x + y| \leq \max\{|x|, |y|\}$.

These absolute values determine a metric and thus a topology on k. We call two absolute values *equivalent* if they define the same topology on k. As notation, we let $\mathscr{P}_k$ be the set of places, $\mathscr{P}_k^\infty$ the set of archimedean (or infinite) places, and $\mathscr{P}_k^0$ the set of non-archimedean (or finite) places. We call classes of absolute values *places*. If K/k is a field extension, we say that a place $w \in \mathscr{P}_K$ *lies over a place* $v \in \mathscr{P}_k$ (denoted $w \mid v$) if the restriction of a representative of w to k is a representative of v.

For the purpose of this book, we will be working over number fields $K/\mathbb{Q}$. We have a nice classification of places on number fields:

EXAMPLE 9.1.1.2. Let K be a number field. Let $\rho_1, \dots, \rho_{r_1} : K \hookrightarrow \mathbb{R}$ be the real embeddings and $\tau_1, \overline{\tau}_1, \dots, \tau_{r_2}, \overline{\tau}_{r_2} : K \hookrightarrow \mathbb{C}$ be the pairs of complex embeddings. Then we have a bijection $\{\rho_1, \dots, \rho_{r_1}, \tau_1, \dots, \tau_{r_2}\} \to \mathscr{P}_K^\infty$ given by $\sigma \mapsto |\sigma(\cdot)|_\infty$ where $|\cdot|_\infty$ denotes the standard absolute value on $\mathbb{R}$ or $\mathbb{C}$.

Now, let $p \in \mathbb{Z}$ a prime and suppose we may write $(p) = \mathfrak{p}_1^e \cdots \mathfrak{p}_g^e$ (the fact that all exponents are the same is a standard fact in algebraic number theory, see ref). Then there is a bijection $\{\mathfrak{p}_1, \dots, \mathfrak{p}_g\} \to \{\text{places lying over } p\}$ given by $\mathfrak{p} \mapsto |\cdot|_\mathfrak{p}$ and $|x|_\mathfrak{p} = p^{-\operatorname{ord}_\mathfrak{p}(x)/e}$.

These represent all of the places in $\mathscr{P}_K$.

To wrap up this review, we will state some important propositions:

PROPOSITION 9.1.1.3 (Degree formula). Let K/k be an extension of number fields and let $v \in \mathscr{P}_k$. Then

$$\sum_{w|v} [K_w : k_v] = [K : k].$$

Proof. [prove later, or cite from somewhere] □

DEFINITION 9.1.1.4. Let K be a global field and $v \in \mathscr{P}_K$. The **local degree** of v is the number $n_v = [K_v : \mathbb{Q}_w]$ where w is the restriction of v to $\mathbb{Q}$. The **normalized absolute value**

$$\|\cdot\|_v := |\cdot|_v^{n_v}.$$

We define the **idelic norm** of $x \in K$ to be

$$\|x\|_{\mathbb{I}_K} := \prod_{v \in \mathscr{P}_K} \|x\|_v$$

REMARK 9.1.1.5. The reason for the term *idelic norm* comes from the construction of the adelic ring of a global field and its corresponding ring of units the ideles. See [Mil20, §V.4] for more.

Proposition 9.1.1.6 (Product formula). Let K be global field and $x \in K^\times$. Then

$$\|x\|_{\mathbb{I}_K} = 1.$$

Proof. Let E/K be finite separable. Then

$$\begin{aligned}\|x\|_{\mathbb{I}_E} &= \prod_{u\in\mathscr{P}_K} \prod_{v\in\mathscr{P}_E,\ v|u} \|x\|_v \\ &= \prod_{u\in\mathscr{P}_K} \prod_{v\in\mathscr{P}_E,\ v|u} \left\|N_{E_v/K_u}(x)\right\|_u \\ &= \prod_{u\in\mathscr{P}_K} \left\|N_{E/K}(x)\right\|_u,\end{aligned}$$

and since the norm carries $E^\times$ into $K^\times$ showing the result for K proves it for E. In particular, it suffices to prove this for $\mathbb{Q}$ or $\mathbb{F}_q(t)$. Furthermore, we see it suffices to verify the result for just the irreducibles of the ring of integers $\mathcal{O}_K$.

- For $K = \mathbb{Q}$, since $\mathcal{O}_\mathbb{Q} = \mathbb{Z}$ we can check it for positive primes p and $\mathbb{Z}^\times = \{\pm 1\}$. By definition $|p|_q = 1$ for all primes $q \neq p$. Thus we only have to multiply the p-adic and infinite place, giving us:

$$\|p\|_{\mathbb{I}_K} = \|p\|_\infty \cdot \|p\|_p = p \cdot \frac{1}{p} = 1.$$

- For $K = \mathbb{F}_q(t)$, we only have to consider irreducible polynomials $p(t) \in \mathbb{F}_q[t]$. Again this has trivial absolute value at all but two places. This gives us:

$$\|p(t)\|_{\mathbb{I}_K} = \|p(t)\|_\infty \cdot \|p(t)\|_{p(t)} = q^{\deg p} \cdot \left(\#\left\{\mathbb{F}_q[t]/\left\langle p(t)\right\rangle\right\}\right)^{-1} = 1.$$

This completes our proof. □

§9.1.2 Heights on Projective Space

Now that we have a way to measure the arithmetic complexity of elements of number fields, we can easily extend this to points in projective space.

Definition 9.1.2.1. Let K be a number field and $P \in \mathbb{P}^n(K)$ with $P = [x_0 : \cdots : x_n]$. The **multiplicative height** of P (with respect to K) is defined as

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

The **logarithmic height** of P is defined as

$$h_K(P) = \log H_K(P) = \sum_{v\in\mathscr{P}_K} -n_v \min\{\log(|x_0|_v), \ldots, \log(|x_n|_v)\}.$$

Lemma 9.1.2.2. (1) The multiplicative (and thus logarithmic) height of a point P is independent of the choice of homogeneous coordinates.

(2) $H_K(P) \geq 1$ for all $P \in \mathbb{P}^n(K)$

(3) Let K'/K be a finite extension. Then

$$H_{K'}(P) = H_K(P)^{[K':K]}$$

Proof. (1) Let $P = [x_0 : \cdots : x_n]$. Then any other representative of P has the form $[cx_0 : \cdots : cx_n]$ for some $c \in K^\times$. So,

$$\begin{aligned}\prod_{v\in\mathscr{P}_K} \max\{\|cx_0\|_v, \dots \|cx_n\|_v\} &= \left(\prod_{v\in\mathscr{P}_K} \|c\|_v\right)\left(\prod_{v\in\mathscr{P}_K} \max\{\|x_0\|_v, \dots \|x_n\|_v\}\right) \\ &= \prod_{v\in\mathscr{P}_K} \max\{\|x_0\|_v, \dots \|x_n\|_v\}\end{aligned}$$

where the final equality follows from the product formula (9.1.1.6). This settles (1).

(2) Again, let $P = [x_0 : \cdots : x_n]$. By part (1), we may assume by choice of coordinate that $x_0 = 1$. Then $\max\{\|x_0\|_v, \dots, \|x_n\|_v\} \geq \|x_0\|_v = 1$ for all v, so the assertion is clear from the definition of H_K.

(3) Fix $v' \in M_{K'}$, let $v \in \mathscr{P}_K$ and $v_0 \in M_{\mathbb{Q}}$ be the places lying below v' in their respective fields. Then by definition **??**, we have that

$$\|\cdot\|_{v'} = \|\cdot\|_v^{[K'_{v'}:K_v]} = |\cdot|_v^{[K'_{v'}:\mathbb{Q}_{v_0}]}.$$

Now, restricting our attention to places lying over a particular $w \in \mathscr{P}_K$, we see that

$$\begin{aligned}\prod_{w'|w} \max\{\|x_0\|_{w'}, \dots, \|x_n\|_{w'}\} &= \prod_{w'|w} \max\{\|x_0\|_w, \dots, \|x_n\|_w\}^{[K'_{w'}:K_w]} \\ &= \max\{\|x_0\|_w, \dots, \|x_n\|_w\}^{\sum_{w'|w}[K'_{w'}:K_w]} \\ &= \max\{\|x_0\|_w, \dots, \|x_n\|_w\}^{[K':K]}\end{aligned}$$

where the last equality follows by the Degree formula (9.1.1.3). Now, taking the product over all $w \in \mathscr{P}_K$, the claim follows.

□

The above lemma allows us to give the following

DEFINITION 9.1.2.3. We define the **absolute (multiplicative) height** on $\mathbb{P}^n$ is the function

$$H : \mathbb{P}^n(\overline{\mathbb{Q}}) \to [1, \infty)$$

such that for each $P \in \mathbb{P}^n(K)$, $H(P) = H_K(P)^{\frac{1}{[K:\mathbb{Q}]}}$. We similarly define the **absolute (logarithmic) height** on $\mathbb{P}^n$ by

$$h(P) = \log H(P) = \frac{1}{[K:\mathbb{Q}]} h_K(P).$$

This is independent of the field K by 9.1.2.2 (3). For notational convenience, for any element $x \in K$, we will define $H_K(x)$ to be $H_K([x:1])$ in $\mathbb{P}^2$ and similarly for $h_K(x), H(x)$, and $h(x)$.

We will show that height is invariant under the action of the absolute Galois group of $\mathbb{Q}$. To begin, we will state some simple results about how the Galois group acts on places. Let K be a finite subextension of $\overline{\mathbb{Q}}/\mathbb{Q}$, i.e., a number field. Any $\sigma \in G_{\overline{\mathbb{Q}}}$ defines an isomorphism $k \xrightarrow{\sim} \sigma(K)$ and defines an absolute value

$\sigma(v)$ on $\sigma(K)$ via $|\sigma(x)|_{\sigma(v)} = |x|_v$ for $x \in K$. It thus identifies

$$\begin{aligned}\sigma : \mathscr{P}_K &\xrightarrow{\sim} \mathscr{P}_{\sigma(K)} \\ v &\mapsto \sigma(v),\end{aligned}$$

via which we have an isomorphism $K_v \xrightarrow{\sim} \sigma(K)_{\sigma(v)}$. In particular, as σ fixes the restriction w of v to $\mathbb{Q}$, this shows that $n_v := [K_v : \mathbb{Q}_w] = [\sigma(K)_v : \mathbb{Q}_w] =: n_{\sigma(v)}$.

PROPOSITION 9.1.2.4. The action of the absolute Galois group $G_{\mathbb{Q}} := \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ on $\mathbb{P}^n(\overline{\mathbb{Q}})$ leaves height invariant. I.e., for $P \in \mathbb{P}^n(\overline{\mathbb{Q}})$, $\sigma \in G_{\mathbb{Q}}$, $H(P) = H(\sigma(P))$.

Proof. Let $K/\mathbb{Q}$ be a number field with $P \in \mathbb{P}^n(K)$ and let $P = [x_0 : \cdots : x_n]$ with $x_i \in K$. By the above,

$$\begin{aligned}H_{\sigma(K)}(\sigma(P)) &= \prod_{v \in M_{\sigma(K)}} \max\{\|\sigma(x_0)\|_v, \ldots, \|\sigma(x_n)\|_v\} \\ &= \prod_{v \in M_{\sigma(K)}} \max\{|\sigma(x_0)|_v, \ldots, |\sigma(x_n)|_v\}^{n_v} \\ &= \prod_{w \in M_{\sigma(K)}} \max\{|\sigma(x_0)|_{\sigma(w)}, \ldots, |\sigma(x_n)|_{\sigma(w)}\}^{n_{\sigma(w)}} \\ &= \prod_{w \in M_{\sigma(K)}} \max\{|x_0|_w, \ldots, |x_n|_w\}^{n_w} \\ &= H_K(P)\end{aligned}$$

Then letting $n = [K : \mathbb{Q}] = [\sigma(K) : \mathbb{Q}]$ and taking nth roots gives us our result. □

DEFINITION 9.1.2.5. Let $P = [x_0 : \cdots : x_n] \in \mathbb{P}(\overline{\mathbb{Q}})$ be a point. The **field of definition** of P is

$$\mathbb{Q}(P) = \mathbb{Q}\left(\frac{x_0}{x_j}, \ldots, \frac{x_n}{x_j}\right)$$

for j such that $x_j \neq 0$.

Add some exposition here later since result is important.

EXAMPLE 9.1.2.6. Let $P \in \mathbb{P}^n(\mathbb{Q})$. Then we may write $P = [x_0 : \cdots : x_n]$ such that $x_i \in \mathbb{Z}$ for all i and $\gcd(x_0, \ldots, x_n) = 1$. Then

$$H(P) = \max\{|x_0|, \ldots, |x_n|\}. \tag{9.1.2-I}$$

EXERCISE 9.1.2.A (Easy). Check equation 9.1.2-I.

THEOREM 9.1.2.7 (Finiteness Theorem). For numbers $B, D \geq 0$, the set

$$\{P \in \mathbb{P}^n(\overline{\mathbb{Q}}) | H(P) \leq B \text{ and } [\mathbb{Q}(P) : \mathbb{Q}] \leq D\}$$

is finite.

Proof. This is obvious if we restrict to points in $\mathbb{P}^n(\mathbb{Q})$ by example 9.1.2.6 above. We will reduce the general case to this.

Let $P = [x_0 : \cdots : x_n]$ and choose homogeneous coordinates such that $x_i = 1$ for some i. Since $H([x_0, \ldots, x_n]) \geq H([1, x_j])$ and $\mathbb{Q}(x_j) \subset \mathbb{Q}(P)$ for all j, we are reduced to the case $P = [1 : x] \in \mathbb{P}^1(\overline{\mathbb{Q}})$. It is clearly equivalent to prove the sets

$$\{P \in \mathbb{P}^n(\overline{\mathbb{Q}}) | H(P) \leq B \text{ and } [\mathbb{Q}(P) : \mathbb{Q}] = d\}$$

are finite for all $d \leq D$, we may assume $[\mathbb{Q}(x) : \mathbb{Q}] = d$. Let $x^{(1)}, \ldots, x^{(d)}$ be the conjugates of x and $1 = s_0, \ldots, s_d$ be the elementary symmetric polynomials in the $x^{(i)}$'s so that the characteristic polynomial of x,

$$F_x(t) = \prod_{i=1}^{d} \left(t - x^{(i)}\right) = \sum_{j=0}^{d} (-1)^j s_j T^{d-j}.$$

Let v be a place on $\mathbb{Q}(x)$. Using the triangle inequality, we have

$$\begin{aligned} |s_j|_v &= \Big| \sum_{1 \leq i_1, < \cdots < i_m \leq d} x^{(i_1)} \cdots x^{(i_j)} \Big|_v \\ &\leq c \max_{1 \leq i_1, < \cdots < i_j \leq d} |x^{(i_1)} \cdots x^{(i_j)}|_v \\ &\leq c \max_{i \in [d]} |x^{(i)}|_v^j \\ &\leq c \prod_{i \in [d]} \max\{|x^{(i)}|_v, 1\}^j \end{aligned}$$

where we may take $c = 2^d$ if v is archimedean (as there are $\binom{d}{j} \leq 2^d$ terms in the sum) and $c = 1$ otherwise. Now, taking the product over all places and taking the dth roots, since there are only finitely many non-archimedean places, there is a suitably large C (dependent only on d) such that

$$H(s_j) \leq C \prod_{i=1}^{d} H(x^{(i)})^j$$

but by Proposition 9.1.2.4, $H(x^{(i)}) = H(x)$ for all i. Thus, we have $H(s_j) \leq CH(x^{(i)})^{dj} \leq CH(x^{(i)})^{d^2}$ for all j.

Finally, suppose x is in the set

$$\{x \in \overline{\mathbb{Q}} \mid H(x) \leq B, [\mathbb{Q}(x) : \mathbb{Q}] = d\}.$$

We have just shown that x is the root of a polynomial $F_x(t) \in \mathbb{Q}[t]$ with coefficients $s_0, \ldots, s_d$ such that $H(s_j) \leq CB^{d^2}$ for all j. However, there are only finitely many such s_j by example 9.1.2.6. Thus there are only finitely many possible polynomials $F_x(t)$ and hence only finitely many possible x. □

We have the following immediate corollary:

> COROLLARY 9.1.2.8 (Kronecker). Let $P = [x_0 : \cdots : x_n] \in \mathbb{P}^n(K)$ for K a number field. Suppose $x_i \neq 0$. Then $H(P) = 1$ if and only if $\frac{x_j}{x_i}$ is either 0 or a root of unity for all j.

Proof. The backward direction is clear as $|x_j|_v = 1$ if x_j is a root of unity.

Conversely, suppose that $H(P) = 1$ and let $P^r := [x_0^r : \cdots : x_n^r]$. It is clear that $H(P^r) = H(P)^r$, so in particular, $H(P^r) = 1$ for all $r \geq 0$. But by theorem 9.1.2.7, the collectoin of points in $\mathbb{P}^n(K)$ with height 1 is finite, so we must have some s, t such that $P^s = P^t$, so $x_i^s = x_i^t$ for all i, which shows that each x_i is either 0 or a $(s-t)$th root of unity. □

There's some extra stuff in Hindry Silverman, but idk how useful it is

§9.1.3 Heights on Projective Varieties

Our goal is to now define heights over a general projective variety V over $\overline{\mathbb{Q}}$. To do so, we start with a morphism to projective space, then use the machinery we built in the previous section.

DEFINITION 9.1.3.1. Let $\phi : V \to \mathbb{P}^n$ be a morphism of projective varieties. We define the (absolute logarithmic) **height on V relative to ϕ** as the function

$$\begin{aligned} h_\phi : V(\overline{\mathbb{Q}}) &\to [0, \infty) \\ P &\mapsto h(\phi(P)) \end{aligned}$$

where $h : \mathbb{P}(\overline{\mathbb{Q}}) \to [0, \infty)$ is the height function from definition 9.1.2.3.

It is a classical result ([Har77, Ch. II, Thm 7.1]) that any such morphism $\phi : V \to \mathbb{P}^n$ is defined by a line bundle, namely $\phi^*(\mathcal{O}(1))$. There can be many morphisms which give rise to the same line bundle, and as such, we would like to know how the height function varies with respect to these morphisms. It turns out that it behaves nicely, made precise in the following:

THEOREM 9.1.3.2. Let $\mathscr{L}$ be a basepoint free line bundle (i.e., $\mathcal{O}_Z \otimes \Gamma(\mathscr{L}^{\otimes 2}) \to \mathscr{L}^{\otimes 2}$ is surjective), and let $\phi : V \to \mathbb{P}^n$ and $\psi : V \to \mathbb{P}^m$ be two maps associated to $\mathscr{L}$. Then,

$$h_\phi = h_\psi + O(1),$$

i.e. $|h_\phi(P) - h_\psi(P)|$ is bounded as P ranges over $V(\overline{\mathbb{Q}})$.

REMARK 9.1.3.3. In terms of Weil divisors, if $\mathscr{L} \simeq \mathcal{O}_V(E)$ for some divisor E, and if $H \subset \mathbb{P}^n$ is a hyperplane not containing $\phi(V)$, then $\mathscr{L}$ is associated to ϕ if and only if $\phi^*(H)$ is linearly equivalent to E. We will use this viewpoint in the proof.

Proof. Let E be an effective divisor in the linear equivalence class of $\mathscr{L}$ (so that $\mathcal{O}_V(E) \simeq \mathscr{L}$). By [Har77, Ch. II, Rmk. 7.8.1], on the complement of E, we may write $\phi = [f_0 : \cdots : f_n]$ and $\psi = [g_0 : \cdots : g_m]$ for $f_i, g_j \in \Gamma(V, \mathscr{L})$, thought of as rational functions such that $(f_i) = D_i - E$ and $(g_j) = D_j' - E$ for effective divisors D_i, D_j'. Furthermore, since $\mathscr{L}$ has no basepoints, the D_i's have no points in common.

Let K be a common number field of definition for V, $f_0, \ldots, f_n$, and $g_0, \ldots, g_m$. Pick any j and consider the ideal $I := (\frac{f_0}{g_j}, \ldots, \frac{f_n}{g_j})$ in the ring $R := K[\frac{f_0}{g_j}, \ldots, \frac{f_n}{g_j}]$. Since $(\frac{f_i}{g_j}) = D_i - D_j'$ and the D_i's have no points in common, we claim that

$$I \text{ is the unit ideal.} \tag{9.1.3-I}$$

To prove claim 9.1.3-I, we suppose it's false for contradiction. Then, there exists a maximal ideal $\mathfrak{m}$ such that $I \subset \mathfrak{m}$. Under the identification of Spec R as an open subscheme of V containing $V - D'_j$, $\mathfrak{m}$ corresponds to a point $Q \notin D'_j$ such that $\left(\frac{f_i}{g_j}\right)(Q) = 0$ for all i. But then Q would lie in the support of D_i for all i, a contradiction.

Now 9.1.3-I tells us that there is a polynomial $P_j \in K[T_0, \ldots, T_n]$ with zero constant term such that

$$P_j(\frac{f_0}{g_j}, \ldots, \frac{f_n}{g_j}) = 1.$$

Taking v-adic absolute values and using the triangle inequality, we find a constant $C_1 > 0$, dependent on v, ϕ, ψ, P_j such that for all $Q \in V - D'_j$,

$$\max\{|\frac{f_0}{g_j}|_v, \ldots, |\frac{f_n}{g_j}|_v\} \geq C_1$$

and further, we can choose $C_1 = 1$ for all but finitely many v, independent of Q.

Multiplying through by $|g_j|_v$ (noting that the inequality still holds when $g_j(Q) = 0$) and taking the maximum over all j, we have

$$\max\{|f_0(Q)|_v, \ldots, |f_n(Q)|_v\} \geq C_2 \max\{|g_0(Q)|_v, \ldots, |g_m(Q)|_v\}$$

for a constant $C_2 > 0$, dependent on v, ϕ, ψ, where Q ranges over the complement of E and we may choose $C_2 = 1$ for all but finitely many v. Raising each of these to the $[K_v : \mathbb{Q}_w]$th power (where $w = v|_{\mathbb{Q}}$), multiplying over all $v \in \mathscr{P}_K$, and taking the $[K : \mathbb{Q}]$th roots gives us

$$H(\phi(Q)) \geq C_3 H(\psi(Q))$$

where $C_3 > 0$ is the product over all $C_2(v)$ (and is well-defined since we chose all but finitely many $C_2(v)$ to be 1), again with Q running over the complement of E in V. Now, since $\mathscr{L}$ has no basepoints, we may choose divisors $E_1, \ldots, E_r$ in the linear system for $\mathscr{L}$ with trivial intersection to get this inequality on all $P \in V$. The other bound follows by symmetry. Finally, taking logarithms gives us the form in the statement of the theorem. □

With this theorem, we can now introduce Weil's height machine, which relates geometric information about line bundles (or divisors) to arithmetic information, their respective height functions. Before we begin, we will introduce a notion which will make our notation easier.

DEFINITION 9.1.3.4. We define the **group of functions mod** $O(1)$ on V,

$$\mathscr{H}(V) := \{\text{functions } V(\overline{\mathbb{Q}}) \to \mathbb{R}\}/\{\text{bounded functions}\}.$$

Now we may set up the following:

THEOREM 9.1.3.5 (Weil's Height Machine). Let K be a number field and V/K a projective variety. Then there exists a unique homomorphism

$$HM : \text{Pic}(V) \to \mathscr{H}(V)$$
$$\mathscr{L} \mapsto h_{\mathscr{L}}$$

such that for any $\mathscr{L}$ has no base points and $\phi : V \to \mathbb{P}^n$ is an associated morphism, then

$$h_{\mathscr{L}} = h_\phi + O(1).$$

Additionally, it has the following properties:

(1) (Functoriality) If $f : V \to W$ is a morphism of projective varieties, and $\mathscr{L}$ is a line bundle on W, then

$$h_{f^*\mathscr{L}} = h_{\mathscr{L}} \circ f + O(1).$$

(2) (Positivity) Let $\mathscr{L}$ be a line bundle on V whose base locus B is not all of V (i.e., there is a rational map $V \to \mathbb{P}^n$ associated to $\mathscr{L}$ which is a morphism on the complement of B.) Then

$$h_{\mathscr{L}} \geq O(1)$$

on the complement of B.

(3) (Algebraic Equivalence) Let $\mathscr{L}$ and $\mathscr{M}$ be algebraically equivalent divisors (see Def 1.2.5.2), and assume that $\mathscr{L}$ is ample. Then for all $\epsilon > 0$,

$$(1-\epsilon)h_{\mathscr{L}} - O(1) \leq h_{\mathscr{M}} \leq (1+\epsilon)h_{\mathscr{L}} + O(1),$$

(here, the $O(1)$ constants depend on ϵ). Note in particular that

$$\lim_{\substack{P \in V(\overline{\mathbb{Q}}) \\ h_{\mathscr{L}}(P) \to \infty}} \frac{h_{\mathscr{M}}(P)}{h_{\mathscr{L}}(P)} = 1.$$

(4) (Finiteness) If $\mathscr{L}$ is ample, then for all constants $B, D \geq 0$, the set

$$\{P \in \mathbb{P}^n(\overline{\mathbb{Q}}) \mid h_{\mathscr{L}}(P) \leq B \text{ and } [\mathbb{Q}(P) : \mathbb{Q}] \leq D\}$$

is finite.

Sanity Check. Let $\mathscr{L} = \mathscr{O}_{\mathbb{P}^n}(H)$ for some hyperplane $H \subset \mathbb{P}^n$, and let h be the absolute height function on $\mathbb{P}^n$ (see Def 9.1.2.3). Verify that

$$h_{\mathscr{L}} = h + O(1).$$

Remark 9.1.3.6. Because we do not assume that V is smooth, when we mention a divisor, we are referring to Cartier divisors.

We will break the proof down into multiple Propositions for ease of writing and reading. We begin by defining the map $h_{\mathscr{L}}$ for general $\mathscr{L} \in \mathrm{Pic}(V)$. If $\mathscr{L}$ has no basepoints, we define it as the obvious function based on our previous work – if $\phi : V \to \mathbb{P}^n$ is a map associated to $\mathscr{L}$, then we define $h_{\mathscr{L}} = h_\phi \in \mathscr{H}(V)$. For general $\mathscr{L}$, we may write $\mathscr{L} = \mathscr{L}_1 \otimes \mathscr{L}_2^{-1}$ where $\mathscr{L}_1, \mathscr{L}_2$ are basepoint free (see give reference), so we define $h_{\mathscr{L}} = h_{\mathscr{L}_1} - h_{\mathscr{L}_2} \in \mathscr{H}(V)$.

PROPOSITION 9.1.3.7 (Well-Definedness). The map

$$HM : \mathrm{Pic}(V) \to \mathscr{H}(V)$$
$$\mathscr{L} \mapsto h_{\mathscr{L}}$$

is well-defined

This is the content of Theorem 9.1.3.2.

PROPOSITION 9.1.3.8 (Additivity). HM is a homomorphism, i.e.

$$h_{\mathscr{L}\otimes\mathscr{M}} = h_{\mathscr{L}} + h_{\mathscr{M}} + O(1)$$

for all $\mathscr{L}, \mathscr{M} \in \mathrm{Pic}(V)$.

It is sufficient to check this on $\mathscr{L}, \mathscr{M}$ with no basepoints by the definition of $h_{\mathscr{L}}$. Let $\phi = [f_0 : \cdots : f_n]$ and $\psi = [g_0 : \cdots : g_m]$ be maps to projective space associated to $\mathscr{L}$ and $\mathscr{M}$ respectively. Then the composition map

$$V \xrightarrow{\delta} V \times V \xrightarrow{\phi\times\psi} \mathbb{P}^n \times \mathbb{P}^m \xrightarrow{s} \mathbb{P}^{nm+n+m}$$

(where $s : \mathbb{P}^n \times \mathbb{P}^m \to \mathbb{P}^{nm+n+m}$ is the Segre embedding) is associated to $\mathscr{L} \otimes \mathscr{M}$. Explicitly this map is given by

$$[\ldots, f_i g_j, \ldots]_{0\le i\le n, 0\le j\le m},$$

which using this description, we immediately get

$$\max\{\ldots, |f_i g_j|, \ldots\} = \max\{f_0, \ldots, f_n\} \max\{g_0, \ldots, g_m\},$$

which proves the proposition.

PROPOSITION 9.1.3.9 (Uniqueness). The homomorphism HM is determined by the property that for $\mathscr{L}$ with no base points and $\phi : V \to \mathbb{P}^n$ an associated morphism, $h_{\mathscr{L}} = h_\phi + O(1)$ (as defined in 9.1.2.1).

EXERCISE 9.1.3.A (Easy). Check this!

PROPOSITION 9.1.3.10 (Functoriality). If $f : V \to W$ is a morphism of projective varieties, and $\mathscr{L}$ is a line bundle on W, then

$$h_{f^*\mathscr{L}} = h_{\mathscr{L}} \circ f + O(1).$$

Again, by additivity, we may check on $\mathscr{L}$ with no basepoints. Here, the result follows easily from the fact that if ϕ is the morphism associated to $\mathscr{L}$, then $\phi \circ f$ is the morphism associated to $f^*\mathscr{L}$.

PROPOSITION 9.1.3.11 (Positivity). Let $\mathscr{L}$ be a line bundle on V whose base locus B is not all of V (i.e., there is a rational map $V \to \mathbb{P}^n$ associated to $\mathscr{L}$ which is a morphism on the complement of B.) Then

$$h_{\mathscr{L}} \ge O(1)$$

on the complement of B.

As before, we may write $\mathscr{L}$ as a difference of basepoint-free line bundles $\mathscr{L} = \mathscr{M}_1 \otimes \mathscr{M}_2^{-1}$ and suppose. Because $B \neq V$, there are effective divisors, call them D, E_1, E_2 in the linear systems for $\mathscr{L}, \mathscr{M}_2, \mathscr{M}_2$ respectively. Consider a basis $\{f_0, \dots, f_n\}$ for the linear system $|E_2|$. By the positivity of D,

$$E_1 + (f_i) = D + E_2 + (f_i) \geq 0,$$

so each $f_i \in |E_1|$ and we may extend $\{f_0, \dots, f_n\}$ to a basis $\{f_0, \dots, f_n, g_1, \dots, g_m\}$ for $|E_1|$. Then, we can associate to $\mathscr{M}_2$ the morphism

$$\phi = (f_0 : \cdots : f_n) : V \to \mathbb{P}^n,$$

and to $\mathscr{M}_1$ the morphism

$$\psi = (f_0 : \cdots : f_n : g_1 : \cdots : g_m) : V \to \mathbb{P}^{n+m}.$$

The functions $f_0, \dots, f_n$ are regular outside of $\mathrm{Supp}(E_1)$ and so for any point $P \in V - \mathrm{Supp}(E_1)$, we have

$$\begin{aligned} h_{\mathscr{L}} &= h_{\mathscr{M}_1} - h_{\mathscr{M}_2} + O(1) \\ h(f_0(P) : \cdots : f_n(P) : g_1(P) : \cdots : g_m(P) - h(f_0(P) : \dots f_n(P)) + O(1) & \\ &\geq O(1) \end{aligned}$$

where the inequality follows from the trivial fact that for all $v \in \mathscr{P}_K$,

$$\max\{\|f_0(P)\|_v, \dots, \|f_n(P)\|_v, \|g_1(P)\|_v, \dots, \|g_m(P)\|_v\} \geq \max\{\|f_0(P)\|_v, \dots, \|f_n(P)\|_v\}.$$

Now, to extend this result to all points in V, choose very ample line bundles $\mathcal{O}_V(D_1), \dots, \mathcal{O}_V(D_r)$ such that $\cap_i \mathrm{Supp}(D_i) = \emptyset$ and $\mathscr{L} \otimes \mathcal{O}_V(D_i)$ are very ample for all i (for the existence of such D_i, see Rmk 9.1.3.12). Now, applying what we did above to the decompositions $\mathscr{L} = (\mathscr{L} \otimes \mathcal{O}_V(D_i)) \otimes \mathcal{O}_V(D_i)^{-1}$, we get our result for all points $P \in V - \mathrm{Supp}(D)$. Finally, varying over all divisors in $|D|$ gives us our result on all $P \in V - B$.

Remark 9.1.3.12. In the proof above, we mentioned that for effective divisor D, we may find D_i's with $\mathcal{O}_V(D_i)$ very ample such that$\cap_i \mathrm{Supp}(D_i) = \emptyset$ and $\mathcal{O}(D) \otimes \mathcal{O}_V(D_i)$ very ample for all i. One way of constructing these is by considering very ample $\mathcal{O}_V(E)$ for which there is some N such that $\mathcal{O}_V(D + NE)$ is very ample (N's existence is clear by using Serre's criterion), then if $\phi : V \to \mathbb{P}^n$ is a morphism associated to $\mathcal{O}_V(D + NE)$, we then take the pullback $\phi^{-1}(\mathcal{O}(H_i))$ where the H_i's are the coordinate hyperplanes.

Proposition 9.1.3.13 (Algebraic Equivalence). Let $\mathscr{L}$ and $\mathscr{M}$ be algebraically equivalent divisors (see Def 1.2.5.2), and assume that $\mathscr{L}$ is ample. Then for all $\epsilon > 0$,

$$(1 - \epsilon)h_{\mathscr{L}} - O(1) \leq h_{\mathscr{M}} \leq (1 + \epsilon)h_{\mathscr{L}} + O(1),$$

(here, the $O(1)$ constants depend on ϵ).

As ampleness is preserved under algebraic equivalence (see Rmk 1.2.5.5), $\mathscr{M}$ is also ample. By Lemma 9.1.3.16 (whose proof is postponed until later), there is some integer e such that for all $n > 0$,

$$(\mathscr{M} \otimes \mathscr{L}^{-1})^{\otimes n} \otimes \mathscr{M}^{\otimes e} \simeq \mathscr{L}_n$$

where $\mathscr{L}_n$ is effective and very ample, and thus base-point free. Note that if $\mathscr{L}, \mathscr{M}$ are defined over a number field K, then so are the $\mathscr{L}_n$ for all n. We have that $\mathscr{M}^{\otimes(n+e)} \simeq \mathscr{L}^{\otimes n} \otimes \mathscr{L}_n$ for all n, and taking heights gives

$$(n+e)h_{\mathscr{M}} = nh_{\mathscr{L}} + h_{\mathscr{L}_n}.$$

Let $0 < \epsilon < 1$ (the $\epsilon > 1$ case clearly follows from this) and choose n large enough that $\frac{n}{n+e} < 1 - \epsilon$. Dividing by $n+e$ and noting that $h_{\mathscr{L}_n} \geq O(1)$ (by positivity and basepoint-freeness) gives us that

$$h_{\mathscr{M}}(P) \geq (1-\epsilon)h_{\mathscr{L}}(P) + O(1)$$

for all $P \in V$. Symmetry gives the other inequality, so we are done.

COROLLARY 9.1.3.14. With $\mathscr{L}, \mathscr{M}$ as in 9.1.3.13,

$$\lim_{\substack{P \in V(\overline{\mathbb{Q}}) \\ h_{\mathscr{L}}(P) \to \infty}} \frac{h_{\mathscr{M}}(P)}{h_{\mathscr{L}}(P)} = 1.$$

In addition, if $\mathscr{N}$ is algebraically equivalent to zero, then

$$\lim_{\substack{P \in V(\overline{\mathbb{Q}}) \\ h_{\mathscr{L}}(P) \to \infty}} \frac{h_{\mathscr{N}}(P)}{h_{\mathscr{L}}(P)} = 0.$$

The first claim is an immediate consequence of Prop 9.1.3.13. For the second, by the definition of algebraic equivalence, there is some connnected K-scheme S and a line bundle $\overline{\mathscr{N}}$ on $S \times V$ such that at some points $s, s' \in S$, $\overline{\mathscr{N}}_s \simeq \mathcal{O}_V$ and $\overline{\mathscr{N}}_{s'} \simeq \mathscr{N}$. Consider the line bundle $\mathrm{pr}_2^*(\mathscr{L}) \otimes \overline{\mathscr{N}}$. It is clear that $(\mathrm{pr}_2^*(\mathscr{L}) \otimes \overline{\mathscr{N}})_s \simeq \mathscr{L}$ and $(\mathrm{pr}_2^*(\mathscr{L}) \otimes \overline{\mathscr{N}})_{s'} \simeq \mathscr{L} \otimes \mathscr{N}$, so in particular $\mathscr{L} \otimes \mathscr{N}$ is algebraically equivalent to $\mathscr{L}$. Thus,

$$\lim_{\substack{P \in V(\overline{\mathbb{Q}}) \\ h_{\mathscr{L}}(P) \to \infty}} \frac{h_{\mathscr{N}}(P)}{h_{\mathscr{L}}(P)} = \lim_{\substack{P \in V(\overline{\mathbb{Q}}) \\ h_{\mathscr{L}}(P) \to \infty}} \frac{h_{\mathscr{L}\otimes\mathscr{N}}(P) - h_{\mathscr{L}}(P)}{h_{\mathscr{L}}(P)} = 1 - 1 = 0$$

by the previous result.

PROPOSITION 9.1.3.15 (Finiteness). If $\mathscr{L}$ is ample, then for all constants $B, D \geq 0$, the set

$$\{P \in \mathbb{P}^n(\overline{\mathbb{Q}}) \mid h_{\mathscr{L}}(P) \leq B \text{ and } [\mathbb{Q}(P) : \mathbb{Q}] \leq D\}$$

is finite.

Firstly note that since different representatives of $h_{\mathscr{L}}$ differ up to a bounded function independent of P, this statement is independent of the choice of representative for $h_{\mathscr{L}}$. Since $h_{\mathscr{L}^{\otimes m}} = mh_{\mathscr{L}} + O(1)$, it suffices to prove this for some tensor power of $\mathscr{L}$. In particular, since $\mathscr{L}$ is ample, then we may assume $\mathscr{L}$ is very ample. Thus, $\mathscr{L}$ defines an embedding $\phi : V \hookrightarrow \mathbb{P}^n$ for some n. But then

$$h_{\mathscr{L}} \circ \phi = h + O(1)$$

by functoriality and the fact that $\mathscr{L} \simeq \phi^*(\mathcal{O}(1))$. Thus, this statement reduces to Theorem 9.1.2.7.

We left a key lemma out of the proof of algebraic equivalence, so record it here.

LEMMA 9.1.3.16. Let V be a non-singular proper variety. Let $\mathscr{L}$ be an ample line bundle on V. Then there exists an integer $e > 0$ such that for any $\mathscr{M}$ algebraically equivalent to 0, $\mathscr{M} \otimes \mathscr{L}^{\otimes e}$ is very ample.

Proof. See Lang's Diophantine Geometry Lemma 4.3.2 □

Now that we have constructed Weil's height machine, we will give some consequences to illustrate its usefulness.

PROPOSITION 9.1.3.17. Let C/K be a smooth projective curve over a number field. Then if $\mathscr{L}, \mathscr{M}$ are line bundles on C with $\deg \mathscr{L} \geq 1$. Then

$$\lim_{\substack{P \in V(\overline{\mathbb{Q}}) \\ h_{\mathscr{L}}(P) \to \infty}} \frac{h_{\mathscr{M}}(P)}{h_{\mathscr{L}}(P)} = \frac{\deg \mathscr{L}}{\deg \mathscr{M}}.$$

and if $f, g \in k(C)$ are rational functions on C, then

$$\lim_{\substack{P \in V(\overline{\mathbb{Q}}) \\ h(f(P)) \to \infty}} \frac{h(f(P))}{h(g(P))} = \frac{\deg f}{\deg g}.$$

Proof. Finish Later! □

§9.1.4 Canonical Heights

Now that we have constructed Weil's height machine (9.1.3.5), we have a way to associate geometric information, line bundles (or divisors if you prefer), to equivalence classes of functions mod $O(1)$ satisfying many nice properties. In some cases, it's possible to pin down a function within this equivalence class which satisfies particularly nice properties. This theory was developed by Néron and Tate. Finish Later!

DEFINITION 9.1.4.1. Let K be a number field, V a projective variety over K, $\mathscr{L}$ a line bundle, and $\phi : V \to V$ a morphism. Suppose further that

$$\phi^* \mathscr{L} \simeq \mathscr{L}^{\otimes \lambda}$$

for some $\lambda \geq 2$. Then we define the **canonical height** on V relative to ϕ and $\mathscr{L}$,

$$\widehat{h}_{V,\phi,\mathscr{L}} : V(\overline{\mathbb{Q}}) \to \mathbb{R},$$

(or $\widehat{h}_{\phi,\mathscr{L}}$ when V is clear) via

$$\widehat{h}_{V,\phi,\mathscr{L}}(P) = \lim_{n \to \infty} \frac{1}{\lambda^n} h_{V,\mathscr{L}}(\phi^n(P)) \tag{9.1.4-I}$$

where ϕ^n is the n-fold iteration of ϕ.

REMARK 9.1.4.2. If we consider the induced linear map $\phi^* : \mathrm{Pic}(V) \otimes_{\mathbb{Z}} \mathbb{R} \to \mathrm{Pic}(V) \otimes_{\mathbb{Z}} \mathbb{R}$, we may think of the condition that $\phi^* \mathscr{L} \simeq \mathscr{L}^{\otimes \lambda}$ as $\mathscr{L}$ being an eigenvector with eigenvalue $\lambda > 1$ for ϕ^*.

We will first check that the limit 9.1.4-I converges for all P. Applying Weil's height machine (9.1.3.5) to the relation $\phi^*\mathscr{L} \simeq \mathscr{L}^{\otimes\lambda}$, we get a constant C such that

$$|h_{\mathscr{L}}(\phi(P)) - \lambda h_{\mathscr{L}}(P)| \leq C \tag{9.1.4-II}$$

for all $P \in V(\overline{\mathbb{Q}})$. (Note that C depends on $V, \mathscr{L}, \phi$, and the choice of $h_{\mathscr{L}}$). Fixing some $Q \in V(\overline{k})$, we will verify that limit 9.1.4-I converges by showing the sequence is Cauchy. Let $n \geq m$ and consider the telescoping sum expansion

$$\begin{aligned}
|\lambda^{-n} h_{\mathscr{L}}(\phi^n(Q)) - \lambda^{-m} h_{\mathscr{L}}(\phi^m(Q))| &= \left| \sum_{i=m+1}^{n} \lambda^{-i} \left(h_{\mathscr{L}}(\phi^i(Q)) - h_{\mathscr{L}}(\phi^{i-1}(Q)) \right) \right| \\
&\leq \sum_{i=m+1}^{n} \lambda^{-i} \left| \left(h_{\mathscr{L}}(\phi^i(Q)) - h_{\mathscr{L}}(\phi^{i-1}(Q)) \right) \right| \\
&\leq \sum_{i=m+1}^{n} \lambda^{-i} C \\
&= \left(\frac{\lambda^{-m} - \lambda^{-n}}{\lambda - 1} \right) C
\end{aligned} \tag{9.1.4-III}$$

where C comes from applying inequality 9.1.4-II to $\phi^i(Q)$ for all i. This clearly goes to 0 as m (and thus n) goes to ∞ and thus the sequence is Cauchy and therefore converges.

Now that we know the canonical height is well-defined, we can show that it is uniquely determined by the following conditions:

THEOREM 9.1.4.3 (Néron-Tate). Suppose we are in the situation of definition 9.1.4.1. Then the canonical height function $\widehat{h}_{\phi,\mathscr{L}}$ is the unique function such that

(1) $\widehat{h}_{\phi,\mathscr{L}} = h_{\mathscr{L}} + O(1)$ and

(2) $\widehat{h}_{\phi,\mathscr{L}} \circ \phi = \lambda h_{\mathscr{L}}$.

Proof. Firstly, $\widehat{h}_{V,\phi,\mathscr{L}}$ satisfies property (1) since taking $m = 1$ and letting $n \to \infty$ in inequality 9.1.4-III, we get

$$\left| \widehat{h}_{\phi,\mathscr{L}}(P) - h_{\mathscr{L}}(P) \right| \leq \frac{C}{\lambda - 1}$$

for all $P \in V(\overline{\mathbb{Q}})$ which explicitly gives us our bound.

For property (2), we have

$$\begin{aligned}
\widehat{h}_{\phi,\mathscr{L}}(\phi(P)) &= \lim_{n\to\infty} \frac{1}{\lambda^n} h_{V,\mathscr{L}}(\phi^{n+1}(P)) \\
&= \lambda \left(\lim_{n\to\infty} \frac{1}{\lambda^{n+1}} h_{V,\mathscr{L}}(\phi^{n+1}(P)) \right) \\
&= \lambda \widehat{h}_{\phi,\mathscr{L}}(P)
\end{aligned}$$

for all $P \in V(\overline{\mathbb{Q}}$.

Now for uniqueness, suppose $\widehat{h}, \widehat{h}'$ are both functions satisfying properties (1) and (2). Let $g = \widehat{h} - \widehat{h}'$. By property (1), $g = O(1)$, i.e. there exists some C' such that $g(P) \leq C'$ for all $P \in V(\overline{\mathbb{Q}})$. By property

(2), $g \circ \phi = \lambda g$ and so iterating, we have that

$$|g(P)| = \frac{|g(\phi^n(P))|}{\lambda^n} \leq \frac{C'}{\lambda^n}.$$

But since $\lambda > 1$, this goes to 0 as $n \to \infty$ and so $|g(P)| = 0$ for all $P \in V(\overline{\mathbb{Q}})$, i.e., $\widehat{h} = \widehat{h}'$. □

As a first show of the usefulness of canonical heights, in the case that $\mathscr{L}$ is ample, we will prove some neat finiteness results. Before we begin, we will introduce some terminology from the theory of dynamical systems.

DEFINITION 9.1.4.4. Let S be a set and let $\phi : S \to S$ be a function and let ϕ^n be the n-fold composite. We call an element $P \in S$ **periodic for** ϕ if $\phi^n(P) = P$ for some $n \geq 1$ and **preperiodic for** ϕ if $\phi^m(P)$ is periodic for ϕ for some $m \geq 1$. Equivalently, P is preperiodic if the forward orbit

$$O_\phi^+(P) = \{P, \phi(P), \phi^2(P), \dots\}$$

is finite.

PROPOSITION 9.1.4.5. Suppose we are in the situation of definition 9.1.4.1 in particular, that there exists some line bundle $\mathscr{L}$ and some $\lambda \geq 2$ such that $\phi^*(\mathscr{L}) = \mathscr{L}^{\otimes \lambda}$, and assume further that $\mathscr{L}$ is ample.

(1) Let $P \in V(\overline{\mathbb{Q}})$. Then $\widehat{h}_{\phi,\mathscr{L}}(P) \geq 0$, and $\widehat{h}_{\phi,\mathscr{L}}(P) = 0$ if and only if P is preperiodic for ϕ.

(2) The set $\{P \in V(K) \mid P \text{ is preperiodic for } \phi\}$ is finite.

Proof. (1) Since $\mathscr{L}$ is ample, we may choose a representative for $h_{\mathscr{L}}$ which is non-negative. But then by the definition of $\widehat{h}_{\phi,\mathscr{L}}$ (9.1.4.1), it is clearly also non-negative.

Let $P \in V(\overline{k})$. By possibly extending K, we may assume that $P \in V(K)$. Now, suppose that P is preperiodic for ϕ. Since the set $\{P, \phi(P), \phi^2(P), \dots\}$ is finite, $h_{\mathscr{L}}(\phi^n(P)) \leq C$ for some $C > 0$ for all n. Thus, $\lambda^{-n} h_{\mathscr{L}}(\phi^n P) \to 0$ as $n \to \infty$. By definition, this implies that $\widehat{h}_{\phi,\mathscr{L}}(P) = 0$.

Conversely, suppose $\widehat{h}_{\phi,\mathscr{L}}(P) = 0$. Then, we have that $\widehat{h}_{\phi,\mathscr{L}}(\phi^n P) = \lambda^n \widehat{h}_{\phi,\mathscr{L}}(P) = 0$ for all n. So in particular for any $B \geq 0$, since $\phi^n(P) \in V(K)$ for all n, we have

$$O_\phi^+(P) = \{P, \phi(P), \phi^2(P), \dots\} \subset \{P \in V(K) \mid \widehat{h}_{\phi,\mathscr{L}}(P) \leq B\},$$

which, since $\mathscr{L}$ is ample, by the finiteness part of Theorem 9.1.3.5, is finite and thus, P is preperiodic for ϕ.

(2) This follows from the finiteness of the set $\{P \in V(K) \mid \widehat{h}_{\phi,\mathscr{L}}(P) \leq B\}$ for $B = 0$, and part (1).

There is a proof for general projective varieties (i.e. ones without the condition that there is a line bundle with is an eigenvector for ϕ of value ≥ 2). This is a classical result given in Northcott Paper.

This proof in HS is punted to a paper which is long and confusing. I will try to decipher it at some point and eit

□

In order to construct canonical heights, we require the existence of a line bundle $\mathscr{L}$ which is an eigenvector for ϕ^* with eigenvalue $\lambda \geq 2$. And furthermore, to apply the above proposition, we require

that $\mathscr{L}$ be ample. These may seem like odd requirements, but we will illustrate some important cases in which this occurs.

EXAMPLE 9.1.4.6. Consider the case $V = \mathbb{P}^n$ and let $\phi : \mathbb{P}^n \to \mathbb{P}^n$ be a morphism of degree $d > 1$. Then $\phi^*\mathscr{O}(1) \simeq \mathscr{O}^{\otimes d}$. Thus by Proposition 9.1.4.5(2),

$$\{P \in \mathbb{P}^n(K) \mid P \text{ is preperiodic for } \phi\}$$

is finite.

As a special case of this, consider the morphism $\phi(x_0 : \cdots : x_n) = (x_0^2 : \cdots : x_n^2)$. One can easily check that the corresponding canonical height function is just the height function defined in 9.1.2.1. Furthermore, the preperiodic points for ϕ are just the points for which for every $x_j \neq 0$, $\frac{x_i}{x_j}$ is a root of unity or 0 for each i. In this case, the finiteness of these point is just Kronecker's Theorem (Cor 9.1.2.8).

EXERCISE 9.1.4.A. Check the details in the above example.

EXAMPLE 9.1.4.7. Let A be an abelian variety over a number field K. Let $\mathscr{L} \in \mathrm{Pic}(A)$ be an ample, symmetric ($[-1]^*\mathscr{L} = \mathscr{L}$) line bundle. (Remember that the existence of such a line bundle is guaranteed by first taking an ample line bundle then considering $\mathscr{L} \otimes [-1]^*\mathscr{L}$). Then by theorem 1.1.4.1, we have that $[m]^*\mathscr{L} \simeq \mathscr{L}^{\otimes m^2}$, so considering $m > 1$, we may apply proposition 9.1.4.5 to conclude that there are only finitely many preperiodic points in $A(K)$ for $[m]$. But by definition, a point $P \in A(K)$ is preperiodic for $[m]$ if $m^iP = m^jP$ for some $i > j$. We claim that these are exactly the torsion points in $A(K)$. Firstly, these are clearly torsion as $(m^i - m^j)P = 0$. Conversely, suppose $P \in A(K)$ is torsion, so there is some n such that $nP = 0$. By the pigeonhole principle, there are some $i > j$ such that $m^i \equiv m^j \bmod n$, so P must be preperiodic with respect to m.

This implies that $A(K)_{\mathrm{tor}}$ is finite! This is significant as it proves the torsion part of the Mordell-Weil theorem.

§9.1.5 Heights on Abelian Varieties

We will begin by discussing basic results in the non-canonical case in the setting of abelian varieties. Before we get into that, we recall a basic definition which will be important throughout this section:

DEFINITION 9.1.5.1. Let A, B be abelian groups and suppose B is uniquely 2-divisible, i.e. 2-multiplication is an isomorphism (e.g. B is a field, in particular $\mathbb{R}$). A **quadratic function** $h : A \to B$ is a function satisfying

$$h(P+Q+R) - h(P+Q) - h(P+R) - h(Q+R) + h(P) + h(Q) + h(R) - h(0) = 0$$

for all $P, Q, R \in A$.

Equivalently, $h : A \to B$ is a quadratic function if the associated pairing $\langle \cdot, \cdot \rangle_h : A \times A \to B$ defined by

$$\langle P, Q \rangle_h = \frac{1}{2}(h(P+Q) - h(P) - h(Q) + h(0)),$$

is bilinear.

REMARK 9.1.5.2. Note that any quadratic function can be written uniquely as a sum of a quadratic form, a linear form, and a constant function. The associated quadratic form is given by $q(P) = \frac{1}{2}(h(P)+h(-P)-2h(0))$, the linear form is given by $\ell(P) = \frac{1}{2}(h(P)+h(-P))$, and of course the constant is given by $h(0)$. We leave the details of checking that these are indeed quadratic and linear to the reader.

Now that we have that out of the way, one might naturally ask how a (non-canonical) height function behaves with respect to the group law. This follows directly from the cubical structure of $\mathscr{L}$ (1.1.3.4):

PROPOSITION 9.1.5.3. Let $A/\overline{\mathbb{Q}}$ be an abelian variety and let $\mathscr{L}$ be a line bundle on A. Then for all $P, Q, R \in A(\overline{\mathbb{Q}})$,

$$h_{\mathscr{L}}(P+Q+R) - h_{\mathscr{L}}(P+Q) - h_{\mathscr{L}}(P+R) - h_{\mathscr{L}}(Q+R) + h_{\mathscr{L}}(P) + h_{\mathscr{L}}(Q) + h_{\mathscr{L}}(R) = O(1).$$

I.e., $h_{\mathscr{L}}$ behaves like a quadratic function with respect to addition on A up to a constant function.

Proof. Defining $m, m_{ij}, \mathrm{pr}_i : A \times A \times A \to A$ as the 3-fold addition, 2-fold addition in the i, j factors, and projection maps respectively (as in the proof of the cubical structure of $\mathscr{L}$), we have that

$$\mathscr{M} = m^*\mathscr{L} \otimes m_{12}^*\mathscr{L}^{-1} \otimes m_{13}^*\mathscr{L}^{-1} \otimes m_{23}^*\mathscr{L}^{-1} \otimes \mathrm{pr}_1^*(\mathscr{L}) \otimes \mathrm{pr}_2^*(\mathscr{L}) \otimes \mathrm{pr}_3^*(\mathscr{L})$$

is trivial. Thus, applying functoriality and additivity of theorem 9.1.3.5, we have

$$O(1) = h_{\mathscr{M}} = h_{\mathscr{L}} \circ m - h_{\mathscr{L}} \circ m_{12} - h_{\mathscr{L}} \circ m_{13} - h_{\mathscr{L}} \circ m_{23} + h_{\mathscr{L}} \circ \mathrm{pr}_1 + h_{\mathscr{L}} \circ \mathrm{pr}_2 + h_{\mathscr{L}} \circ \mathrm{pr}_3 .$$

plugging in (P, Q, R) into the above equation gives us our desired result. □

Either as a corollary of this result, or by our work in chapter 1, we may describe how height functions interact with $[n]$. We have already seen the special case of $\mathscr{L}$ symmetric in example 9.1.4.7!

COROLLARY 9.1.5.4. Let $[n] : A \to A$ be the multiplication by n map. Then

$$h_{\mathscr{L}} \circ [n] = \left(\frac{n^2+n}{2}\right) h_{\mathscr{L}} + \left(\frac{n^2-n}{2}\right) h_{\mathscr{L}} \circ [-1] + O(1).$$

In particular, for $\mathscr{L}$ symmetric (i.e., $[-1]^*\mathscr{L} = \mathscr{L}$),

$$h_{\mathscr{L}} \circ [n] = n^2 h_{\mathscr{L}} + O(1),$$

and if $\mathscr{L}$ is anti-symmetric (i.e., $[-1]^*\mathscr{L} = \mathscr{L}^{-1}$),

$$h_{\mathscr{L}} \circ [n] = n h_{\mathscr{L}} + O(1).$$

Proof. Apply theorem 1.1.4.1 and functoriality and additivity of theorem 9.1.3.5. Alternatively, apply proposition 9.1.5.3 to $Q = (n+1)P$ and $R = -P$ and induct upwards and downwards. □

To finish our discussion of non-canonical heights, we will give a relation which will be key proving important properties about associated canonical heights.

COROLLARY 9.1.5.5. Let $\mathscr{L}$ be a line bundle on A and $P, Q \in A(\overline{\mathbb{Q}})$.

(1) If $\mathscr{L}$ is symmetric then $h_{\mathscr{L}}$ satisfies the *parallelogram law*:

$$h_{\mathscr{L}}(P+Q) + h_{\mathscr{L}}(P-Q) = 2h_{\mathscr{L}}(P) + 2h_{\mathscr{L}}(Q) + O(1);$$

(2) if $\mathscr{L}$ is anti-symmetric then $h_{\mathscr{L}}$ is linear:

$$h_{\mathscr{L}}(P+Q) + h_{\mathscr{L}}(P-Q) = h_{\mathscr{L}}(P) + h_{\mathscr{L}}(Q) + O(1).$$

Proof. (1) Plugging $R = -Q$ in prop 9.1.5.3, we get

$$h_{\mathscr{L}}(P) - h_{\mathscr{L}}(P+Q) - h_{\mathscr{L}}(P-Q) - h_{\mathscr{L}}(0) + h_{\mathscr{L}}(P) + h_{\mathscr{L}}(Q) + h_{\mathscr{L}}(-Q) = O(1).$$

Noting that $h_{\mathscr{L}}(0)$ is constant and by functoriality in theorem 9.1.3.5, since $\mathscr{L}$ is symmetric, we have that $h_{\mathscr{L}}(-P) = h_{[-1]^*\mathscr{L}}(P) + O(1) = h_{\mathscr{L}}(P) + O(1)$. Thus, rearranging the above equation gives us the desired result.

(2) This time, plugging $R = -P - Q$ into prop 9.1.5.3, we get

$$h_{\mathscr{L}}(0) - h_{\mathscr{L}}(P+Q) - h_{\mathscr{L}}(-Q) - h_{\mathscr{L}}(-P) + h_{\mathscr{L}}(P) + h_{\mathscr{L}}(Q) + h_{\mathscr{L}}(-P-Q) = O(1);$$

By functoriality in theorem 9.1.3.5, since $\mathscr{L}$ is anti-symmetric, we have that $h_{\mathscr{L}}(-P) = h_{[-1]^*\mathscr{L}}(P) + O(1) = -h_{\mathscr{L}}(P) + O(1)$, so rearranging gives us our result. □

Now that we have set up some basic properties in the general Weil height setting, we can now explore associated canonical heights. All basic properties will be summarized in the following:

THEOREM 9.1.5.6 (Néron-Tate). Let A be an abelian variety defined over a number field K and let $\mathscr{L}$ be a *symmetric* line bundle on A. Then there is a height function

$$\widehat{h}_{\mathscr{L}} : A(\overline{\mathbb{Q}}) \to \mathbb{R}$$

called the **canonical height** on A relative to $\mathscr{L}$ which is the unique function such that:

(1) $\widehat{h}_{\mathscr{L}} = h_{\mathscr{L}} + O(1)$ and

(2) for all $m \in \mathbb{Z}$, $\widehat{h}_{\mathscr{L}} \circ [m] = m^2 h_{\mathscr{L}} + O(1)$.

Furthermore, it satisfies the following

(3) (Parallelogram Law) For all $P, Q \in A(\overline{\mathbb{Q}})$,

$$\widehat{h}_{\mathscr{L}}(P+Q) + \widehat{h}_{\mathscr{L}}(P-Q) = 2\widehat{h}_{\mathscr{L}}(P) + 2\widehat{h}_{\mathscr{L}}(Q)$$

(4) $\widehat{h}_{\mathscr{L}}$ is a quadratic form and thus has an associated bilinear form which we denote $\langle \cdot, \cdot \rangle_{\mathscr{L}} : A(\overline{\mathbb{Q}}) \times A(\overline{\mathbb{Q}}) \to \mathbb{R}$ defined by

$$\langle P, Q \rangle_{\mathscr{L}} = \frac{1}{2}\left(\widehat{h}_{\mathscr{L}}(P+Q) - \widehat{h}_{\mathscr{L}}(P) - \widehat{h}_{\mathscr{L}}(Q)\right).$$

REMARK 9.1.5.7. We will later do a similar construction for anti-symmetric line bundles and since for every line bundle $\mathscr{L}$, we can write $\mathscr{L}^{\otimes 2}$ as a tensor of a symmetric and anti-symmetric line bundle (namely $\mathscr{L} \otimes [-1]^*\mathscr{L}$ and $\mathscr{L} \otimes [-1]^*\mathscr{L}^{-1}$), we will extend to all line bundles via linearity.

Proof. Consider $\widehat{h}_{\mathscr{L}} := \widehat{h}_{[m],\mathscr{L}}$. Properties (1) and (2) for m as well as uniqueness follow immediately from theorem 9.1.4.3 (noting that $[m]^*\mathscr{L} \simeq \mathscr{L}^{\otimes m^2}$ by theorem 1.1.4.1). So, we just need to show that the $\widehat{h}_{[m],\mathscr{L}}$'s coincide for all $m \in \mathbb{Z}$.

To show this, first define $\widehat{h}_{\mathscr{L}} := \widehat{h}_{[2],\mathscr{L}}$ (we choose 2 simply for ease of notation). From corollary 9.1.5.4, we get that $h_{\mathscr{L}}([m]P) = m^2 h_{\mathscr{L}}(P) + O(1)$. Using the definition of $\widehat{h}_{[2],\mathscr{L}}$, we get

$$\begin{aligned}\widehat{h}_{\mathscr{L}}([m]P) &= \lim_{n\to\infty} \frac{1}{4^n} h_{\mathscr{L}}([2^n][m]P) \\ &= \lim_{n\to\infty} \frac{1}{4^n} \left(m^2 h_{\mathscr{L}}([2^n]P) + O(1)\right) \\ &= m^2 \widehat{h}_{\mathscr{L}}(P).\end{aligned}$$

Note that we crucially use that $[m]$ and $[2^n]$ commute. Thus, we get (2) for general m.

Now for (3), we use the parallelogram law for non-canonical heights:

$$h_{\mathscr{L}}(P+Q) + h_{\mathscr{L}}(P-Q) = 2h_{\mathscr{L}}(P) + 2h_{\mathscr{L}}(Q) + O(1).$$

As the constant in $O(1)$ does not depend on P, Q, we may replace them by $[2^n]P, [2^n]Q$ respectively, divide by $\frac{1}{4^n}$ and take the limit as $n \to \infty$, getting rid of the $O(1)$ term and giving us our desired result.

To prove (4), we use the identity from proposition 9.1.5.3:

$$h_{\mathscr{L}}(P+Q+R) - h_{\mathscr{L}}(P+Q) - h_{\mathscr{L}}(P+R) - h_{\mathscr{L}}(Q+R) + h_{\mathscr{L}}(P) + h_{\mathscr{L}}(Q) + h_{\mathscr{L}}(R) = O(1).$$

Again, since the $O(1)$ term doesn't depend on P, Q, R, replacing them by $[2^n]P, [2^n]Q, [2^n]R$ respectively, multiplying by $\frac{1}{4^n}$, and taking the limit as $n \to \infty$ gets rid of the $O(1)$ term and shows that $\widehat{h}_{\mathscr{L}}$ is a quadratic *function*. The fact that it is a quadratic *form* then follows from the property that $\widehat{h}_{\mathscr{L}} \circ [m] = m^2 \widehat{h}_{\mathscr{L}}$ for all m.

□

PROPOSITION 9.1.5.8. Suppose we are in the situation of theorem 9.1.5.6 and suppose further that $\mathscr{L}$ is also ample. Then

(1) $\widehat{h}_{\mathscr{L}} \geq 0$,

(2) $\widehat{h}_{\mathscr{L}}(P) = 0$ if and only if P is torsion

(3) and more generally, $\widehat{h}_{\mathscr{L}}$ extends to a positive definite quadratic form on $A(\overline{\mathbb{Q}}) \otimes \mathbb{R}$.

REMARK 9.1.5.9. Property (3) does not follow from (1) and (2). For example, consider the group $\Lambda := \mathbb{Z} + \sqrt{2}\mathbb{Z}$ and let $q : \Lambda \to \mathbb{R}$ be the square of the absolute value as a subset of $\mathbb{R}$, i.e. $a + b\sqrt{2} \mapsto a^2 + 2b^2 + 2ab\sqrt{2}$. This is clearly positive definite since $\sqrt{2}$ is irrational, but the induced map $q : \lambda \otimes \mathbb{R} \to \mathbb{R}$

is not: consider $(a, b) = (-\sqrt{2}, 1)$. This arises as an issue because the image $q(\Lambda) \subset \mathbb{R}$ is not discrete (see cor 9.1.5.11).

Proof. Properties (1) and (2) follow immediately from example 9.1.4.7.

To show property (3), assume that $x \in A(\overline{\mathbb{Q}}) \otimes \mathbb{R}$ such that $\widehat{h}_{\mathscr{L}}(x) = 0$. Then $x = \sum_{i=1}^{n} a_i P_i$ with $P_i \in A(\overline{\mathbb{Q}})$. We know from properties (1) and (2) that $\widehat{h}_{\mathscr{L}}$ is positive definite on the lattice $\Lambda = \mathrm{Span}_{\mathbb{Z}}\{P_i\}_{i=1}^{r}$ since torsion points are in the kernel of the map $A(\overline{\mathbb{Q}}) \to A(\overline{\mathbb{Q}}) \otimes \mathbb{R}$. In addition, by enlarging K such that $P_i \in A(K)$ for all i, we know that for each $B \geq 0$,

$$\{Q \in \Lambda \mid \widehat{h}_{\mathscr{L}}(Q) \leq B\} \subset \{Q \in A(K) \mid \widehat{h}_{\mathscr{L}}(Q) \leq B\}$$

is finite. Thus, by a corollary of Minkowski's theorem (9.1.5.11), $\widehat{h}_{\mathscr{L}}$ is positive definite on $\Lambda \otimes \mathbb{R}$ and thus $x = 0$. □

To finish the proof above, we need Minkowski's Theorem on lattice points in symmetric domains (the proof of which is a very classic result, but we include the proof for completion) and a corollary of it about quadratic forms. This is a very classical result, but we include the proof for completion.

> THEOREM 9.1.5.10 (Minkowski's Theorem). Let Λ be a lattice in $\mathbb{R}^n$, F a fundamental domain for $\mathbb{R}^n/\Lambda$ and let $U \subset \mathbb{R}^n$ be a symmetric convex set (i.e., $U = -U$ and U contains the line between any two points in it). Then if $\mathrm{Vol}(U) > 2^n \mathrm{Vol}(F)$, $U \cap \Lambda$ contains a non-zero lattice point.

Proof. It suffices to show that there are lattice points $\lambda_1, \lambda_2 \in \Lambda$ such that

$$(\frac{1}{2}U + \lambda_1) \cap (\frac{1}{2}U + \lambda_2) \neq \emptyset.$$

This is because if there is a point in that set, i.e. points $x_1, x_2 \in U$ such that $\frac{1}{2}x_1 + \lambda_1 = \frac{1}{2}x_2 + \lambda_2$, then we have a lattice point $\lambda = \lambda_1 - \lambda_2 = \frac{1}{2}x_1 - \frac{1}{2}x_2$, which is the midpoint of the line between x_1 and x_2.

Now, to show that such λ_1, λ_2 exist, suppose the opposite so that $\frac{1}{2} + \lambda$ are pairwise disjoint for all $\lambda \in \Lambda$. Thus, so would their intersection with F, so that we have

$$\begin{aligned} \mathrm{Vol}\, F &\geq \sum_{\lambda \in \Lambda} \mathrm{Vol}(F \cap (\frac{1}{2}U + \lambda) \\ &= \sum_{\lambda \in \Lambda} \mathrm{Vol}\left((F - \lambda) \cap \frac{1}{2}U\right) \\ &= \mathrm{Vol}(\frac{1}{2}U) \qquad \text{(9.1.5-I)} \\ &= \frac{1}{2^n}\mathrm{Vol}(U). \end{aligned}$$

where the equality 9.1.5-I follows from the fact that $\cup_{\lambda \in \Lambda}(F - \lambda) = \mathbb{R}^n$ by the definition of fundamental domains. This contradicts the assumption. □

From this, we get the following fact about quadratic forms which we applied above to finish the proof of prop 9.1.5.8.

> COROLLARY 9.1.5.11. Let Λ be a free abelian group of finite rank. Let $q : \Lambda \to \mathbb{R}$ be a positive definite

quadratic form such that for all $B \geq 0$,

$$\{\lambda \in \Lambda \mid q(\lambda) \leq B\}$$

is finite. Then q extends to a positive definite quadratic form on $\Lambda \otimes \mathbb{R}$.

Proof. Let $\overline{q} : A \otimes \mathbb{R} \to \mathbb{R}$ denote the extension of q to $A \otimes \mathbb{R}$ and define $r = \dim A \otimes \mathbb{R}$. It is a standard result in linear algebra that any quadratic form over $\mathbb{R}$ can be written in the form

$$q(x) = x_1^2 + \cdots + x_s^2 - x_{s+1}^2 - \cdots - x_{s+t}^2$$

where $s + t \leq r$. Note that q is positive definite if and only if $s = r$.

Let

$$L = \inf\{q(\lambda) \mid \lambda \in \Lambda, \lambda \neq 0\}.$$

Let $\delta, \epsilon > 0$ and consider the sets

$$U_{\delta,\epsilon} = \left\{ (x_1, \ldots, x_r) \mid \sum_{i=1}^{s} x_i^2 \leq \delta \text{ and } \sum_{j=1}^{t} x_{s+j}^2 \leq \epsilon \right\}.$$

This is clearly a symmetric convex set. Furthermore, from the definition of L, it is clear that $U_{L/2,\epsilon} \cap \Lambda = \{0\}$ for all $\epsilon > 0$. Thus by Minkowski's theorem (9.1.5.10), $\mathrm{Vol}(U_{L/2,\epsilon})$ is bounded as $\epsilon \to \infty$. But, if $s + t < r$, $U_{L/2,\epsilon}$ is infinite for all ϵ and if $s + t = r$, $\mathrm{Vol}(U_{L/2,\epsilon})$ grows like $\epsilon^{t/2}$ as $\epsilon \to \infty$. Thus, we must have $s = r$, i.e. $\overline{q}$ is positive definite. □

We have just developed the theory of canonical heights in detail for symmetric line bundles. Corollary 9.1.5.4 indicates that we can apply the theory of canonical heights to develop a similar theory for anti-symmetric line bundles. Since the square of any line bundle can be written as a product of a symmetric and anti-symmetric line bundles, we can then combine these two theories to give us a theory of heights for general line bundles on abelian varieties via linearity.

THEOREM 9.1.5.12. Let A be an abelian variety defined over a number field K and let $\mathscr{L}$ be an *anti-symmetric* line bundle on A. Then there is a height function

$$\widehat{h}_{\mathscr{L}} : A(\overline{\mathbb{Q}}) \to \mathbb{R}$$

called the **canonical height** on A relative to $\mathscr{L}$ which is the unique function such that:

(1) $\widehat{h}_{\mathscr{L}} = h_{\mathscr{L}} + O(1)$ and

(2) $\widehat{h}_{\mathscr{L}}(P + Q) = \widehat{h}_{\mathscr{L}}(P) + \widehat{h}_{\mathscr{L}}(Q)$.

Proof. This proof is essentially the same as the proof of theorem 9.1.5.6, so we leave it to the reader! Note in particular that property (2) implies $\widehat{h}_{\mathscr{L}}([m]P) = m\widehat{h}_{\mathscr{L}}(P)$ for all $m \in \mathbb{Z}$, allowing us to use the Néron-Tate canonical heights construction. □

EXERCISE 9.1.5.A. Prove theorem 9.1.5.12.

Now we may combine the two in the following:

THEOREM 9.1.5.13. Let A/K be an abelian variety defined over a number field, and let $\mathscr{L} \in \mathrm{Pic}(A)$. Then there exists a homomorphism

$$\begin{aligned} \mathrm{Pic}(A) &\to \{\text{quadratic functions } A(\overline{\mathbb{Q}}) \to \mathbb{R}\} \\ \mathscr{L} &\mapsto \widehat{h}_{\mathscr{L}} \end{aligned}$$

which is the unique homomorphism through which the map HM (see theorem 9.1.3.5) factors through. $\widehat{h}_{\mathscr{L}}$ is called the **canonical height function** on A with respect to $\mathscr{L}$.

Furthermore, it satisfies the following properties:

(1) (Functoriality) Let B/K be another abelian variety and $\phi : B \to A$ a morphism (not necessarily a homomorphism!). Then

$$\widehat{h}_{B,\phi^*\mathscr{L}} = \widehat{h}_{A,\mathscr{L}} \circ \phi - \widehat{h}_{A,\mathscr{L}}(\phi(0)),$$

where the subscript A, B are added for emphasis of the domain of the respective functions.

(2) (Decomposition) There is a unique quadratic form $\widehat{q}_{\mathscr{L}} : A(\overline{\mathbb{Q}}) \to \mathbb{R}$ and a unique linear form $\widehat{\ell}_{\mathscr{L}} : A(\overline{\mathbb{Q}}) \to \mathbb{R}$ such that

$$\widehat{h}_{\mathscr{L}} = \widehat{q}_{\mathscr{L}} + \widehat{\ell}_{\mathscr{L}}.$$

More precisely, using the canonical heights for symmetric and anti-symmetric line bundles (theorems 9.1.5.6 and 9.1.5.12 resp.), we have

$$\widehat{q}_{\mathscr{L}} = \frac{1}{2}\widehat{h}_{\mathscr{L}\otimes[-1]^*\mathscr{L}} \quad \text{and} \quad \widehat{q}_{\mathscr{L}} = \frac{1}{2}\widehat{h}_{\mathscr{L}\otimes[-1]^*\mathscr{L}^{-1}}.$$

Hence, if $\mathscr{L}$ is symmetric, $\widehat{h}_{\mathscr{L}} = \widehat{q}_{\mathscr{L}}$ and if $\mathscr{L}$ is anti-symmetric $\widehat{h}_{\mathscr{L}} = \widehat{\ell}_{\mathscr{L}}$.

Proof. We claim that $\widehat{h}_{\mathscr{L}}$ is uniquely categorized by the following three properties (each of which is clearly implied by the theorem statement):

(i) $\widehat{h}_{\mathscr{L}}$ is a quadratic function,

(ii) $\widehat{h}_{\mathscr{L}} = h_{\mathscr{L}} + O(1)$, and

(iii) $\widehat{h}_{\mathscr{L}}(0) = 0$.

We will start by proving existence and uniqueness of functions satisfying these properties.

To begin, we will construct $\widehat{h}_{\mathscr{L}}$. Consider the line bundles $\mathscr{L}^+ = \mathscr{L} \otimes [-1]^*\mathscr{L}$ and $\mathscr{L}^- = \mathscr{L} \otimes [-1]^*\mathscr{L}^{-1}$ which are symmetric and anti-symmetric respectively. Theorems 9.1.5.6 and 9.1.5.12 give us unique functions $\widehat{h}_{\mathscr{L}^+}, \widehat{h}_{\mathscr{L}^-}$ which are respectively a quadratic form and a linear form on $A(\overline{\mathbb{Q}}$. Thus, inspired by the relation $\mathscr{L}^{\otimes 2} = \mathscr{L}^+ \otimes \mathscr{L}^-$, we define

$$\widehat{h}_{\mathscr{L}} := \frac{1}{2}(\widehat{h}_{\mathscr{L}^+} + \widehat{h}_{\mathscr{L}^-}).$$

This is clearly quadratic and by the additivity of non-canonical heights in theorem 9.1.3.5,

$$\begin{aligned} 2\widehat{h}_{\mathscr{L}} &= \widehat{h}_{\mathscr{L}^+} + \widehat{h}_{\mathscr{L}^-} \\ &= \widehat{h}_{\mathscr{L}^+\otimes\mathscr{L}^-} + O(1) \\ &= 2h_{\mathscr{L}} + O(1). \end{aligned}$$

Finally, since $\widehat{h}_{\mathscr{L}^+}(0) = \widehat{h}_{\mathscr{L}^-}(0) = 0$, $\widehat{h}_{\mathscr{L}}(0) = 0$. Thus, we have existence of a function satisfying properties (i) - (iii).

For uniqueness, suppose $\widehat{h}'_{\mathscr{L}}$ is another quadratic function with $\widehat{h}'_{\mathscr{L}} = h_{\mathscr{L}} + O(1)$. Then the difference $f := \widehat{h}'_{\mathscr{L}} - \widehat{h}_{\mathscr{L}}$ is a bounded quadratic function. In particular, the associated bilinear form $\langle \cdot, \cdot \rangle_f$ (see def 9.1.5.1) is bounded. However, we have that $\langle nP, nQ \rangle_f = n^2 \langle P, Q \rangle_f$ for all n, so we must have $\langle P, Q \rangle_f = 0$ for all $P, Q \in A(\overline{\mathbb{Q}})$. By definition $\langle P, Q \rangle_f = \frac{1}{2}(f(P+Q) - f(P) - f(Q) + f(0))$, so this implies $f(P+Q) + f(0) = f(P) + f(Q)$ and since $f(0) = 0$, $f(P+Q) = f(P) + f(Q)$ for all $P, Q \in A(\overline{\mathbb{Q}})$. In particular, $f(nP) = nf(P)$ for all n, which implies $f(P) = 0$ for all $P \in A(\overline{\mathbb{Q}})$, so $\widehat{h}'_{\mathscr{L}} = \widehat{h}_{\mathscr{L}}$ which shows uniqueness.

Now we will show that the map $\mathscr{L} \mapsto \widehat{h}_{\mathscr{L}}$ is a well-defined homomorphism. For well-definedness, suppose $\mathscr{L} \simeq \mathscr{L}'$. By well-definedenss of non-canonical heights (see prop 9.1.3.7), we have that $\widehat{h}_{\mathscr{L}'} = h_{\mathscr{L}} + O(1)$, so satisfies property (ii). It clearly also satisfies properties (i) and (iii), so by uniqueness $\widehat{h}_{\mathscr{L}'} = \widehat{h}_{\mathscr{L}}$.

For additivity, let $\mathscr{L}, \mathscr{M} \in \mathrm{Pic}(A)$. We note that by additivity of non-canonical heights (see prop 9.1.3.8), $\widehat{h}_{\mathscr{L}} + \widehat{h}_{\mathscr{M}} = h_{\mathscr{L}\otimes\mathscr{M}} + O(1)$ and since the sum of quadratic functions is quadratic and $\widehat{h}_{\mathscr{L}}(0) + \widehat{h}_{\mathscr{M}}(0) = 0$, this satisfies the properties (i)-(iii) for the function $\widehat{h}_{\mathscr{L}\otimes\mathscr{M}}$, so by uniqueness, we have $\widehat{h}_{\mathscr{L}} + \widehat{h}_{\mathscr{M}} = \widehat{h}_{\mathscr{L}\otimes\mathscr{M}}$.

For functoriality, let $\phi : B \to A$ be a morphism of abelian varieties. note that by functoriality of non-canonical heights (thm 9.1.3.5), we have

$$f := \widehat{h}_{\mathscr{L}} \circ \phi - \widehat{h}_{\mathscr{L}}(\phi(0)) = h_{\phi^*\mathscr{L}} + O(1).$$

By construction, we have $f(0) = 0$ and finally, by corollary 1.1.1.3, ϕ is the composition of a homomorphism and a translation, so f is a quadratic form. Thus, as it satisfies properties (i)-(iii) with respect to $\phi^*\mathscr{L}$, by the uniqueness of the canonical height function, $\widehat{h}_{\phi^*\mathscr{L}} = \widehat{h}_{\mathscr{L}} \circ \phi - \widehat{h}_{\mathscr{L}}(\phi(0))$.

Finally, for the decomposition into linear and quadratic parts, we have already exhibited the decomposition above and by the uniqueness of the decomposition of quadratic functions into quadratic forms, linear forms, and constants (noted in 9.1.5.2), this is unique. Since when $\mathscr{L}$ is symmetric, $\mathscr{L}^+ = 2\mathscr{L}$ and $\mathscr{L}^- = \mathscr{O}_A$, and when it's anti-symmetric, $\mathscr{L}^- = 2\mathscr{L}$ and $\mathscr{L}^+ = \mathscr{O}_A$, the final claims are clear. □

Recall that all line bundles in $\mathrm{Pic}^0_{A/K}(K) = \widehat{A}(K)$ are anti-symmetric (see remark 1.2.3.4). Thus by theorem 9.1.5.13, the canonical height $\widehat{h}_{\mathscr{L}}$ is a homomorphism $A(\overline{\mathbb{Q}}) \to \mathbb{R}$. The following theorem will relate the natural pairing $A(\overline{\mathbb{Q}}) \times \widehat{A}(\overline{\mathbb{Q}}) \to \mathbb{R}$ given by evaluation of heights $(P, \mathscr{L}) \to \widehat{h}_{\mathscr{L}}(P)$ to the Poincaré bundle, giving us another relation between the Poincaré bundle and bilinear pairings (see sections 1.3.1, 2.1.3 for the other related bilinear forms we've explored).

THEOREM 9.1.5.14. Define the **canonical height pairing**

$$[\cdot, \cdot]_A : A(\overline{\mathbb{Q}}) \times \widehat{A}(\overline{\mathbb{Q}}) \to \mathbb{R}$$
$$(P, \mathscr{L}) \mapsto \widehat{h}_{\mathscr{L}}(P).$$

Then $[\cdot, \cdot]_A$ is bilinear and the kernel on either side consists of torsion points. Furthermore if $\mathscr{P}_A \in$

Pic$(A \times \widehat{A})$ is the Poincaré bundle (see definition 1.2.1.1) then,

$$[P, \mathscr{L}]_A = \widehat{h}_{A \times \widehat{A}, \mathscr{P}_A}(P, \mathscr{L})$$

for all $(P, \mathscr{L}) \in (A \times \widehat{A})(\overline{\mathbb{Q}}) = A(\overline{\mathbb{Q}}) \times \widehat{A}(\overline{\mathbb{Q}})$.

Proof. Since line bundles in $\widehat{A}(\overline{\mathbb{Q}})$ are anti-symmetric by remark 1.2.3.4, theorem 9.1.5.13(2) gives us linearity in the first coordinate and linearity in the second follows from the fact that the map in that theorem is a homomorphism.

Now to handle the claim that the kernel on both sides are the torsion points. Firstly, we note that by 1.2.3.1, any line bundle $\mathscr{L} \in \widehat{A}(\overline{\mathbb{Q}})$ can be written in the form

$$\mathscr{L} \simeq t_P^* \mathscr{N} \otimes \mathscr{N}^{-1}$$

for (any) ample $\mathscr{N}$ and some $P \in A(\overline{\mathbb{Q}})$. Conversely, any line bundle of the form $t_P^* \mathscr{N} \otimes \mathscr{N}^{-1}$ is contained in $\widehat{A}(\overline{\mathbb{Q}})$. So, fix a *symmetric* ample $\mathscr{N}$.

Assume that $P \in A(\overline{\mathbb{Q}})$ such that $[P, \mathscr{L}]_A = 0$ for all $\mathscr{L} \in \widehat{A}(\overline{\mathbb{Q}})$. Then $[P, t_Q^* \mathscr{N} \otimes \mathscr{N}^{-1}]_A = 0$ for all Q. Using this and the standard properties of the canonical height function (from theorem 9.1.5.13), we have

$$\begin{aligned} 0 &= [P, t_Q^* \mathscr{N} \otimes \mathscr{N}^{-1}]_A \\ &= \widehat{h}_{t_Q^* \mathscr{N} \otimes \mathscr{N}^{-1}}(P) \\ &= \widehat{h}_{t_Q^* \mathscr{N}}(P) - \widehat{h}_{\mathscr{N}}(P) \\ &= \widehat{h}_{\mathscr{N}}(P+Q) - \widehat{h}_{\mathscr{N}}(Q) - \widehat{h}_{\mathscr{N}}(P). \end{aligned}$$

This holds for all $Q \in A(\overline{\mathbb{Q}})$. In particular, we may take $Q = P$. Since $\mathscr{N}$ is assumed to be symmetric, $\widehat{h}_{\mathscr{N}}(2P) = 4\widehat{h}_{\mathscr{N}}(P)$. Combining this with the above equation gives us $2\widehat{h}_{\mathscr{N}}(P) = 0$. But since $\mathscr{N}$ is ample, proposition 9.1.5.8(2) tells us that P is torsion.

For the other side, fix a line bundle $\mathscr{L} \in \widehat{A}(\overline{\mathbb{Q}})$ and suppose $[P, \mathscr{L}]_A = 0$ for all $P \in A(\overline{\mathbb{Q}})$. By the surjectivity of $\phi_{\mathscr{N}}$, we may write $\mathscr{L} = t_Q^* \mathscr{N} \otimes \mathscr{N}^{-1}$. By the same calculation as above, we have that

$$0 = [P, \mathscr{L}]_A = \widehat{h}_{\mathscr{N}}(P+Q) - \widehat{h}_{\mathscr{N}}(Q) - \widehat{h}_{\mathscr{N}}(P)$$

for all $P \in A(\overline{\mathbb{Q}})$. As before, letting $P = Q$ gives us $\widehat{h}_{\mathscr{N}}(Q) = 0$, so Q is torsion. But $\phi_{\mathscr{N}}$ is a homomorphism, so the image $\phi_{\mathscr{N}}(Q) = \mathscr{L}$ must also be torsion.

We will now prove

$$[P, \mathscr{L}]_A \stackrel{?}{=} \widehat{h}_{A \times \widehat{A}, \mathscr{P}_A}(P, \mathscr{L}).$$

Let $\iota_{\mathscr{L}} : A \to A \times \widehat{A}$ be the inclusion via $a \mapsto (a, \mathscr{L})$ for $a \in A$. By the definition of $\mathscr{P}_A$, $\iota_{\mathscr{L}}^* \mathscr{P}_A \simeq \mathscr{L}$. Thus, via the standard properties of canonical heights,

$$\begin{aligned} [P, \mathscr{L}]_A &= \widehat{h}_{\mathscr{L}}(P) \\ &= \widehat{h}_{\iota_{\mathscr{L}}^* \mathscr{P}_A} \\ &= \widehat{h}_{\mathscr{P}_A}(\iota_{\mathscr{L}}(P)) - \widehat{h}_{\mathscr{P}_A}(\iota_{\mathscr{L}}(0)) \\ &= \widehat{h}_{\mathscr{P}_A}(P, \mathscr{L}) - \widehat{h}_{\mathscr{P}_A}(0, \mathscr{L}). \end{aligned}$$

so we just need to show $\widehat{h}_{\mathscr{P}_A}(0, \mathscr{L}) \stackrel{?}{=} 0$. First, define $\widehat{\iota}_0 : \widehat{A} \to A \times \widehat{A}$ via $\mathscr{L} \mapsto (0, \mathscr{L})$. Again by the

definition of $\mathscr{P}_A$, $\widehat{\iota}_0^*\mathscr{P}_A = 0$ (thought of as a line bundle on $\widehat{A}$ by duality). Thus we have

$$\begin{aligned} 0 &= \widehat{h}_{\widehat{\iota}_0^*\mathscr{P}_A}(\mathscr{L}) \\ &= \widehat{h}_{\mathscr{P}_A}(\widehat{\iota}_0(\mathscr{L})) - \widehat{h}_{\mathscr{P}_A}(\widehat{\iota}_0(\mathscr{O}_A)) \\ &= \widehat{h}_{\mathscr{P}_A}(0, \mathscr{L}) - \widehat{h}_{\mathscr{P}_A}(0, \mathscr{O}_A) \\ &= \widehat{h}_{\mathscr{P}_A}(0, \mathscr{L}). \end{aligned}$$

□

We finish this section with a neat application of canonical heights on abelian varieties back to Weil heights on general smooth projective varieties which is a generalization of corollary 9.1.3.14.

PROPOSITION 9.1.5.15. Let V/K be a smooth projective variety over a number field K. Let $\mathscr{L}, \mathscr{M} \in \mathrm{Pic}(V)$ with $\mathscr{L}$ ample and $\mathscr{M}$ algebraically equivalent to $\mathscr{O}_V$. Then there is a constant $c > 0$ (dependent on the choice of $h_{\mathscr{L}}$ and $h_{\mathscr{M}}$) such that

$$h_{\mathscr{L}} \leq c\sqrt{h_{\mathscr{M}} + 1}.$$

Proof. Finish Later! □

Section 9.2

The Mordell-Weil Theorem

Now that we have developed the theory of heights on abelian varieties, we are equipped to prove the Mordell-Weil theorem. We will do this in two steps: the weak Mordell-Weil theorem, then the full theorem deduced from the weak version. These will be given in reverse order simply because the full Mordell-Weil theorem is an easy corollary of the weak version using our results in the previous section on heights. Thus, it is a natural jumping off point to apply all of the work we just did! On the other hand, the proof of the weak Mordell-Weil theorem is much longer and will rely heavily on Galois cohomology machinery rather than the topics built up throughout the main parts of this book.

§9.2.1 Mordell-Weil from Weak Mordell-Weil

Now that we have developed the theory of heights on abelian varieties, we are now equipped to prove the Mordell-Weil theorem given the *weak* Mordell-Weil theorem:

> **Theorem 9.2.1.1 (Weak Mordell-Weil).** Let A be an abelian variety over a number field K. Then the abelian group
>
> $$A(K)/nA(K)$$
>
> is finite for all $n \in \mathbb{Z}$.

We will prove this in the following subsection (9.2.2). Assuming this, it is not difficult to prove:

> **Theorem 9.2.1.2 (Mordell-Weil).** Let A be an abelian variety over a number field K. Then $A(K)$ is a finitely generated abelian group.

Proof. Choose a symmetric ample line bundle $\mathscr{L}$ on A (e.g. choose an ample line bundle $\mathscr{M}$ and consider $\mathscr{L} = \mathscr{M} \otimes [-1]^*\mathscr{M}$). Then we know the Néron-Tate canonical height $\widehat{h}_{\mathscr{L}}$ restricts to a non-negative quadratic form on $A(K)$ with the property that for all $B > 0$, the set

$$\{P \in A(K) \mid \widehat{h}(P) \leq B\}$$

is finite (see Theorem 9.1.5.6). Then the result will follow from Lemma 9.2.1.3 below, which generalizes this situation. □

> **Lemma 9.2.1.3 (Descent Lemma).** Let A be an abelian group and $q : A \to \mathbb{R}$ be a quadratic form on A with the property that for all $B > 0$, the set
>
> $$\{a \in A \mid q(g) \leq B\}.$$
>
> Assume that there is some $n \geq 2$ such that A/nA is finite. Then A is finitely generated.
>
> More precisely, let $g_1, \ldots, g_m$ be a set of generators for A/nA and let $B_0 = \max_i\{q(g_i)\}$. Then A

is generated by the finite set
$$S = \{a \in A \mid q(a) \leq B_0\}.$$

Proof. The idea of this proof is to show that if $q(a) > B_0$, then we can subtract off some copies of g_i to make it an m-multiple of something smaller, and repeat this process until we find an element inside of S which, together with our g_i's generates g.

Note that we must have that $q(a) \geq 0$ for all $a \in A$ as otherwise, if there was some b such that $q(b) < 0$, then $\{b, 2b, \dots\} \subset \{a \in A \mid q(a) \leq 0\}$ must be finite, so b is torsion, but any torsion point is killed by q (suppose k kills b, then $q(h) = \frac{1}{k^2}q(kh) = \frac{1}{k^2}q(0) = 0$). Thus, q is positive semidefinite on A.

Define $|a| = \sqrt{q(a)}$ for all $a \in A$. And for ease of notation, let $B_0' := \sqrt{B_0} = \max_i\{|g_i|\}$ and note that $S = \{a \in A \mid |a| \leq B_0'\}$. Now consider some $a_0 \in A$. If $a_0 \in S$, we are done. If not, then $|a_0| > B_0'$, and if g_j is the coset representative of a_0 in A/nA, we may write $a_0 = g_{j_0} + na_1$ for some $a_1 \in A$. Then we have

$$\begin{aligned} n|a_1| &= |a_0 - g_{j_0}| \\ &\leq |a_0| + |g_{j_0}| \\ &< 2|a_0| \end{aligned} \tag{9.2.1-I}$$

Where inequality 9.2.1-I follows from the triangle inequality (which holds for $|\cdot|$ by 9.2.1.4 below). Thus, since we assumed $n \leq 2$, we have that $|a_1| < |a_0|$. If $a_1 \in S$, then we have a way to write a_0 as a linear combination of elements in S, so we are done. If not, then running the same procedure on a_1 gives us some a_2 such that $a_1 = g_{j_1} + na_2$ and $|a_2| < |a_1|$. Continuing in this fashion, we find a chain

$$|a_0| > |a_1| > |a_2| > \dots$$

which must be finite since A has finitely many elements of bounded size by assumption. Thus, the sequence terminates with an element in S. Expanding this back out gives us an expression for a_0 as a linear combination of elements in S. □

LEMMA 9.2.1.4. If $q : A \to \mathbb{R}$ is a positive semidefinite quadratic form, then $|\cdot| := \sqrt{q}$ satisfies the triangle inequality.

Proof. It is a classic result in linear algebra that the square root of a positive semidefinite quadratic form is a seminorm and thus this is obvious on the extension of q to $A \otimes_{\mathbb{Z}} \mathbb{R}$. Thus, fixing an isomorphism $A \simeq A_{\text{tf}} \oplus A_{\text{tors}}$ where A_{tf} is torsion-free and A_{tors} is torsion, we have that $|\cdot|$ satisfies the triangle inequality when restricted to A_{tf}. Thus, it remains to check it for torsion points. But, for a decomposition of an element $a = a_{\text{tf}} + a_{\text{tors}}$ in A, $q(a_{\text{tf}} + a_{\text{tors}}) = q(a_{\text{tf}})$. This is because there is some $n > 0$ such that $na_{\text{tors}} = 0$ and so
$$q(a_{\text{tf}} + a_{\text{tors}}) = \frac{1}{n^2}q(n(a_{\text{tf}} + a_{\text{tors}})) = \frac{1}{n^2}q(na_{\text{tf}}) = q(a_{\text{tf}}),$$
so in particular $|a| = |a_{\text{tf}}|$. Then for $a, b \in G$
$$|a + b| = |a_{\text{tf}} + b_{\text{tf}}| \leq |a_{\text{tf}}| + |b_{\text{tf}}| = |a| + |b|.$$
□

Add exposition later on uses of this/ work towards actually computing number of points of a size/ rank

§9.2.2 Weak Mordell-Weil

The goal of this section is to prove Theorem 9.2.1.1. We will take a standard approach to proving the weak Mordell-Weil Theorem which will make heavy use of Galois cohomology. We point the reader to Appendix B to review, learn, or find references on this topic. There are many different ways to prove this theorem (and multiple proofs for many of the theorems in this), but we chose this way because it gives the most constructive way to bound the size of $A(K)/nA(K)$ and introduce objects of much interest in number theory: Selmer groups and Tate-Shafarevich groups. If the reader is curious about other proofs, they can see [HS00, Sections C.1, C.2].

The rough idea of the proof is to inject $A(K)/nA(K)$ into a subgroup of some Galois cohomology group which we will then prove is finite. The upside of this method is that, by studying this cohomology group, we may bound the size of this group and thus the rank of $A(K)$. However, this is in general very difficult due to the difficulty of explicitly computing Galois cohomology groups.

We begin by recalling some basic facts from Galois cohoology, in the context of our problem.

PROPOSITION 9.2.2.1. Let $f : A \to B$ be an isogeny of abelian varieties over a general field k (for the purpose of this section, it will either be a local or global field). Then we have exact sequence

$$0 \to \ker(f) \xrightarrow{\iota} A \xrightarrow{f} B \to 0,$$

which induces the long exact sequence

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \ker(f)(k) & \longrightarrow & A(k) & \longrightarrow & B(k) & & \\ & & & & & & \searrow^{\delta} & & \\ & & H^1(G,\ker(f)) & \longrightarrow & H^1(G,A) & \longrightarrow & H^1(G,B) & \longrightarrow & \cdots \end{array}$$

Here, by $H^1(G,S)$ for S a group scheme over K, we mean $H^1(\mathrm{Gal}(\overline{k}/k), S(\overline{k}))$.

The connecting homomorphism δ is defined as follows: Let $x \in B(k)$ and $y \in A(\overline{k})$ such that $f(y) = x$. Then $\delta(x)$ is the cohomology class associated to the map

$$\begin{aligned} \delta(x) : G &\to \ker(f) \\ \sigma &\mapsto y^\sigma - y. \end{aligned}$$

Furthermore, the long exact sequence gives rise to the fundamental short exact sequence

$$0 \to B(k)/f(A(k)) \xrightarrow{\delta} H^1(G,\ker(f)) \to H^1(G,A)[f] \to 0,$$

where $H^1(G,A)[f] = \ker(H^1(G,A) \to H^1(G,B))$.

Proof. The long exact sequence follows from the definition of Galois cohomology and the explicit description of δ follows from ref result in appendix. □

We will now fix some notation. Throughout this section, we let K be a number field, $G_K = \mathrm{Gal}(\overline{\mathbb{Q}}/K)$ be the absolute galois group, and A/K be an abelian variety.

We would like to show the finiteness of $B(K)/f(A(K))$ via the injection $B(K)/f(A(K)) \hookrightarrow H^1(G,\ker(f))$. It would be nice if we could show $H^1(G_K,\ker(f))$ is finite, but this is in general not true, e.g. in the case K a number field and $f = [m]$, $m \geq 2$. Thus, we need to gain finer control of the image of $B(K)/f(A(K))$

in $H^1(G_K, \ker(f))$. The method we will use to do this is via localization at the places of K. Essentially, we will have that the image satisfies nice ramification properties, which will make it land in a finite subgroup.

Let $v \in \mathscr{P}_K$ be a place, let K_v be the complete local ring with respect to v, and let $G_v := \mathrm{Gal}(\overline{K}_v/K_v)$ be the absolute Galois group of K_v. We have that G_v injects into G_K, so we have the restriction map $H^1(G_K, \cdot) \xrightarrow{\text{res}} H^1(G_v, \cdot)$. This, along with proposition 9.2.2.1 applied to K_v, gives us the following commutative diagram:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & B(K)/f(A(K)) & \xrightarrow{\delta} & H^1(G_K, \ker(f)(\overline{\mathbb{Q}})) & \xrightarrow{\iota} & H^1(G_K, A(\overline{\mathbb{Q}}))[f] & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & B(K_v)/f(A(K_v)) & \xrightarrow{\delta_v} & H^1(G_v, \ker(f)(\overline{K}_v)) & \xrightarrow{\iota_v} & H^1(G_v, A(\overline{K}_v))[f] & \longrightarrow & 0 \end{array}$$

Now, observe that by diagram chasing, if we have some $[a] \in A(K)/f(B(K))$, then $\delta([a]) \in \ker(\iota)$ and thus $\delta_v([a]) \in \ker(\iota_v)$. Thus, each place v gives us some information about the image of δ. This motivates the following definitions:

DEFINITION 9.2.2.2. Let $f : A \to B$ be an isogeny of abelian varieties defined over a number field K. Then the **Selmer group** of A with respect to f is the group

$$\mathrm{Sel}^{(f)}(A/K) := \bigcap_{v \in \mathscr{P}_K} \ker\left(H^1(G_K, \ker(f)) \to H^1(G_v, A)[f]\right)$$

and the **Tate-Shafarevich group** of A is the group

$$Ш(A/k) := \bigcap_{v \in \mathscr{P}_K} \ker\left(H^1(G_K, A) \to H^1(G_v, A)\right)$$

We can now immediately deduce from proposition 9.2.2.1 that the following sequence is exact:

$$0 \to B(K)/f(A(K)) \to \mathrm{Sel}^{(f)}(A/K) \to Ш(A/K)[f] \to 0,$$

where $Ш(A/K)[f] := \ker\left(Ш(A/K) \xrightarrow{f} Ш(B/K)\right)$.

We claim that the Selmer group $\mathrm{Sel}^{(f)}(A/K)$ is finite for any isogeny f, which will then show that both $B(K)/f(A(K))$ and $Ш(A/K)[f]$ are finite. Hence, letting $f = [n]_A$ be the multiplication-by-n map, we will have then proven the weak Mordell-Weil theorem.

Before we begin this proof, we give a geometric interpretation of the Selmer and Tate-Shafarevich groups. As seen in reference result in galois cohomology appendix, an element of $H^1(G_K, A)$ corresponds to an A-torsor and corresponds to the trivial element of H^1 if and only if it admits a K-point. Then, a $\ker(f)$-torsor is contained in $\mathrm{Sel}^{(f)}(A/K)$ if and only if it has K_v-points for all places $v \in \mathscr{P}_K$. Similarly, a non-trivial element in $Ш(A/K)$ corresponds to a torsor which has K_v-points for all $v \in \mathscr{P}_K$, but neverless has no K-points, i.e., the Hasse principle fails. it would be nice to give examples for which the TS group is nontrivial.

maybe add some more exposition about this stuff, including what we currently know about these groups

We will now prove that the Selmer group is finite, which will then conclude our proof of the weak Mordell-Weil theorem (9.2.1.1). We will do this via a "ramification" argument. So, to begin, we define what it means for a cohomology class to be ramified.

DEFINITION 9.2.2.3. Let $v \in \mathscr{P}_K$ be a place of K, $I_v \subset G_K$ the inertia group of v, M a G_K-module, and $\phi \in H^1(G_K, M)$ a cohomology class. We say ϕ is **unramified at** v if its restriction to $H^1(I_v, M)$ is trivial. If it is non-trivial, we say ϕ is **ramified at** v.

Furthermore, if $S \subset \mathscr{P}_K$ is a finite set of places, then we define the notation

$$H^1_S(G_K, M) = \bigcap_{v \in \mathscr{P}_K - S} \ker\left(H^1(G_K, M) \xrightarrow{\text{res}} H^1(I_v, M)\right),$$

i.e., $H^1_S(G_K, M)$ is the subgroup of cohomology classes unramified outside of S.

We will state the key propositions which will conclude our proof of the weak Mordell-Weil theorem. After stating them as our goals for the following section, for each proposition, we will prove a series of results for each.

PROPOSITION 9.2.2.4. Let M be a finite G_K-module, and let S be a finite set of places. Then the group $H^1_S(G_K, M)$ of cohomology classes unramified outside of S is finite.

Proof. To begin, note that by the inflation-restriction exact sequence, for any finite extension L/K, we have that

$$0 \to H^1(G_{L/K}, M^{G_{L/K}}) \to H^1(G_K, M) \to H^1(G_L, M)$$

is exact. Thus, it is sufficient to prove the result for L, i.e. we may enlarge K by a finite extension. In addition, since M is finite and the action of G_K is continuous, we have that an open subgroup of G_K acts trivially. This corresponds to a finite extension L, so by replacing K with L, we may assume G_K acts trivially on M. Now since M is finite, we may assume it has some exponent m. Thus, $H^1_S(G_K, M)$ consists of elements of exponent m unramified outside S. Thus, we have an injection $H^1_S(G_K, M) \hookrightarrow H^1_S(G_{L/K}, M)$ where $L = L_{m,S}$ is the maximal abelian extension of exponent m unramified outside of S (this injection is clear by thinking of elements of H^1 as twisted homomorphisms). Thus, if $L_{m,S}$ is finite for any finite set S and any m, we are done.

Thus, since $L_{m,S}$ is the compositum of all extensions of K of degree dividing m and unramified outside S, this is settled by proposition C.3.0.1, a classical result in algebraic number theory, proven by Hermite. □

PROPOSITION 9.2.2.5. Let $f : A \to B$ be an isogeny of abelian varieties over K. Let S be a set containing the following places:

- all archimedean places of K,
- all places of bad reduction for A and B (which can be shown to be the same set),
- all places dividing $\deg(f)$.

Then $\text{Sel}^{(f)}(A/K) \subset H^1_S(G_K, \ker(f))$. In particular, by proposition 9.2.2.4, $\text{Sel}^{(f)}(A/K)$ is finite.

REMARK 9.2.2.6. Letting $f = [n]_A$ concludes the proof of the weak Mordell-Weil theorem.

Proof. Let $\phi \in \text{Sel}^{(f)}(A/K)$ and let v be a place not in S. Choose a point $y \in A(\overline{K}_v)$ such that $\phi(\sigma) = \sigma(y) - y$ for $\sigma \in G_v$. For $\sigma \in I_v$, we have that the reduction to the residue field k

$$\overline{\sigma(y) - y} = \overline{\sigma(y)} - \overline{y} = 0.$$

But $\sigma(y) - y$ is a torsion point (as it is contained in $\ker(f)$). Additionally, since we assumed v is not in S, we have that v is a finite place of good reduction not dividing $\deg(f)$, theorem 9.2.2.7 tells us that $\sigma(y) - y = 0$. This shows that $\phi(\sigma) = 0$ for all $\sigma \in I_v$, i.e. ϕ is unramified outside S. □

THEOREM 9.2.2.7. Let A be an abelian variety over K, a number field, and let v be a finite place at which A has good reduction, k the residue field of K at v, and p the characteristic of k. Then for any $m \geq 1$ with $p \nmid m$, then the reduction map

$$A[m](K) \to A_k(k)$$

is injective. I.e. the reduction map is injective on the prime-to-p torsion subgroup of $A(K)$.

We will give two different proofs, dependent on what the reader is more comfortable with. The first proof will be more geometric and will use the results of section 8.1. The second will be a nice trick using formal groups, which we explored in chapter 4.

Proof (Geometric). exercise for Rawley □

Proof (Formal Groups). exercise for Aaryan, outline here:

LEMMA 9.2.2.8. Let R be a local ring, $K = \text{Frac}(R)$ and A/K have good reduction at $\mathfrak{m}$. Then the coefficients of the power series of the formal group associated to A lie in the ring R.

PROPOSITION 9.2.2.9. Let R be a complete local ring with maximal ideal $\mathfrak{m}$ and residue characteristic p, and let F be a formal group over R. Then the group $F(\mathfrak{m})$ has no prime-to-p torsion.

PROPOSITION 9.2.2.10. Let R be a complete local ring with maximal ideal $\mathfrak{m}$, fraction field K, and residue field k. Let A/K be an abelian variety with good reduction at $\mathfrak{m}$, and let

$$A_1(K) := \ker(A(K) \xrightarrow{\text{red}} A_k(k))$$

be the kernel of reduction. Let F be the formal group of A described above, which by lemma 9.2.2.8 is defined over R. Then there is an isomorphism

$$F(\mathfrak{m}) \xrightarrow{\sim} A_1(K).$$

□

We note that the conclusion of the weak Mordell-Weil theorem implies the finiteness of $Ш(A/K)[n]$ for all n. In fact, it is conjectured that $Ш(A/K)$ is finite. This is known in a few special cases. research more about this and say a bit more

§9.2.3 Bounds on the Rank of $A(K)$

Now that we have proven the Mordell-Weil theorem, a natural question is to ask whether we can bound the rank of $A(K)$. As mentioned before, the Selmer group gives a method to compute a bound on the n-torsion via counting torsors with K_v points using Hensel's lemma.

We should mention more about this/compute example or smth .

Here, we will give a more elementary bound by instead bounding the size of $H^1_S(G_K, A[n])$ with S as in proposition 9.2.2.5. If we let $L_{n,S}$ be the maximal abelian extension of exponent n unramified outside S as in the proof of proposition 9.2.2.4, then as we noted, cocycles in $H^1_S(G_K, A[n])$ factor through $H^1(G_{L_{n,S}/K}, A[n](\overline{K}))$. Therefore, we can get a rough bound via computing $G_{L_{n,S}/K}$. We of course know the structure $A[n](\overline{K}) \simeq (\mathbb{Z}/n\mathbb{Z})^{2g}$ (see 1.1.4.11).

Finish Later!

this stuff is the only thing which requires Kummer theory, so idk if we need to include an appendix on it

Part IV

Appendices

.

Appendix A

Reviewing “Nice" Properties of Schemes and Morphisms

"*...as Grothendieck taught us, the objects of a category do not play a big role, it is the morphisms that are essential.*"

Translation from *Motifs*
Jean-Pierre Serre

The power of algebraic geometry lies in the large class of nice possible properties that morphisms in the category of schemes have, and their useful consequences. This class is large enough that even the best of mathematicians cannot reasonably remember all the details, conditions and uses of these properties. For this purposes we collect here a laundry list of such nice properties and the useful ways in which they interact that will be used everywhere throughout this text. Not much in this Appendix will be proven in great detail; indeed, many of these results would be (at most) homework exercises in an introductory class on scheme-theoretic algebraic geometry. It is simply recorded here for convenience of reference.

Section A.1

Smoothness

§A.1.1 Flatness

We begin with a slightly non-standard (but equivalent!) definition of flatness, which we will see has better geometric motivation.

Definition A.1.1.1. For A a commutative ring, an A-module M is **flat** if the canonical map

$$I \otimes_A M \to IM$$
$$i \otimes m \mapsto im$$

is an isomorphism for every ideal $I \trianglelefteq A$. Moreover M is **faithfully flat** if $IM \neq M$ for any *proper* ideal I.

Exercise A.1.1.A. Show the above definition is equivalent to the "standard" definition of flatness: that tensoring with M preserves exact sequences in the category of A-modules.

If one adopts the definition in A.1.1.A, one can show flatness implies various nice properties with some algebraic formalism, but the definition itself seems rather disjoint from those properties, and unmotivated. However our definition lends itself much more nicely to geometric intution: if M denotes the module of global sections of some $\mathcal{O}_{\mathrm{Spec}(A)}$-module, then IM is precisely the sections vanishing on the closed subscheme $Z := V(I)$. Flatness of M tells us that these sections on IM have a nice description: they are precisely formal linear combinations of pure tensors $i \otimes m$, with no non-trivial (i.e. not induced by A-module structure) relations.

Key Idea. Geometrically, the sections of a flat module M "are not too discontinuous". This is why, for example, flatness is often an essential stipulation in various moduli problems, as it results in "nice" behavior of fibrations over some base scheme.

Example A.1.1.2. If A a field, then any vector space is flat. If A is a Dedekind domain then flatness if (visibly, by unqiue prime factorization of ideals) equivalent to M containing no torsion.

Exercise A.1.1.B. Show the following properties of flatness, for an A-module M:

(a) M flat if and only if its localization $M_{\mathfrak{p}}$ is flat for all $\mathfrak{p} \in \mathrm{Spec}(A)$.

(b) If A is a local ring and M is finitely generated then M is flat if and only if it is free.

(c) Show *any* localization $S^{-1}A$ and *any* polynomial algebra $A[x_1, \ldots x_n]$ are flat A-modules (given the standard A-module structure induced by inclusion of A).

In general we say a ring map $R \to S$ is (faithfully) flat if S is flat as an R-module (induced by the morphism).

Definition A.1.1.3. A morphism of schemes $f : X \to Y$ is **flat** at $x \in X$ if the induced morphism of stalks $\mathcal{O}_{Y,f(x)} \to \mathcal{O}_{X,x}$ is flat. If it is flat at all $x \in X$ we say f is a flat morphism. It is **faithfully flat** if f is also surjective.

Exercise A.1.1.C. If $X = \mathrm{Spec}(R), Y = \mathrm{Spec}(S)$, show $f : X \to Y$ being (faithfully) flat if and only if the induced $S \to R$ is a (faithfully) flat morphism of rings.

§A.1.2 Differential Forms

In this section we define the sheaf of relative differential forms of one scheme over another, via the purely algebraic method of derivations. In all that follows, let R be a ring and A an R-algebra.

Definition A.1.2.1. An R**-derivation** of A into an A-module M is an R-linear map $d : A \to M$ such that

$$d(fg) = fd(g) + gd(f) \ \forall f, g \in A.$$

We denote by $\mathrm{Der}_R(A, M)$ the set of all R-derivations of A into M.

We note that $d(1) = d(1) + d(1) \implies d(1) = 0$, and R-linearity lets us deduce that $d(r) = 0$ for all $r \in R$. Moreover we can also (canonically!) endow $\mathrm{Der}_R(A, M)$ with the structure of an A-module.

Example A.1.2.2. Consider the differentiation map

$$d : \mathscr{C}^\infty(\mathbb{R}) \to \mathscr{C}^\infty(\mathbb{R})$$
$$f \mapsto f',$$

which we know from single-variable calculus satisfies $d(fg) = fd(g) + gd(f)$, and hence is an $\mathbb{R}$-derivation. For this classical reason the defining identity of a derivation given in Defn. A.1.2.1 is referred to as the **Leibnitz Rule**.

Definition A.1.2.3. The **module of relative differential forms** (of degree 1) of A over R is the data of an A-module $\Omega^1_{A/R}$ along with an R-derivation $\mathrm{d}_{A/R} : A \to \Omega^1_{A/R}$ that is universal in the sense that the canonical map

$$\mathrm{Hom}_A(\Omega^1_{A/R}, M) \to \mathrm{Der}_R(A, M)$$
$$\phi \mapsto \phi \circ \mathrm{d}_{A/R}$$

is bijection for any A-module M. The map $\mathrm{d}_{A/R}$ is called the **exterior differential**.

If such a module exists, its universal property immediately tells us it must be unique up to (canonical!) isomorphism. The question of existence is addressed in the following proposition

Proposition A.1.2.4. For any R-algebra A, the module of relative differential forms of degree 1 exists.

Proof. Let $m : A \otimes_R A \to A$ be the map induced by multiplication in A, and set $I := \ker(m)$. We may naturally give I/I^2 the structure of an $(A \otimes_R A)/I$-module, and, via the (canonical!) isomorphism $(A \otimes_R A)/I \xrightarrow{\sim} A$, the structure of an A-module too. Consider the map

$$\begin{aligned} \mathrm{d} : \ A &\to I/I^2 \\ a &\mapsto 1 \otimes a - a \otimes 1 \bmod (I^2), \end{aligned}$$

which is visibly an R-derivation of A. We claim the pair $(I/I^2, \mathrm{d})$ is precisely our desired module of relative differential forms. To see it satisfies the universal property, take any A-module M; we wish to show precomposing with d yields a bijection $\mathrm{Hom}_A(I/I^2, M) \xrightarrow{\sim} \mathrm{Der}_R(A, M)$.
To show injectivity, it suffices to observe that the A-span of $\{\mathrm{d}(\ a)\}_{a \in A}$ actually generates I/I^2, hence precomposition by d must act injectively on the set of morphisms out of I/I^2. This follows from the identity

$$x \otimes y = xy \otimes 1 - x(y \otimes 1 - 1 \otimes y)$$

For surjectivity, take any R-derivation $D : A \to M$, and consider the map

$$\begin{aligned} \phi_D : A \otimes_R A &\to M \\ a \otimes a' &\mapsto aD(a'), \end{aligned}$$

which is visibly A-linear. Then we have

$$\begin{aligned} \phi_D \circ \mathrm{d}(\ a) &= \phi_D(1 \otimes a - a \otimes 1) \\ &= D(a) - aD(1) = D(a), \end{aligned}$$

so D has a preimage under precomposition with d , which proves surjectivity. □

Remark A.1.2.5. The construction in the above proof is useful for verifying certain functorial properties of the sheaf of relative differentials, but the reader might wonder about a much more "obvious" description in terms of formal linear combinations of the $\mathrm{d}a$ symbols:

$$\Omega^1_{A/R} := \left(\bigoplus_{a \in A} A \, \mathrm{d}a \right) \Big/ \langle \, \mathrm{d}(fg) \ - (f \, \mathrm{d}g \ + g \, \mathrm{d}f \ ), \ f, g \in A \rangle$$

This, by construction, is an A-module satisfying the universal property (thus isomorphic to our above description). In particular, if $A \cong R[x_1, \dots x_n]/(f_1, \dots f_r)$, then

$$\Omega_{A/R} \cong \left(\bigoplus_{i \in [n]} A \, \mathrm{d}x_i \right) \Big/ \langle \, \mathrm{d}f_1 \ , \dots , \ \mathrm{d}f_r \ \rangle .$$

This can be interpreted as the cokernel of the Jacobian map $J : A^{\oplus r} \to A^{\oplus n}$.

The universal property of the module of relative differentials implies various functorial properties

Proposition A.1.2.6. Let R be a ring and take any morphism $\phi : A \to B$ of R-algebras. The module of degree 1 relative differential forms $\Omega^1_{A/R}$ satisfies the following properties

(a) There is an induced A-linear map of Ω^1's given by

$$\Omega^1_{A/R} \to \Omega^1_{B/R}$$
$$\sum f_i \mathrm{d_{A/R}(g_i)} \mapsto \phi(f_i)\, \mathrm{d}_{B/R}\ (\phi(g_i)).$$

(b) There is a canonical exact sequence

$$\Omega^1_{A/R} \otimes_A B \to \Omega^1_{B/R} \to \Omega^1_{B/A} \to 0.$$

(c) If $B = A/\mathfrak{a}$ is a residue algebra of A, the (canonical) B-linear map

$$\delta : \mathfrak{a}/\mathfrak{a}^2 \to \Omega^1_{A/R} \otimes_A B$$
$$\overline{a} \mapsto \mathrm{d}_{A/R}\ (a) \otimes 1$$

induces the exact sequence

$$\mathfrak{a}/\mathfrak{a}^2 \to \Omega^1_{A/R} \otimes_A B \to \Omega^1_{B/R} \to 0.$$

EXERCISE A.1.2.A. Prove the above proposition.

We now want to globalize the construction of the module of relative differential forms in terms of sheaves over schemes. Let S be a base scheme, and take some $X \in \underline{\mathrm{Sch}}\,/S$. The diagonal map

$$\Delta : X \to X \times_S X$$

gives an isomorphism of X onto its image, which is a locally closed subscheme of an open subscheme $W \subseteq X \times_S X$. Let $\mathcal{I}$ be the defining sheaf of ideals of $\Delta(X)$ as a subscheme of W.

DEFINITION A.1.2.7. With notation as above, the **sheaf of relative differential one-forms** (of X over S) is defined as the pullback

$$\Omega^1_{X/S} := \Delta^*(\mathcal{I}/\mathcal{I}^2).$$

Since $\mathcal{I}/\mathcal{I}^2$ has the canonical structure of an $\mathcal{O}_{\Delta_X}$-module (as discussed in the proof of A.1.2.4), the pullback has the structure of an $\mathcal{O}_X$-module, which we see is quasi-coherent (in any case) and finite type if X is locally of finite type over S.
If pr_1, pr_2 denote the projection maps of $X \times_S X$ onto the first and second factor respectively, then the map

$$\mathrm{d}_{X/S} : \mathcal{O}_X \to \Omega^1_{X/S}$$
$$f \mapsto \Delta^*((p_2^* f - p_1^* f) + \mathcal{I}^2)$$

is the **exterior differential**. Locally speaking, if $\mathrm{Spec}(R) \subseteq S$ is an affine open and $U = \mathrm{Spec}(A) \subseteq X$ is an open lying over $\mathrm{Spec}(R)$, then $\Omega^1_{X/S}(U) = \Omega^1_{A/R}$, and the differential map is precisely the one constructed in the proof of A.1.2.4.

The sheaf of relative differential forms also obeys a series of nice functorial properties which we recount here. We first remark that if we have a morphism of S-schemes $f : X \to Y$, then we can "pull back"

differential forms as follows. Given affine opens $U = \mathrm{Spec}(B) \subseteq X, V = \mathrm{Spec}(B) \subseteq Y$ with $f(U) \subseteq V$, we have

$$f^*\Omega^1_{Y/S}(U) = B \otimes_A \Omega^1_{A/\mathcal{O}_S}.$$

We define the pullback map by

$$\begin{aligned} f^*\Omega^1_{Y/S}(U) &\to \Omega^1_{X/S}(U) (= \Omega^1_{B/\mathcal{O}_S}) \\ b \otimes \mathrm{d}a &\mapsto b \cdot \mathrm{d}(f^\#(a)) \ . \end{aligned}$$

Gluing these maps together gives us a morphism of $\mathcal{O}_X$-modules $f^*\Omega^1_{Y/S} \to \Omega^1_{X/S}$.

PROPOSITION A.1.2.8. Take $X, Y \in \underline{\mathrm{Sch}}\,/S$, and let $f : X \to Y$ be a morphism of S-schemes. Then the sheaves of relative differential one-forms satisfy the following properties:

(a) The sequence of $\mathcal{O}_X$-modules

$$f^*\Omega^1_{Y/S} \to \Omega^1_{X/S} \to \Omega^1_{X/Y} \to 0$$

is exact (here maps are induced by the sequence described in part (b) of A.1.2.6).

(b) If $\iota : X \hookrightarrow Y$ is an immersion of S-schemes and $\mathcal{I}$ is the sheaf of ideals defining X as a subscheme of Y, the canonical sequence of $\mathcal{O}_X$-modules

$$\mathcal{I}/\mathcal{I}^2 \to \iota^*\Omega^1_{Y/S} \to \Omega^1_{X/S} \to 0$$

is exact (here the maps are the one induced by part (c) of A.1.2.6).

(c) Let $p : X_Y \to X$ be the projection from the base change of X to Y. Then the canonical map

$$p^*\Omega^1_{X/S} \to \Omega^1_{X_Y/Y}$$

is an isomorphism.

(d) If p_X, p_Y denote the projections from $X \times_S Y$ to their components, the canonical map

$$p_X^*\Omega^1_{X/S} \oplus p_Y^*\Omega^1_{Y/S} \to \Omega^1_{(X\times_S Y)/S}$$

is an isomorphism.

Proof. Parts (a) and (b) follow essentially immediately from A.1.2.6. We leave the proofs of (c), (d) as exercises. [Finish proof] □

EXERCISE A.1.2.B. Use A.1.2.8 and the Jacobian description in A.1.2.5 to show $\Omega^1_{\mathbb{A}^n_S}$ is free for any scheme S.

§A.1.3 Unramified, Smooth and Étale Morphisms

DEFINITION A.1.3.1. A morphism of schemes $f : X \to S$ is called **unramified** at a point $x \in X$ if there exist an open neighborhood $U \ni x$ and an S-immersion

$$\iota : U \hookrightarrow \mathbb{A}^n_S$$

such that the following conditions are satisfied:

(a) f is locally of finite presentation.
(b) The differential forms of type $\mathrm{d}g$ with sections g of the sheaf of ideals $\mathcal{I}$ of $\iota(U)$ generate $\Omega^1_{\mathbb{A}^n_S/S}$ at $j(x)$.

REMARK A.1.3.2. Condition (a) says that an unramified morphism $f : X \to Y$ is locally of finite presentation. If we remove this and only impose (b), we say f is **formally unramified**.
We note that there is some lack of consensus on this definition; some texts replace condition (a) with being locally of finite *type*. In a lot of cases (e.g. when when working in Noetherian schemes) these definitions are equivalent, but it is worth being careful nonetheless.

PROPOSITION A.1.3.3. Let $f : X \to S$ be locally of finite presentation, let x be a point of X, and set $s = f(x)$. Then TFAE:

(a) f is unramified at x.
(b) $\Omega^1_{X/S,x} = 0$.
(c) The diagonal morphism $\Delta : X \to X \times_S X$ is a local isomorphism at x.
(d) The $\kappa(s)$ scheme $X_s := X \times_S \mathrm{Spec}(\kappa(s))$ is unramified over $\kappa(s)$ at x.
(e) The maximal ideal $\mathfrak{m}_x$ of $\mathcal{O}_{X,x}$ is generated by the maximal ideal $\mathfrak{m}_s \trianglelefteq \mathcal{O}_{S,s}$ and $\kappa(x)/\kappa(s)$ is a finite separable extension.

Proof. [finish] □

EXAMPLE A.1.3.4. Let L/K be an extension of number fields (hence finite separable!). Then the induced morphism $f : \mathrm{Spec}(\mathcal{O}_L) \to \mathrm{Spec}(\mathcal{O}_K)$ is unramified at x iff (by part (e) of the above Proposition) $\kappa(x)/\kappa(f(x))$ is finite separable and $\mathfrak{m}_{f(x)}\mathcal{O}_{X,x} = \mathfrak{m}_x \trianglelefteq \mathcal{O}_{X,x}$. But this is precisely the condition that x is an unramified prime over s. Thus the unramified locus of the morphism are *precisely* the primes of $\mathcal{O}_L$ which are unramified over $\mathcal{O}_K$.

EXERCISE A.1.3.A. Show that unramified morphisms are stable under composition, products and base changes.

DEFINITION A.1.3.5. A morphism $f : X \to S$ is called **smooth of relative dimension** r (we also say X is a smooth over S) at a point $x \in X$ if there exists an open neighborhood $U \ni x$ and an S-immersion

$$\iota : U \hookrightarrow \mathbb{A}^n_S$$

satisfying:

(a) Locally at $s := \iota(x)$, the sheaf of ideals defining $\iota(U)$ as a subscheme is generated by $(n-r)$ sections $g_{r+1}, \dots g_n$.

(b) The differentials $\mathrm{d}g_{r+1}\,(y), \dots \mathrm{d}g_n\,(y)$ are linearly independent in $\Omega^1_{\mathbb{A}^n_S/S} \otimes \kappa(y)$.

If f is smooth of relative dimension 0 (at a point x), we say it is **étale**.

EXERCISE A.1.3.B. Verify that smooth and étale morphisms are stable under composition, products and base changes.

LEMMA A.1.3.6. An immersion $f : X \to S$ is étale if and only if f is open.

Proof. Suppose f is an open immersion. Then it is locally an isomorphism, and thus smooth of relative dimension 0 (i.e étale), so this direction is clear.
For this other direction, [Finish] □

PROPOSITION A.1.3.7. Let $f : X \to Y$ is a smooth morphism of schemes. Then

(a) $\Omega^1_{X/Y}$ is locally free, and its rank at $x \in X$ is equal to the relative dimension of f at x.

(b) If X, Y are both S-schemes the canonical sequence of $\mathcal{O}_X$-modules

$$0 \to f^*\Omega^1_{Y/S} \to \Omega^1_{X/S} \to \Omega^1_{X/Y} \to 0$$

is exact and locally split.

There is also an infinitesimal lifting criterion for smoothness, unramifiedness and étaleness.

PROPOSITION A.1.3.8. For $f : X \to S$ locally of finite presentation, the following conditions are equivalent.

(a) f is unramified (resp. smooth, resp. étale).

(b) For any affine S-schemes Y and for all closed subschemes $Y_0 \subseteq Y$ defined by sheaves of ideals $\mathcal{I} \trianglelefteq \mathcal{O}_Y$ with $\mathcal{I}^2 = 0$, the canonical map

$$\mathrm{Hom}_S(X, Y) \to \mathrm{Hom}_S(Y_0, X)$$

is injective (resp. surjective, resp. bijective).

Proposition A.1.3.9. Let X, Z be S-schemes and $\iota : X \hookrightarrow Z$ a closed immersion locally of finite presentation. Let $\mathcal{I}$ be the sheaf of ideals defining X as a subscheme of Z. Take $x \in X, z := \iota(x)$. Suppose Z is S-smooth at z of relative dimension n. Then TFAE:

(a) X is smooth over S of relative dimension r.

(b) The canonical sequence of $\mathcal{O}_X$-modules

$$0 \to \mathcal{I}/\mathcal{I}^2 \to \iota^*\Omega^1_{Z/S} \to \Omega^1_{X/S} \to 0$$

is split-exact at x, and $r = \dim\left(\Omega^1_{X/S} \otimes \kappa(x)\right)$.

(c) If $\mathrm{d}z_1, \dots \mathrm{d}z_n$ is a basis of $(\Omega^1_{Z/S})_z$, and if $g_1, \dots g_N$ are local sections of $\mathcal{O}_Z$ generating $\mathcal{I}_z$, there exists a re-indexing of the $z_1, \dots z_n$ and of the $g_1, \dots g_N$ such that $g_{r+1}, \dots g_n$ generate $\mathcal{I}$ at z and such that $\mathrm{d}z_1, \dots \mathrm{d}z_r, \mathrm{d}g_{r+1}, \dots \mathrm{d}g_n$ generate $(\Omega^1_{Z/S})_z$.

(d) There exist local sections $g_{r+1}, \dots g_n$ of $\mathcal{O}_Z$ generating $\mathcal{I}_z$ such that the differentials $\mathrm{d}g_{r+1}(z), \dots \mathrm{d}g_n(z)$ are linearly independent in $\Omega^1_{Z/S} \otimes \kappa(z)$.

Proposition A.1.3.10. Let X be locally of finite-type over a field k, and take some point $x \in X$. Then TFAE:

(a) X is smooth over k at x.

(b) $(\Omega^1_{X/k})_x$ is generated by $\dim_x X$ elements (and thus free).

(c) There exists some open neighborhood $U \ni x$ and a perfect field extension k'/k such that $U \otimes_k k'$ is regular.

(d) There exists an open neighborhood $U \ni x$ such that $U \otimes_k k'$ is regular for all field extensions k'/k.

Section A.2

Henselian Schemes and Rings

The notion of étale morphism is akin to the scenario in differential topology of being an isomorphism on tangent spaces. In diff-top land, this is equivalent to being a local isomorphism, but this need not be the case for an étale map. To this end we make the following definition.

Definition A.2.0.1. An affine local scheme $S = \mathrm{Spec}(R)$ is **Henselian** if each étale map $X \to S$ is a local isomorphism at all points x over s with trivial residue field extension $\kappa(x)/\kappa(s)$. If $\kappa(s)$ is separably closed for all s we say S is **strictly Henselian**.

Proposition A.2.0.2. Let R be a local henselian ring with residue field k, and X a smooth R-scheme. Then the canonical map $X(R) \to X(k)$ is surjective. In particular if R is strictly henselian then the set of k-valued points of X_k that lift to R-valued points of X is dense in X_k.

Section A.3

Blowups

The content of this section is primarily due to Brian Conrad's exposition in his 2025 Fall Algebraic Geometry Topics Course, which itself was an adaption and expansion of some notes of Deligne.

§A.3.1 Basic construction

Blowing up a variety along a closed subvariety, or a scheme along a closed subscheme, is one of the more important geometric operations in algebraic geometry. It arose historically in the study of "resolution of singularities", but also occurs in quite different settings. There are at least 3 initial descriptions of blow-ups, most of which look uncomputable at first sight.

(1) (Classical definition): Take the schematic closure of $X - Z$ in an auxiliary variety. This method is typically only discussed in the classical setting of varieties (see [**H**]). It is a bit mysterious at first, and doesn't initially look amenable to computations.

(2) (Relative Proj construction): This is also unintuitive, and looks even more uncomputable, but it is technically useful. We will get to it in §A.3.2. This is the approach [Har77, Ch. II, Section 7] and in the abstruse [EGA, II, §8].

(3) (Universal property): This approach is due to Hironaka (generalizing the classical blowing up at a point), and is rather computable if X is not too "bad" along Z. The blowup is the final object in a specific category of "nice" objects over X. We adopt this final approach as our starting point.

The setup is as follows: let $Z \hookrightarrow X$ be a closed immersion of schemes. Then consider those morphisms $f : Y \to X$ for which the base change closed immersion $f^{-1}Z \hookrightarrow Y$ is Cartier (where $f^{-1}Z := Y \times_X Z$). This says the quasi-coherent ideal sheaf $\mathcal{I}_Z\mathcal{O}_Y$ is invertible. (A boring example of this type is to take Y to be empty.)

We seek a *final object* $Y \xrightarrow{f} X$ among such X-schemes. It will be covered by affine open subschemes reminiscent of the open covering of $\mathbf{P}^n_X$ covered by the affine spaces $D_+(t_j)'s$ which each refine the mapping property of $\mathbf{P}^n_X$. However, similarly to normalizations, the mapping property for the construction we seek will only be as a final object among certain sufficiently nice X-schemes, not with a description of its functor of points on general X-schemes (so this is less robust than the mapping property for projective spaces).

> Definition A.3.1.1 (Blowups via Universal Property). The **blow-up** of X along Z is a morphism $\mathrm{Bl}_Z(X) \xrightarrow{\pi} X$ which is final among maps $f : Y \to X$ with $f^{-1}Z \to \mathrm{Bl}_Z X$ Cartier. The Cartier pullback $E = \pi^{-1}Z \hookrightarrow \mathrm{Bl}_Z(X)$ is called the **exceptional divisor**.

How do we build this? Working locally, we focus on the affine case: $X = \operatorname{Spec} A$, $Z = \operatorname{Spec}(A/I)$. Consider any $Y \to X = \operatorname{Spec}(A)$ such that the quasi-coherent ideal sheaf $\mathcal{I} = I\mathcal{O}_Y \subset \mathcal{O}_Y$ is invertible. Zariski-locally on Y, we can find an $\mathcal{O}_Y$-basis of $\mathcal{I}$ given by elements $a \in I$ (i.e. on members of an open cover, we can trivialize $\mathcal{I}$ with a basis coming from I itself). Indeed, it is enough to check the y-stalks $\mathcal{I}_Y$

as an $\mathcal{O}_y$-module for all $y \in Y$, since a basis of a stalk of a locally free sheaf of finite rank spreads out to a basis over an open neighborhood. Now use the following lemma:

> LEMMA A.3.1.2. Let M be a free module of rank $d > 0$ over a local ring $(R, \mathfrak{m})$, and $\Sigma \subset M$ a (possibly infinite) generating set. Then there exists a basis of M in Σ.

Proof. Let $k = R/\mathfrak{m}$, so the map $\Sigma \to M/\mathfrak{m}M \simeq k^{\oplus d}$ has image that spans. Then by linear algebra, there is subset $\sigma_1, \dots \sigma_d \in \Sigma$ such that $\{\sigma_i\} \subset M$ reduces to a k-basis. This is then an R-basis of the free module M by Nakayama's lemma. □

In particular, Y has the open cover $\{Y_a\}_{a \in I}$ where

$$Y_a = \{y \in Y : a \text{ is an } \mathcal{O}_{Y,y}\text{-basis of } I\mathcal{O}_{Y,y}\}.$$

We can do better when $I = (a_1, \dots, a_m)$ is finitely generated (as is usually the case): then the finite collection of opens $\{Y_{a_1}, \dots, Y_{a_m}\}$ covers Y (this should be regarded as an analogue of $\{D_+(t_j)\}$ covering $\mathbf{P}^n_A$, for reasons which will be clearer towards the end of this section).

Let us now choose $a \in I$ and impose on $Y \xrightarrow{f} X$ the refined condition that:

$$I\mathcal{O}_Y \text{ is an } \mathcal{O}_Y\text{-free with basis } \{a\}. \tag{*a}$$

(this is analogous in the case of $\mathbf{P}^n$ to the universal property of a single $D_+(t_i)$: it represents the condition on surjections $\mathcal{O}^{n+1} \twoheadrightarrow \mathscr{L}$ that the ith standard basis section of $\mathcal{O}^{n+1}$ for $0 \le i \le n$ is carried to a global basis of $\mathscr{L}$). We seek to find the final such Y subject to this refined condition. It is crucial that we are not merely imposing principality with generator a but really *freeness* with generator a.

> PROPOSITION A.3.1.3. Consider the A-subsalgebra $A[Ia^{-1}] \subset A_a$ generated by elements a'/a for $a' \in I$. Then $\operatorname{Spec}(A[Ia^{-1}]) \to \operatorname{Spec} A$ is final for (*a).

Proof. Setting $A' = A[Ia^{-1}] \subset A_a$, we see that in A' the element a is *not* a zero-divisor since A' is a subring of A_a in which a is a unit, and for any $a' \in I$ we can write $a' = (a'/a)a$ with $a'/a \in A'$, so IA' is A'-free on $\{a\}$. So $\operatorname{Spec} A'$ at least belongs to the desired category of A-schemes. We seek to show it is the final object in this category.

That is, for any $f : Y \to \operatorname{Spec} A$ for which $I\mathcal{O}_Y$ is free with basis a, we seek to uniquely fill in a commutative diagram

$$\begin{array}{ccc} Y & \xrightarrow{\exists!?} & \operatorname{Spec} A' \\ & \searrow^{f} & \downarrow \\ & & \operatorname{Spec} A, \end{array}$$

Thanks to the uniqueness aspect, for this purpose we we can work *locally* on Y. Hence, it suffices to treat the affine case $Y = \operatorname{Spec} R$ (any two affine opens in the original Y have overlap covered by affine opens!). Note that by the hypothesis $IR = Ra$ is *free* with basis a, the element $a \in R$ is not a zero-divisor (even though a may well be a zero-divisor in A) This will be crucial.

Since a is not a zero divisor in R, the natural map $R \to R_a$ is *injective*. Then by considering the composite map $A \to R \hookrightarrow R_a$, we get an associated unique map of A-algebra $h : A_a \to R_a$. Let's assemble

this information in a commutative diagram:

$$\begin{array}{ccccc} A' = A[Ia^{-1}] & \overset{?}{\dashrightarrow} & R & \overset{!!}{\hookrightarrow} & R_a \\ & \nwarrow & \uparrow & & \uparrow{\scriptstyle \exists! h} \\ & & A & \longrightarrow & A_a \end{array}$$

The key point is that the A-subalgebra $h(A') \subset R_a$ lies in R precisely because $IR = Ra$!

Thus, we get an A-algebra map $A' \xrightarrow{\varphi} R$, and to check uniqueness consider the commutative diagram of A-algebras

$$\begin{array}{ccc} A' & \xrightarrow{\varphi} & R \\ \downarrow & & \downarrow \\ A'_a = A_a & \longrightarrow & R_a \end{array}$$

where the bottom map is φ_a. But we also noted that $h : A_a \to R_a$ is the unique A-algebra map, so the bottom horizontal map φ_a is really h, no matter what φ may exist. The vertical maps are injections, so we have the desired uniqueness for φ. □

Definition A.3.1.4. For $a \in I$, we define $D_+(a) = \mathrm{Spec}(A[Ia^{-1}])$.

This notation is meant to be suggestive of $D_+(f) \subset \operatorname{Proj} S$, and below it will relate nicely to the standard affine charts of $\mathbf{P}^N_A$ when I is finitely generated. For $a, a' \in I$, we have the diagram of A-schemes

$$D_+(a) \subset D_+(a)_{a'/a} \underset{\simeq}{\overset{\alpha}{\to}} D_+(a')_{a/a'} \supset D_+(a').$$

where α over A between the open subschemes $D_+(a)_{a'/a}$ and $D_+(a')_{a/a'}$ arises from each being final for the same condition on an A-schemes Y: that $I\mathcal{O}_Y$ has each of a and a' as a *basis* (check this shared finality property!). This isomorphism α just says concretely that inside the localization $A_{aa'}$, we have an actual equality of A-subalgebras

$$A[Ia^{-1}]_{a'/a} = A[Ia'^{-1}]_{a/a'}.$$

Checking the triple overlap condition is easy since it expresses an equality among three A-subalgebras of $A_{aa'a''}$, so we can glue the $D_+(a)$'s along such α's.

Proposition A.3.1.5. This gluing is a final object for the Cartier condition on $I\mathcal{O}_Y$ which defines $\mathrm{Bl}_Z(X) = \mathrm{Bl}_I(A)$.

Proof. The argument is very similar to how the universal property of the gluing $\mathbf{P}^n_A$ is established using the meaning of each $D_+(t_j) \subset \mathbf{P}^n_A$ and their overlaps (check!). □

§A.3.2 Computing blow-ups

If $I = (a_0, \ldots, a_m)$ then $\mathrm{Bl}_I(A)$ is covered by the affine open subschemes $D_+(a_j) = \operatorname{Spec} A[Ia_j^{-1}]$. Soon we will see (Example A.3.2.2) that there is a closed immersion

$$i : \mathrm{Bl}_I(A) \hookrightarrow \mathbf{P}^m_A$$

over A such that $i^{-1}(D_+(t_j)) = D_+(a_j)$.

In general there are then two ways of trying to "compute" $D_+(a)$.

Method 1: If $I = (a_1, \ldots, a_m)$, we aim to describe the $D_+(a_j)$ in terms of a ring

$$A_j = A[T_{ij}]_{i \neq j}/(a_i - T_{ij}a_j)_{i \neq j},$$

the initial A-algebra where each a_i for $i \neq j$ is a *specified* multiple of a_j. We want a *unique* multipler. The issue is that the a_j may be a zero divisor in this quotient ring (even if A is a domain). To fix this, we kill all a_j-power torsion: $A_j/(a_j^\infty\text{-torsion})$. (Beware it is not enough to kill a_j-torsion, since there could be a_j^2-torsion which gives rise to nonzero a_j-torsion in the quotient by the a_j-torsion. For instance, the quotient of the ring $R = \mathbf{Z}/(27)$ by its 3-torsion ideal $9R$ is $\mathbf{Z}/(9)$ in which there is nonzero 3-torsion.)

This quotient is precisely $A[Ia_j^{-1}]$, since they satisfy the same initial property among A-algebras (initial on the ring side expresses finality on the scheme side). The description as such a quotient of A_j is related to the classical description of the blowup, which we'll discuss tomorrow. How does one compute this a_j^∞-torsion concretely? It can be pretty difficult in general, and the difficulty of that computation is one way of expressing the difficulty of computing the schematic closure in the classical description of blow-ups (to be seen a bit later below).

Method 2: We claim that $\left(\bigoplus_{n \geq 0} I^n\right)_{a_1} = A[Ia^{-1}]$, where $a_1 := a \in I$ is viewed in degree 1 (hence the subscript $_1$).

Proof. We have $\left(\bigoplus_{n \geq 0} I^n\right)_{a_1} \subset \left(\bigoplus_{n \geq 0} A\right)_{a_1} = A[T]_{(aT)}$, and this is the A-subalgebra generated by degree-0 fractions b/a^r for $b \in I^r$. But any such b is a finite A-linear combination of products of r elements of I, so this A-subalgebra of A_a is generated by the fractions a'/a for $a' \in I$. This in turn is exactly how we defined $A[Ia^{-1}]$. □

In the A-scheme $B = \mathrm{Proj}\left(\bigoplus_{n \geq 0} I^n\right)$, the preceding description provides open subschemes $D_+(a) = \mathrm{Spec}(A[Ia^{-1}])$ for $a \in I$ (viewed in degree 1 for the $\mathbb{N}$-graded A-algebra $\bigoplus_{n \geq 0} I^n$), and these are certainly a cover (since $\bigoplus_{n \geq 0} I^n$ is generated by its degree-1 part I over its degree-0 part A). Moreover, their overlaps are the *same* as in the gluing $\mathrm{Bl}_I A$. This shows:

THEOREM A.3.2.1. $\mathrm{Bl}_I(A) = \mathrm{Proj}(\bigoplus_{n \geq 0} I^n)$.

This Proj description is the *definition* of the blow-up in [**H**] (and in other references). It is efficient, but quite unmotivated as an initial definition. The approach we have taken hopefully provides a more concrete understand of why this Proj construction has the desired finality property for the blow-up. Next time we will globalize beyond the case of affine X and discuss some properties and work out some calculations with the "affine chart" perspective developed above.

Here is a nice application of the Proj description:

EXAMPLE A.3.2.2. In the case $I = (a_1, \ldots, a_m)$, we have a closed immersion $\mathrm{Bl}_I A \hookrightarrow \mathbf{P}_A^{m-1}$ over A via the surjection of $\mathbf{N}$-graded A-algebras $A[T_1, \ldots, T_m] \twoheadrightarrow \bigoplus_{n \geq 0} I^n$ given by $T_j \mapsto (a_j)_1$.

By definition $\pi : \mathrm{Bl}_Z(X) \to X$ is a relative Proj, namely $\mathrm{Proj}(\mathscr{S})$ for the $\mathbf{N}$-graded quasi-coherent $\mathcal{O}_X$-algebra $\mathscr{S} = \bigoplus_{n \geq 0} \mathcal{I}^n$ for the quasi-coherent ideal $\mathcal{I} \subset \mathcal{O}_X$ of $Z \subset X$. The compatibility of relative

Proj with any base change thereby gives an explicit description of the exceptional divisor $E = \pi^{-1}(Z)$:

$$E = \mathrm{Proj}((\mathcal{O}_X/\mathcal{I}) \otimes_{\mathcal{O}_X} (\bigoplus_{n\geq 0} \mathcal{I}^n)) = \mathrm{Proj}(\bigoplus_{n\geq 0} \mathcal{I}^n/\mathcal{I}^{n+1}).$$

In the special case that X is smooth over a field k and Z is also smooth over k, this takes on a more concrete form, as we now explain.

LEMMA A.3.2.3. If X is smooth over a field k and Z is smooth over k as well then the natural map

$$S(\mathcal{I}/\mathcal{I}^2) \to \bigoplus_{n\geq 0} \mathcal{I}^n/\mathcal{I}^{n+1}$$

between graded $\mathcal{O}_Z$-algebras arising from the universal property of symmetric algebras and the identity map on $\mathcal{I}/\mathcal{I}^2$ in degree 1 is an isomorphism. Moreover, $\mathcal{I}/\mathcal{I}^2$ is a vector bundle over $\mathcal{O}_Z$.

Proof. We may and do extend scalars to arrange that k is algebraically closed (so all closed points are k-rational). The problem is Zariski-local on X, so we can assume $X = \mathrm{Spec}(A)$ is affine. Let $I \subset A$ be the ideal of Z, so by hypothesis both A and A/I are regular. Our aim is to show that the natural map

$$S(I/I^2) \to \bigoplus_{n\geq 0} I^n/I^{n+1}$$

of A/I-algebras is an isomorphism, which is to say that $f_n : S^n(I/I^2) \to I^n/I^{n+1}$ is an isomorphism for all $n \geq 2$ (the cases $n = 0, 1$ are trivial), and that I/I^2 is locally free as an A/I-module (or equivalently by the noetherian property: it has free stalks at maximal ideals). The map is easily seen to be surjective, but injectivity is rather less clear.

We are now going to carry out an algebraic analogue of the method of passing to "local coordinates" in differential geometry to reduce to a concrete situation where one can compute things. This rests on the technique of completion for noetherian rings along ideals, which is developed in [Mat, §8], especially Theorem 8.8 there.

It suffices to check after localizing at each maximal ideal $\mathfrak{m}$ of A/I. For both the freeness (equivalently flatness) of $(I/I^2)_{\mathfrak{m}})$ over $(A/I)_{\mathfrak{m}}$ and the isomorphism property of $(f_n)_{\mathfrak{m}}$ between finitely generated $(A/I)_{\mathfrak{m}}$-modules, it suffices to check tafter faithfully flat base change to the completion of this local noetherian ring.

That completion is R/J where $R = A^{\wedge}_{\mathfrak{m}}$ is regular and $J = IR$ is an ideal of R for which R/J is also regular. The induced map of R-modules in degree n is $S^n(J/J^2) \to J^n/J^{n+1}$. Since R is a *complete* regular local noetherian k-algebra with residue field k, we have $R \simeq k[\![t_1, \ldots, t_d]\!]$ is a formal power series ring over k. Indeed, by regularity the maximal ideal has a set of d generators where $d = \dim R$, so that defines a local k-algebra map $k[t_1, \ldots, t_d]_{(t_1,\ldots,t_d)} \to R$, and by completeness of the target this factors through the completion as a local k-algebra map $k[\![t_1, \ldots, t_d]\!] \to R$. By completeness and separatedness of the target and completeness of the source, this is surjective due to a successive approximation argument (check, building up from the isomorphism property modulo squares of the maximal ideals). But both source and target are domains with the same dimension d, so the prime kernel ideal must be (0).)

The regular quotient R/J is also a formal power series ring, and we can arrange the choice of t_j's (which lifted any basis of $\mathfrak{m}/\mathfrak{m}^2$) so that the first m of them map to formal variables for a power series description of R/J. By surjectivity we can then easily change the other t_j's so they are all killed by the map, which

is to say $J = (t_{m+1}, \ldots, t_d)$. But now the map of interest $S^n(J/J^2) \to J^n/J^{n+1}$ over $R/J = k[\![t_1, \ldots, t_m]\!]$ is obtained by passing to completion at the origin on the analogous map for the usual polynomial ring $k[t_1, \ldots, t_d]$ and its ideal $(t_{m+1}, \ldots, t_d)$. In other words, by completion of a situation with abstract smooth k-schemes we have reduced our task to the special case of affine space and a linear subspace defined by the vanishing of some standard coordinates.

To summarize, we have reduced the general problem to the special case $X = \mathbf{A}^d_k$ (i.e., $A = k[t_1, \ldots, t_d]$) and $Z = \{t_{m+1} = \cdots = t_d = 0\} = \mathbf{A}^m_k$ (i.e., $I = (t_{m+1}, \ldots, t_d)$). In this case by inspection I/I^2 as a module over $A/I \simeq k[t_1, \ldots, t_m]$ is *free* of rank $d - m$, with basis $t_{m+1}, \ldots, t_d$. Thus, $S^n(I/I^2)$ is also a free A/I-module, with basis given by monomials of degree n in $t_{m+1}, \ldots, t_d$. But it is clear by inspection that I^n/I^{n+1} is not only spanned over A/I by those monomials but is even *freely* spanned, since the subring $k[t_1, \ldots, t_m]$ of A maps isomorphically onto A/I. □

By definition $\mathbf{P}(\mathscr{E}) := \mathrm{Proj}(S(\mathscr{E}))$, so we immediately obtain:

PROPOSITION A.3.2.4. If X is smooth over a field k with dimension d and Z is also smooth over k with dimension $m < d$ then the exceptional divisor of $\mathrm{Bl}_Z(X) \to X$ is a $\mathbf{P}^{d-m-1}$-bundle over Z given by $\mathbf{P}(\mathcal{I}/\mathcal{I}^2)$, with $\mathcal{I}/\mathcal{I}^2$ a vector bundle over Z with rank $d - m$.

EXAMPLE A.3.2.5. Let's work out explicitly the most classical case of all: the blow-up of $\mathbf{A}^n_k$ at the origin. In effect, we will show that our definition of blow-up recovers the ad-hoc definition given near the end of [**H**]. Since the ideal $I \subset k[t_1, \ldots, t_n]$ for the origin is $(t_1, \ldots, t_n)$, which has n generators, we have seen via Example A.3.2.2 that $\mathrm{Bl}_0(\mathbf{A}^n_k)$ is a closed subscheme of $\mathbf{P}^{n-1}_{k[t_1,\ldots,t_n]} = \mathbf{A}^n_k \times_k \mathbf{P}^{n-1}_k$. In [**H**] the blow-up at the origin is *defined* by the system of homogenous linear equations over $k[t_1, \ldots, t_n]$ given by $t_i x_j = t_j x_i$ for $i \neq j$ where $x_1, \ldots, x_n$ are the homogeneous coordinates on $\mathbf{P}^{n-1}_k$. So we will show that closed subscheme is exactly $\mathrm{Bl}_0(\mathbf{A}^n_k)$ as such a closed subscheme.

By design, inside the chart

$$D_+(x_i) = \mathbf{A}^{n-1}_{k[t_1,\ldots,t_n]} = \mathrm{Spec}(k[t_1, \ldots, t_n][x_1/x_i, \ldots, x_n/x_i])$$

of projective $(n-1)$-space, the closed subscheme for the blow-up is given by first imposing the relations $t_j = (x_j/x_i)t_i$ for $j \neq i$ (as x_j/x_i for $j \neq i$ are the standard coordinates on the affine $(n-1)$-space $D_+(x_i)$ over the base ring $k[t_1, \ldots, t_n]$) and then killing all t_i^∞-torsion. The relations $t_j = (x_j/x_i)t_i$ for $j \neq i$ are exactly the x_i-dehomogenization of the relations $t_i x_j = t_j x_i$, so it remains to check that upon imposing these dehomogenized relations the element t_i is not a zero-divisor (so there is no nonzero t_i^∞-torsion to kill).

In other words, we want to show that t_i isn't a zero-divisor in the ring

$$k[t_1, \ldots, t_n, x_1/x_i, \ldots, x_n/x_i]/(t_j - (x_j/x_i)t_i)_{j \neq i}.$$

As a quotient of the polynomial ring, the imposed relations make the t_j's disappear for $j \neq i$, leaving us with $k[t_i, x_1/x_i, \ldots, x_n/x_i]$ that is visibly a domain in which t_i is nonzero, so indeed it isn't a zero-divisor in this ring.

§A.3.3 Strict transform

Let X be a scheme, $Z \subset X$ a closed subscheme with associated quasi-coherent ideal sheaf $\mathcal{I}$, $\pi : \widetilde{X} = \mathrm{Bl}_Z(X) \to X$ the corresponding blow-up, and $Z' := \pi^{-1}(Z) = Z \times_X \widetilde{X} \subset \widetilde{X}$ the exceptional divisor (i.e., the closed subscheme corresponding to the invertible ideal $\mathcal{I} \cdot \mathcal{O}_{\widetilde{X}}$ on $\widetilde{X}$ via the universal property of blow-up). The restriction $\pi : \widetilde{X} - Z' \to U := X - Z$ is an isomorphism (since $\mathcal{I}|_U = \mathcal{O}_U$ is trivially invertible). In this way we view U as an open subscheme of $\widetilde{X}$ whose complement is defined by an invertible ideal. The inclusion $j : U \hookrightarrow \widetilde{X}$ is quasi-compact due to:

> LEMMA A.3.3.1. Let $j : V \hookrightarrow Y$ be an open immersion of schemes whose complement is the zero locus of a quasi-coherent ideal of finte type (as is automatic when Y is locally noetherian). Then j is quasi-compact.

Proof. The task is local on Y, so we can assume $Y = \mathrm{Spec}(A)$ is affine. Our task then reduces to showing that V is quasi-compact. The complement of V in Y is defined by a quasi-coherent ideal of finite type, and such an ideal has the form $\widetilde{I}$ for an ideal $I \subset A$ that is finitely generated. If $\{a_1, \ldots, a_m\}$ is a finite generating set for I then the open complement V of $\mathrm{Spec}(A/I)$ in $\mathrm{Spec}(A) = Y$ is covered by the open subsets $\mathrm{Spec}(A_{a_j})$ for $1 \leq j \leq m$. Hence, V is covered by finitely many affine open subschemes, so V is quasi-compact. □

Since $j : U \hookrightarrow \widetilde{X}$ is quasi-compact, it has a meaningful scheme-theoretic image whose underlying space is the topological closure of U in $\widetilde{X}$ and whose corresponding quasi-coherent ideal sheaf is $\ker(\mathcal{O}_{\widetilde{X}} \to j_*(\mathcal{O}_U))$. But this kernel vanishes, since over open affines on $\widetilde{X}$ where the invertible ideal of the exceptional divisor is free of rank 1 the map j is $\mathrm{Spec}(R_r) \to \mathrm{Spec}(R)$ for $r \in R$ that isn't a zero-divisor, and then $\ker(R \to R_r) = 0$.

> DEFINITION A.3.3.2. For a closed subscheme Y in X the **strict transform** $\widetilde{Y}$ of Y in X (relative to Z) is defined to be $\mathrm{Bl}_{\mathcal{I} \cdot \mathcal{O}_Y}(Y) = \mathrm{Bl}_{(\mathcal{I}+\mathcal{I}_Y)/\mathcal{I}_Y}(Y) = \mathrm{Bl}_{Z \cap Y}(Y)$.

In the (uninteresting) case $|Y| \subset |Z|$ as underlying topological spaces inside $|X|$, the ideal $\mathcal{I} \cdot \mathcal{O}_Y$ consists of locally nilpotent sections, so the blow-up along this is empty (as the only way an ideal of nilpotent elements can be invertible is if one is in the zero ring). Our definition of "strict transform" is *not* the usual definition of such terminology; in Lemma A.3.3.3 we will prove the equivalence with another construction that is the usual definition.

The first key point is that the canonical map $\widetilde{Y} \to \widetilde{X}$ over $Y \hookrightarrow X$ (from functoriality of $\mathrm{Bl}_Z(X)$ in the pair (X, Z) as described more broadly in the diagram (A.3.4-I) early in Section A.3.4) is a closed immersion. Indeed, it suffices to check this over an open affine covering of X, and if $X = \mathrm{Spec}(A)$ is affine then $Z = \mathrm{Spec}(A/I)$ and $Y = \mathrm{Spec}(A/J)$ for ideals I and J in A, so $\widetilde{Y} = \mathrm{Bl}_{I(A/J)}(\mathrm{Spec}(A/J)) = \mathrm{Bl}_{(I+J)/J}(\mathrm{Spec}(A/J)) = \mathrm{Bl}_{Y \cap Z}(Y)$. The natural map $\widetilde{Y} \to \widetilde{X}$ is (check!) induced by Proj "functoriality" applied to the map of graded A-algebras

$$\bigoplus_{n \geq 0} I^n \to \bigoplus_{n \geq 0} ((I + J)/J)^n = \bigoplus_{n \geq 0} (I^n + J)/J$$

that is surjective by inspection. This establishes that $\widetilde{Y} \to \widetilde{X}$ is a closed immersion.

We have shown that the strict transform of Y in X (relative to Z) is a closed subscheme of $\mathrm{Bl}_Z(X)$. Here is a description of this closed subscheme in a more direct manner (which is the usual definition of "strict transform").

LEMMA A.3.3.3. Let X be a scheme, Y and Z closed subschemes of X, and $U = X - Z$ the open complement of Z in X. We view $Y \cap U$ as a closed subscheme of U, which in turn is an open subscheme of $\widetilde{X} = \mathrm{Bl}_Z(X)$. The immersion $j : Y \cap U \to \widetilde{X}$ is quasi-compact (even affine) and its scheme-theoretic closure in $\widetilde{X}$ is the closed subscheme $\widetilde{Y} := \mathrm{Bl}_{Z\cap Y}(Y)$.

In the special case that $Y \cap U$ is empty, which is to say $|Y| \subset |Z|$, this recovers the earlier observation that $\mathrm{Bl}_{Y\cap Z}(Y)$ is empty when $|Y| \subset |Z|$.

Proof. This can be checked locally on X, so we may assume $X = \mathrm{Spec}(A)$, $Z = \mathrm{Spec}(A/I)$, $Y = \mathrm{Spec}(A/J)$, so $\widetilde{X} = \mathrm{Bl}_Z(X)$. The scheme $Y - (Z \cap Y) = Y \cap U$ is is covered by open affines of the form $\mathrm{Spec}((A/J)_a)$, with $a \in I$. For $a \in I$, we let $Y_a = \mathrm{Spec}((A/J)_a)$ and $X_a = \mathrm{Spec}(A[Ia^{-1}])$. We have a canonical map

$$Y_a \to X_a \hookrightarrow \widetilde{X}$$

(since $Y_a \subset U \subset \widetilde{X}$ and the ideal I contains a which becomes a unit on Y_a, with X_a representing the condition that a freely generates $I\mathcal{O}_{\widetilde{X}}$).

Let's show $j^{-1}(X_a) = Y_a$ (so j is affine). We compute that for any $a' \in I$, $j^{-1}(X_a) \cap Y_{a'} \simeq Y_{a'} \times_{X_{a'}} (X_{a'} \cap X_a)$ is exactly

$$\mathrm{Spec}((A/J)_{a'} \otimes_{A[Ia'^{-1}]} A[Ia'^{-1}]_{a/a'}) = \mathrm{Spec}(((A/J)_{a'})_{a/a'}) = \mathrm{Spec}((A/J)_{aa'}).$$

This is exactly the open subscheme $Y_{a'} \cap Y_a$ inside of $Y \cap U$. Since $Y \cap U$ is covered by the open subschemes $Y_{a'}$ for $a' \in I$, taking the union over all $a' \in I$ yields $j^{-1}(X_a) = Y_a$, as desired.

The kernel of $\mathcal{O}_{\widetilde{X}} \to j_*(\mathcal{O}_{Y\cap U})$ over X_a is exactly the quasi-coherent sheaf associated to the kernel of $A[Ia^{-1}] \to (A/J)_a = A_a/J_a$. Since this kernel is just $J_a \cap A[Ia^{-1}]$ (intersection inside A_a), we just need to show that it agrees with the kernel of $A[Ia^{-1}] \to (A/J)[((I + J)/J)a^{-1}]$. The target ring is a subring of $(A/J)_a = A_a/J_a$, so we're done. □

We may also speak of the strict transform more generally, of both quasi-coherent modules and S-schemes.

DEFINITION A.3.3.4. Let Z be a closed subscheme of X, and let $E \subseteq \mathrm{Bl}_Z(X)$ be the exceptional divisor. Let $f : Y \to X$ a morphism of schemes. Then the **strict transform** of Y (with respect to f and $\mathrm{Bl}_Z(X)$) is the closed subscheme

$$Y' \subseteq Bl_Z(X) \times_X Y$$

cut out by the ideal of sections supported on $\mathrm{pr}^{-1}_{\mathrm{Bl}_Z(X)} E$.

EXERCISE A.3.3.A. Show that Y' is the blowup of Y along $f^{-1}(Z)$.

§A.3.4 Functoriality and base change for blow-ups

For a map of schemes $f : X' \to X$ and the closed subscheme $Z' := f^{-1}(Z) = Z \times_X X' \hookrightarrow X'$, consider the blow-ups

$$\pi' : \widetilde{X'} = \mathrm{Bl}_{Z'}(X') \to X', \quad \pi : \widetilde{X} = \mathrm{Bl}_Z(X) \to X.$$

By design, the closed subschemes $E' := \pi'^{-1}(Z') \subset \widetilde{X'}$ and $E := \pi^{-1}(Z) \subset \widetilde{X}$ are Cartier (the "exceptional divisor" for each blow-up). In particular, since $Z' = f^{-1}(Z)$, so $\pi'^{-1}(Z') = \pi'^{-1}(f^{-1}(Z)) = (f \circ \pi')^{-1}(Z)$, we see that under the composite map

$$f \circ \pi' : \widetilde{X'} \to X' \to X$$

the pullback of $Z \hookrightarrow X$ is Cartier. Hence, by the *universality property* (really the finality) of $\pi : \widetilde{X} \to X$ there is a unique map $\widetilde{f} : \widetilde{X'} \to \widetilde{X}$ fitting into a commutative diagram

$$\begin{array}{ccc} \widetilde{X'} & \xrightarrow{\widetilde{f}} & \widetilde{X} \\ {\scriptstyle \pi'}\downarrow & & \downarrow{\scriptstyle \pi} \\ X' & \xrightarrow[f]{} & X \end{array} \tag{A.3.4-I}$$

and by commutativity

$$\widetilde{f}^{-1}(E) = \widetilde{f}^{-1}(\pi^{-1}(Z)) = (\pi \circ \widetilde{f})^{-1}(Z) = (f \circ \pi')^{-1}(Z) = \pi'^{-1}(f^{-1}(Z)) = \pi'^{-1}(Z') = E'.$$

In particular, the commutativity defines an X'-morphism

$$h : \widetilde{X'} \to X' \times_X \widetilde{X}$$

satisfying $h^{-1}(X' \times_X E) = h^{-1}(\mathrm{pr}_1^{-1}(E)) = (\mathrm{pr}_1 \circ h)^{-1}(E) = \widetilde{f}^{-1}(E) = E'$.

It is natural to ask when h is an isomorphism, in which case we may say that the formation of the blow-up $\mathrm{Bl}_Z(X)$ *commutes with base change along* f. When f is an open immersion then h is an isomorphism; this expresses the evident fact from the construction that the formation of the blow-up is Zariski-local on X. A generalization of an open immersion is a flat map:

> PROPOSITION A.3.4.1. If $f : X' \to X$ is flat then $h : \widetilde{X'} \to X' \times_X \widetilde{X}$ is an isomorphism. In general, for arbitrary f, h is a closed immersion.

Proof. The problem is Zariski-local on X and then on X', so we may assume each is affine: $X = \mathrm{Spec}\, A$ and $X' = \mathrm{Spec}\, A'$ (so A' is an A-algebra), with Z corresponding to an ideal $I \subset A$ and Z' corresponding to the ideal $I' = IA' \subset A'$. the evident ring map $\varphi : S = \bigoplus_{n \geq 0} I^n \to \bigoplus_{n \geq 0} I'^n = S'$ over $A \to A'$ makes $\varphi(S_+)$ generate S'_+ over $S'_0 = A'$ since $I' = IA'$, so the "functoriality" of Proj yields an everywhere-defined morphism

$$\widetilde{X'} = \mathrm{Proj}(S') \to \mathrm{Proj}(S) = \widetilde{X}$$

over $\mathrm{Spec}(A') \to \mathrm{Spec}(A)$. By the uniqueness of $\widetilde{f}$, this map of Proj's must be $\widetilde{f}$. Hence, since $X' \times_X \widetilde{X} = \mathrm{Proj}(A' \otimes_A (\bigoplus_{n \geq 0} I^n)) = \mathrm{Proj}(\bigoplus_{n \geq}(A' \otimes_A I^n))$, h is likewise the map of Proj's induced by the map of

graded A'-algebras

$$\bigoplus_{n\geq 0}(A'\otimes_A I^n)\to\bigoplus_{n\geq 0}I'^n.$$

But this map of graded rings is surjective since $A'\otimes_A I\to IA'=I'$ is surjective, so the associated map h between their Proj's is a closed immersion.

It remains to show that h is an isomorphsm when f is flat, and for that purpose it suffices to show the map of graded A'-algebras giving rise to h is an isomorphism when A' is A-flat. Since $I'^n=I^nA'$, we just need to show that for any ideal $J\subset A$, if A' is A-flat then the natural map $A'\otimes_A J\to JA'$ is an isomorphism. This latter map is certainly surjective, and for injectivity it suffices to analyze it after composing with the inclusion of JA' into A'. But that composite map is the result of applying the *exact* functor $A'\otimes_A(\cdot)$ to the A-linear injection of J into A, so we are done. □

For a blow-up $q:\widetilde{X}=\mathrm{Bl}_{\mathcal{I}}(X)\to X$ and a section s of $\mathcal{I}$ over an open subscheme $W\subset X$, we define $U(s)$ to be the open subscheme of $q^{-1}(W)$ over which $q^*(s)$, viewed as a section of $\mathcal{I}\cdot\mathcal{O}_{\widetilde{X}}$, is a local generator. (If $W=X=\mathrm{Spec}(A)$ is affine, so $\mathcal{I}=\widetilde{I}$ for an ideal $I\subset A$, then $U(s)$ corresponds to what is called $D_+(s)$ in the construction of blow-ups in the affine case for an element $s\in I$.) We will sometimes abuse notation and speak of s, rather than $q^*(s)$, as a section of $\mathcal{I}\cdot\mathcal{O}_{\widetilde{X}}$ when there is no possibility of confusion.

EXERCISE A.3.4.A. Let $\mathcal{I}_1$ and $\mathcal{I}_2$ be two quasi-coherent ideal sheaves on X and let $\widetilde{X}_i=\mathrm{Bl}_{\mathcal{I}_i}(X)$. Show a natural commutative diagram

$$\begin{array}{ccc}\widetilde{X}_{12}&\longrightarrow&\widetilde{X}_1\\ \downarrow&&\downarrow\\ \widetilde{X}_2&\longrightarrow&X\end{array}\qquad\text{(A.3.4-II)}$$

in which $\widetilde{X}_{12}=\mathrm{Bl}_{\mathcal{I}_1\mathcal{I}_2}(X)$. *[Hint: Show that $\mathcal{I}_1\cdot\mathcal{O}_{\widetilde{X}_{12}}$ is invertible on $\widetilde{X}_{12}$, and likewise with $\mathcal{I}_2$ replacing $\mathcal{I}_1$, and use the universal property of blowups.]*

In particular, by induction and the above Exercise there are unique X-isomorphisms $\mathrm{Bl}_{\mathcal{I}^n}(X)\xrightarrow{\sim}\mathrm{Bl}_{\mathcal{I}}(X)$ for all $n\geq 1$; these isomorphisms are induced by applying relative Proj to the obvious explicit map on the level of quasi-coherent graded $\mathcal{O}_X$-algebras.

PROPOSITION A.3.4.2. For closed subschemes Y and Z in a scheme X, the strict transform $\widetilde{Y}$ of Y in $\widetilde{X}=\mathrm{Bl}_{Y\cap Z}(X)$ lies inside of the union of the $U(s)$'s as s ranges through all sections of $\mathcal{I}_Z$ over all opens.

Moreover, if $Y_1,\ldots,Y_n$ are closed subschemes of X then the strict transforms $\widetilde{Y}_j$ in $\mathrm{Bl}_{\cap Y_i}(X)$ satisfy $\cap\widetilde{Y}_j=\emptyset$.

Proof. We may assume $X=\mathrm{Spec}(A)$ is affine. For the first claim, let $J=\Gamma(X,\mathcal{I}_Z)$ and $I=\Gamma(X,\mathcal{I}_Y)$. Clearly $\widetilde{X}$ is covered by $U(s)$ for $s\in I\cup J$ (as this union generates $I+J$), so it is enough to show that for $s\in I$, $U(s)$ is disjoint from the strict transform $\widetilde{Y}=\widetilde{V(I)}$ (as then $\widetilde{Y}$ is covered by the other $U(s)$'s, namely for $s\in J$). In terms of commutative algebra, the overlap of the A_s-scheme $U(s)=\mathrm{Spec}(A[(I+J)s^{-1}])$

and the A-scheme $\widetilde{V(I)}$ in $\widetilde{X}$ is the affine scheme attached to the ring

$$A[(I+J)s^{-1}]/(I_s \cap A[(I+J)s^{-1}]),$$

with intersection $I_s \cap A[(I+J)s^{-1}]$ taken inside of A_s. But $1 \in I_s \cap A[(I+J)s^{-1}]$ since $s \in I$, so the overlap is empty.

For the second claim, let $I_j = \Gamma(X, \mathcal{I}_{Y_j})$, so since $\mathcal{I}_{\cap Y_j} = \sum \mathcal{I}_{Y_j}$ we have that the $U(s)$'s for $s \in \cup I_j$ cover $\widetilde{X} = \mathrm{Bl}_{\cap Y_j}(X)$. It therefore suffices to show that for $s \in I_1$ the open set $U(s)$ is disjoint from $\widetilde{V(I_1)}$ (as then anything in every $\widetilde{V(I_j)}$ is outside of $U(s)$ for all $s \in \cup I_j$, yet such $U(s)$'s cover $\widetilde{X}$ and so there are no points in every $\widetilde{Y}_j$). This goes exactly as in the preceding calculation, setting $I = I_1$ and $J = I_2 + \cdots + I_n$ (so $J = 0$ if $n = 1$). □

Appendix B

Galois Cohomology

"*It is less than four years since cohomological methods (i.e. methods of Homological Algebra) were introduced into Algebraic Geometry in Serre's fundamental paper,... and it seems certain that they are to overflow the part of mathematics in the coming years, from the foundations up to the most advanced parts..*"

The Cohomology Theory of Abstract Algebraic Varieties
Alexander Grothendieck

Section B.1

Primer – Profinite Group Cohomology

§B.1.1 Definining Cohomology

Definition B.1.1.1. For G a profinite-group, a **discrete G-module** is an abelian group A endowed with a group action $\alpha : G \times A \to A$ that is continuous when A is given the discrete topology. The category of all such modules is denoted $\mathrm{disc}(G)$.

Warning. Often in the literature the above are simply referred to as "G-modules" in the context of group cohomology, (which of course has a more general definition as simply an abelian group endowed with an action without any assumption of G even being a topological group).

Exercise B.1.1.A. Show that we may equivalently define a discrete G-module as a G-module A such that the stabilizer of each element is open.

It is for the objects of $\mathrm{disc}(G)$ that group cohomology is defined. There are a couple of different approaches one can take for this definition, both equivalent and ultimately whichever one prefers amounts to a matter of taste. We will give an overview of both and a rough proof of their equivalence.

For any $A \in \mathrm{disc}(G)$, we define the group of **n-cochains** $C^n(G, A)$ to be group of continuous (i.e. locally constant, due to discreteness of A) functions $G^n \to A$. For example, we have $C^0(G, A) \cong A$. We may then define a **coboundary** homomorphism (familiar to any prior cohomology theory the reader may have learned) by

$$\begin{aligned} \delta_n : C^n(G, A) &\to C^{n+1}(G, A) \\ f(g_1, \ldots, g_n) &\mapsto g_1 \cdot f(g_2, \ldots, g_{n+1}) \\ &\quad + \sum_{i=1}^{i=n} (-1)^i f(g_1, \ldots, g_i g_{i+1}, \ldots, g_{n+1}) \\ &\quad + (-1)^{n+1} f(g_1, \ldots, g_n). \end{aligned}$$

One can verify without too much effort that in fact $\delta_{n+1} \circ \delta_n = 0$, and thus we obtain a complex $C^\bullet(G, A)$, and thus we may define

$$H^q(G, A)$$

to be the cohomology groups of that complex.

Example B.1.1.2. The group $H^1(G, A)$ can be identified with the group of **crossed homomorphisms** $G \to A$. These are group homomorphisms $\phi G \to A$ twisted by a g-action, i.e.

$$\phi(gh) = \phi(g) + g \cdot \phi(h) \in A.$$

If $\{G_i\}_{i \in I}$ is an inverse system of profinite groups, and $\{A_i\}_{i \in I}$ is a corresponding directed system of discrete G_i-modules (with the homomorphisms $A_i \to A_j$ being compatible with morphisms $G_i \to G_j$) then

it is not too hard to see that the canonical homomorphism

$$\varinjlim_{i\in} C^\bullet(G_i, A_i) \to C^\bullet(G, A) \tag{B.1.1-I}$$

is an isomorphism of chain complexes. We leave the reader to verify this.

EXERCISE B.1.1.B. Verify the isomorphism in Eqn. B.1.1-I, and its following corollaries:

$$H^q(G, A) = \varinjlim_{\text{open } U \trianglelefteq G} H^q(G/U, A^U) = \varinjlim_{\text{fin-gen } B \subseteq A} H^q(G, B),\ \forall q \geq 0.$$

KEY IDEA. Various facts and computations in group cohomology can be reduced to the case of finite groups, via the above result.

While the definition from cochains is a nice explicit description, it can also be psychologically useful, even if materially unnecessary, to give a much more categorial description. To this end, for any $A \in \text{disc}(G)$, define

$$A^G := \{a \in A : ga = a\ \forall g \in G\}$$

to be the **invariants** of A. Note by our prior definition that we have

$$\begin{aligned} H^0(G, A) &:= \ker\left(\delta_0 : A \to C^1(G, A)\right) \\ &= \ker(a \mapsto (g \mapsto g \cdot a - a)) \\ &= A^G. \end{aligned}$$

Thus the functor of invariants $F : A \mapsto A^G$ is *precisely* the 0-th cohomology functor we defined earlier. One verifies directly that this functor is left-exact on $\text{disc}(G)$, and thus, if we can show that $\text{disc}(G)$ has enough injectives, we can define the cohomomology groups by the right-derived functors of this, i.e.

$$H^i(G, A) := R^i F(A).$$

LEMMA B.1.1.3. The category of discrete G modules $\text{disc}(G)$ has enough injectives.

Proof. Take any $A \in \text{disc}(G)$, we want to construct a monomorphism into an injective object. We know that, as an object in the category of abelian groups, A injects into some injective I. So consider

$$I_G := C(G, I) := \{\text{cont. } f : G \to I\},$$

where I is given the discrete topology. We give this the G-action given by

$$g \cdot f(x) = f(xg).$$

Note we have an isomorphism of hom-sets

$$\begin{aligned} \text{Hom}_{\underline{\text{Ab}}}(A, I) &\xrightarrow{\sim} \text{Hom}_{\text{disc}(G)}(A, I_G) \\ \phi &\mapsto a \mapsto (x \mapsto \phi(x \cdot a)). \end{aligned}$$

Thus we see that the contravariant functor $I \mapsto I_G$ is a right adjoint to the forgetful functor from $\text{disc}(G)$ to $\underline{\text{Ab}}$. Note that if $\phi : A \hookrightarrow I$ is the injection into the injective object in the category of abelian groups,

then ϕ_G (the image of ϕ) is also monomorphism, since

$$\phi_G(a)(e_G) = \phi(a) \in I.$$

Thus it remains to show that this object I_G is injective in $\mathrm{disc}(G)$. Since $\mathrm{disc}(G)$ is an abelian category (since all cokernel, kernel constructions follow from those on the underlying abelian groups), it suffices to show that the Hom-functor $X \mapsto \mathrm{Hom}_{\mathrm{disc}(G)}(X, I_G)$ is exact. But by our adjunction, we have

$$\mathrm{Hom}_{\mathrm{disc}(G)}(X, I_G) \cong \mathrm{Hom}_{\underline{\mathrm{Ab}}}(X, I),$$

and the RHS above is exact in X by injectivity of I. Thus I_G is injective, completing our proof. □

Remark B.1.1.4 (Unimportant but interesting). In contrast to the lemma above, the reader may be interested to know that $\mathrm{disc}(G)$ does not always have enough projectives, it in fact has *no non-zero projectives* if G is infinite! The proof isn't too difficult but is a tad too lengthy and tangential to our exposition here so we offset it to [CK23].

It thus follows from the above Lemma that we can define the right-derived functors of the functor of invariants $A \to A^G$ and set those to be our cohomology groups. We want to check that this aligns with our prior constructions with continuous cochains. We will resolve this with a sledgehammer, courtesy of Grothendieck. Recall an additive functor between abelian cateogires $F : \mathcal{C} \to \mathcal{D}$ is said to be **erasable** (also often called **effaceable**) if for any $A \in \mathcal{C}$ there exists a monomorphism $u : A \hookrightarrow I$ such that $F(u) = 0$.

Theorem B.1.1.5. Let $\{T^i\}_{i\in\mathbb{N}} : \mathcal{C} \to \mathcal{D}$ be a collection of additive functors between additive categories that form a cohomological delta functor, with the property that T^i is *erasable* for all $i > 0$. Then $\{T^i\}_{i\in\mathbb{N}}$ is a universal delta functor.

Proof. See Grothendiecks' original "Tohoku paper", [Gro57]. □

Another way of saying this is that any additive functor has a unique extension to a δ functor if we impose the condition that all of the higher order functors be erasable.

Naturally, any right-derived functor is erasable, since they send all injectives to the zero object, so to conclude the equality of the two cohomologies we would need to show that the explicitly describe higher cohomology functors $H^q(G, A)$ are erasable.

Exercise B.1.1.C. For any $A \in \mathrm{disc}(G)$ and the corresponding I_G we constructed in the proof of Lemma B.1.1.3, show that $H^q(G, I_G)$ vanishes for all $q \geq 0$, for the cochain definition of cohomology. Conclude – using Theorem B.1.1.5 – that our two different definitions of cohomology coincide.

Now that we have established our cohomology theory over profinite groups, we discuss some important classes of morphisms between cohomology groups.

Suppose G, G' are two profinite groups, and suppose we have a homomorphism $\phi : G \to G'$. Given $A \in \mathrm{disc}(G), A' \in \mathrm{disc}(G')$, and a morphism $\psi : A' \to A$ which is compatible with ϕ (i.e. is a G-module homomorphism if we regard A' as a G-module via ϕ), we can pass to cohomology to obtain a homomorphism

$$H^q(G', A') \to H^q(G, A).$$

DEFINITION B.1.1.6. For H a closed subgroup of a profinite group G, and $A \in \mathrm{disc}(G)$, if we view A as a discrete H module then the identity map id_A and canonical injection $H \hookrightarrow A$ induce the **restriction homomorphism**

$$\mathrm{Res} : H^q(G, A) \to H^q(H, A) \ \forall q \geq 0.$$

We want to write a sort of reverse-direction map for the restriction homomorphism. In order to define this, we will need to discus the notion of *coinduced modules*, which will be similar to that of representation theory.

DEFINITION B.1.1.7. Suppose H is a closed subgroup of a profinite group G, and take some $A \in \mathrm{disc}(H)$. The **coinduced module** (sometimes just the **induced module** in some conventions, though ideally this terminology should be used for homology) of A with respect to H in G is defined by

$$M_G^H(A) := \{\text{cont. } \alpha : G \to A \text{ s.t. } \alpha(hg) = h \cdot \alpha(g) \ \forall h \in H, g \in G\}.$$

It is given a G-action by setting $g \cdot \alpha(x) := \alpha(xg)$. In the case that $H = \{e_G\}$ one often just writes $M_G(A)$.

REMARK B.1.1.8. Often the notation $\mathrm{CoInd}_G^H(A)$ is also used for $M_G^H(A)$, but we opt for less transparent but easier to read option.

Evaluating at the identity of G, we obtain a homomorphism

$$\begin{aligned} M_G^H(A) &\to A \\ \alpha &\mapsto \alpha(e_G) \end{aligned}$$

which is compatible with the injection $H \hookrightarrow G$. Thus we get an induced homomorphism on cohomology:

$$H^q(M_G^H(A)) \to H^q(H, A) \tag{B.1.1-II}$$

PROPOSITION B.1.1.9 (Shapiro's Lemma). The map in Eqn. B.1.1-II is an isomorphism for all $q \geq 0$. Thus all the higher cohomologies of any $M_G(A)$ are 0.

Proof. Note for any $A \in \mathrm{disc}(H)$ and $X \in \mathrm{disc}(G)$ we have a (visibly functorial) bijection of homsets

$$\begin{aligned} \mathrm{Hom}_{\mathrm{disc}(G)}(X, M_G^H(A)) &\xrightarrow{\sim} \mathrm{Hom}_{\mathrm{disc}(H)}(X, A) \\ \phi &\mapsto (x \mapsto \phi(x)(e_G)). \end{aligned}$$

It follows then that the forgetful functor from discrete G-modules to discrete H-modules is left-adjoint to the coinduction functor

$$A \mapsto M_G^H(A).$$

Thus, as a right-adjoint functor, coinduction is left-exact, and in particular preserves injective resolutions. In particular, picking any injective resolution $A \to I^\bullet$ in $\mathrm{disc}(H)$, the corresponding

$$M_G^H(A) \to M_G^H(I^\bullet)$$

is an injective resolution in $\mathrm{disc}(G)$. Moreover note for any coinduced module we have

$$M_G^H(A)^G \cong A^H.$$

Indeed, any G-invariant $(f : G \to A) \in M_G^H(A)$ must just be constant, and the H-equivariance condition imposed means any such f can only take values in A^H. Thus the applying the functor of G-invariants to the co-induced injective resolution yields

$$A^H \to (I^H)^\bullet.$$

Taking cohomology on one side of the isomorphism yields $H^\bullet(G, M_G^H(A)$, and taking cohomology on the other side yields $H^\bullet(H, A)$, giving us the desired isomorphism. □

Now suppose H is an open subgroup of G, and take some $A \in \mathrm{disc}(G)$. We can define a surjective G-homomorphism

$$\begin{aligned} \pi : M_G^H(A) &\to A \\ \alpha &\mapsto \sum_{\overline{g} \in G/H} g^{-1} \cdot \alpha(g). \end{aligned}$$

Sanity Check. Verify the sum in the above equation is independent of choice of representative of $g \in \overline{g} \in G/H$.

Upon passing to cohomology, this map π gives rise to a homomorphism

$$H^q(G, M_G^H(A))(= H^q(H, A)) \to H^q(G, A) \tag{B.1.1-III}$$

Definition B.1.1.10. For H an open subgroup of a profinite group G, and for any $A \in \mathrm{disc}(G)$, we define the **corestriction** map (also sometimes called the **inflation** map)

$$\mathrm{Cor} : H^q(H, A) \to H^q(G, A)$$

to be the one in Eqn. B.1.1-III.

Proposition B.1.1.11. For any open subgroup H of a profinite group G, with index $[G : H] = n$, and any $A \in \mathrm{disc}(G)$, the composition $\mathrm{Cor} \circ \mathrm{Res}$ is multiplication by n on all $H^q(G, A)$.

Proof. Recall Shapiro's Lemma (B.1.1.9) yields an isomorphism

$$H^q(G, M_G^H(A)) \xrightarrow{\sim} H^q(H, A)$$

After a little bit of thought, it is not too hard to verify that the map

$$\begin{aligned} \phi : A &\to M_G^H(A) \\ a &\mapsto (g \mapsto g \cdot a) \end{aligned}$$

induces a homomorphism on cohomology $H^q(G, A) \to H^q(M_G^H(A))$ that, when composed with the Shapiro isomorphism, is exactly the restriction map. Moreover we also know by definition the corestriction map is the one induced by

$$\psi : M_G^H(A) \to A$$
$$\alpha \mapsto \sum_{\overline{g} \in G/H} g^{-1} \cdot \alpha(g).$$

Thus the composition $\text{Cor} \circ \text{Res}$ must be induced by the map $\psi \circ \phi : A \to A$, which is

$$a \mapsto \sum_{\overline{g} \in G/H} g^{-1} \cdot g \cdot a = [G : H]a,$$

i.e. this is just multiplication by n. But then one checks without much effort that this must therefore induce the multiplication-by-n map on cohomology groups as well, which was what we wanted. □

This fact can be used to prove surprising things, such as the following result.

EXERCISE B.1.1.D. If the index $[G : H] = n < \infty$, show the kernel of the restriction homomorphism is n-torsion. Conclude that, in general, if $[G : H]$ (possible non-finite!) is coprime to p a prime number, then restriction is injective on the p-primary component of $H^q(G, A)$. [*Hint: Reduce to the finite case by writing H as an intersection of open subgroups, then use Exercise B.1.1.B*]

§B.1.2 Cohomological Dimension

The notion of dimension will be a measure of "cohomological complexity" of our group G. Intuitively, it asks how high in degree we have to climb in order for all cohomological obstructions vanish regardless of the module of coefficients. We begin by leaving the reader to verify a useful technical lemma.

EXERCISE B.1.2.A. If M is any torsion group, let $M(p)$ denote its p-torsion component. Show that for any $A \in \text{disc}(G)$

$$H^q(G, A(p)) \cong H^q(G, A)(p).$$

DEFINITION B.1.2.1. For G a profinite group and p a prime, the **p-cohomological dimension** of G, denoted $\text{cdim}_p(G)$, to be the smallest integer n satisfying the following condition.

(*) For every torsion $A \in \text{disc}(G)$, and every $q > n$, we have $H^q(G, A)(p) = 0$.

We then define the **cohomological dimension** of G to be

$$\text{cdim}(G) := \sup_{p \text{ prime}} \text{cdim}_p(G).$$

SANITY CHECK. Note if

Note of course that the (p)-cohomological dimension may very well be $+\infty$.

PROPOSITION B.1.2.2. For G profinite, p prime, and n an integer, the following are equivalent:

(a) $\text{cdim}_p(G) \leq n$.

(b) $H^q(G, A) = 0$ for all $q > n$ and every p-torsion A.

(c) $H^{n+1}(G, A) = 0$ when A is simple and killed by p.

Proof. The equivalence of (i) and (ii) is a direct consequence of Exercise B.1.2.A. The fact that (ii) implies (iii) follows essentially from definitions, so it suffices to show that (iii) $\implies$ (ii). [complete later] □

PROPOSITION B.1.2.3. Suppose $\mathrm{cdim}_p(G) \leq n$, and suppose $A \in \mathrm{disc}(G)$ is p-divisible. Then $H^q(G, A)(p) = 0$ for $q > n$. As a corollary, if A is divisible and $\mathrm{cdim}(G) \leq n$, then $H^q(G, A) = 0$ for $g > n$.

Proof. Consider the exact sequence

$$0 \to A(p) \to A \xrightarrow{p} A \to 0,$$

which induces the exact sequence on cohomology groups

$$H^q(G, A(p)) \to H^q(G, A) \xrightarrow{p} H^q(G, A).$$

For $q > n$, we have $H^q(G, A(p)) = 0$ by hypothesis, and thus it follows that multiplication by p is injective on $H^q(G, A)$, hence it must have no p-torsion, as desired. □

We also have the notion of *strict* cohomological dimension, which tests the vanishing of higher cohomologies on *all* discrete G-modules (i.e. not just the torsion ones).

DEFINITION B.1.2.4. For G a profinite group and p a prime, the **strict p-cohomological dimension** of G, denoted $\mathrm{scdim}_p(G)$, to be the smallest integer n satisfying the following condition.

(*) For every $A \in \mathrm{disc}(G)$, and every $q > n$, we have $H^q(G, A)(p) = 0$.

We then define the **strict cohomological dimension** of G to be

$$\mathrm{scdim}(G) := \sup_{p \text{ prime}} \mathrm{scdim}_p(G).$$

Since strict cohomological dimension tests on a larger category of coefficient modules, it is clear that $\mathrm{scdim}_p(G) \geq \mathrm{cdim}_p(G)$ for all p (and thus $\mathrm{scdim}(G) \geq \mathrm{cdim}(G)$). But in fact we actually have a much stronger statement.

PROPOSITION B.1.2.5. For any profinite G and any prime p, $\mathrm{scdim}_p(G)$ equals $\mathrm{cdim}_p(G)$ or $\mathrm{cdim}_p(G)+1$.

Proof. It suffices to show $\mathrm{scdim}_p(G) \leq \mathrm{cdim}_p(G) + 1$. Take any $q > \mathrm{cdim}_p(G) + 1$. For any $A \in \mathrm{disc}(G)$, since $A(p)$ and A/pA are both p-torsion groups, it follows that

$$H^q(G, A(p)) = H^q(G, A/pA) = 0.$$

Then, the exact sequence

$$0 \to A(p) \to A \to pA \to 0,$$

which induces a corresponding exact sequence on cohomology, it follows that the map $H^q(G, A) \to H^q(G, pA)$ is injective. Applying similar logic to the exact sequence

$$0 \to pA \to A \to A/pA \to 0,$$

we deduce the map $H^q(G, A) \to H^q(G, A/pA)$ is injective as well. Since the composition $A \to pA \to A$ is multiplication by p, composition of maps induced on cohomology is also multiplication by p. Injectivity of that composition then implies that $H^q(G, A)$ has no p-torsion, i.e. $H^q(G, A)(p) = 0$. Applying this over all $A \in \mathrm{disc}(G)$, it follows that $\mathrm{scdim}_p(G) \leq \mathrm{cdim}(G) + 1$, as desired. □

Note that Shapiro's Lemma (B.1.1.9) tells us that (strict) cohomological dimension does not increase upon passage to a closed subgroup $H \leq G$. But in fact we can show something stronger.

EXERCISE B.1.2.B. For H a closed subgroup of a profinite group G, and p a prime, show that $\mathrm{cdim}_p(H) = \mathrm{cdim}_p(G)$ and $\mathrm{scdim}(H) = \mathrm{scdim}_p(G)$ when either of the following conditions hold:

(a) The (possibly infinite!) index $[G : H]$ is coprime to p.

(b) H is open in G, and $\mathrm{cdim}_p(G) < \infty$. *[Hint: It might be helpful to prove that the corestriction map $H^n(H, A) \to H^n(G, A)$ is surjective on the p-torsion components.]*

The above exercises have the following corollaries.

COROLLARY B.1.2.6. For G a profinite group, the following hold:

(a) If H is a Sylow p-subgroup of G, then

$$\mathrm{cdim}_p(G) = \mathrm{cdim}_p(H) = \mathrm{cdim}(H),$$

and the same equalities hold for strict (p)-cohomological dimension.

(b) The order of G is coprime to p if and only if $\mathrm{cdim}_p(G) = 0$.

(c) In $\mathrm{cdim}_p(G) \neq 0$ or ∞, the exponent of p in the order of G is infinite.

(d) $\mathrm{cdim}_p(G) = n < \infty$ if and only if every open subgroup $H \leq G$ satisfies $H^{n+1}(H, \mathbb{Z})(p) = 0$.

Proof. [lot of it is fairly straightforward, perhaps just expound on some of the parts later.] □

LEMMA B.1.2.7. Let G be any pro-p group. Then, up to isomorphism, $\mathbb{Z}/p\mathbb{Z}$ is the only simple p-torsion discrete G-module.

Proof. [later] □

PROPOSITION B.1.2.8. For G a pro-p group and n and integer, we have $\mathrm{cdim}(G) \leq n$ if and only if $H^{n+1}(G, \mathbb{Z}/p\mathbb{Z}) = 0$

Proof. Note the cohomology groups cannot have any ℓ-torsion for primes $\ell \neq p$ [prove maybe?], thus $\mathrm{cdim}(G) = \mathrm{cdim}_p(G)$. The result in then a direct consequence of part (c) of Proposition B.1.2.2 and Lemma B.1.2.7. □

Section B.2

Commutative Galois Cohomology

§B.2.1 First Definitions and Examples

Let k be a field, and K/k a Galois extension, then we know the Galois group $\mathrm{Gal}(K/k)$ is a profinite group (built as the inverse limit of the system of all finite intermediary extensions!), and thus we may apply the methods and results of the previous section here – if A is a discrete group acted on by the Galois group it makes sense to discuss the cohomology groups

$$H^q(\mathrm{Gal}(K/k), A).$$

In general it is more convenient to not fix a field extension a priori and frame things in the following way. We fix a **ground field** k, and a covariant functor

$$K \to A(K)$$

from the category of algebraic separable extension of k to abelian groups [maybe handle non-abelian case in future draft] satisfying the following conditions:

(i) $A(K) = \varinjlim_{k-\text{finite} K_i \subseteq K} A(K_i)$.

(ii) If $K \to K'$ is an injection then so is $A(K) \to A(K')$.

(iii) If K'/K is a Galois extension, then $A(K) = H^0(\mathrm{Gal}(K'/K), A(K'))$.

Sanity Check. Note the final condition always makes sense since first condition specifies that $\mathrm{Gal}(K'/K)$ must act continusly and functorially on $A(K')$.

> Exercise B.2.1.A. Suppose A is a locally finite-type commutative k-group scheme, then show its functor of points satisfies all the above conditions.

The most basic examples include $\mathbb{G}_a : K \mapsto K$ and $\mathbb{G}_m : K \mapsto K^\times$.

> Proposition B.2.1.1. For every Galois K/k, we have $H^1(K/k, \mathbb{G}_m) = 0$ and $H^q(K/k, \mathbb{G}_a) = 0$ for all $q \geq 1$.

Proof. We first deal with $\mathbb{G}_a$. Since any Galois extension is built as the colimit of finite Galois extensions, it suffices to show the result for any finite Galois K/k. We recall the **normal basis theorem**, which says there exists some $b \in K$ such that its Galois conjugates form a k-basis of K. In other words, we can write $K \cong k[\mathrm{Gal}(K/k)]$. But note

$$k[\mathrm{Gal}(K/k)] = \{f : \mathrm{Gal}(K/k) \to k\} = M_{\mathrm{Gal}(K/k)}(k),$$

i.e. K is *coinduced* from K. But then Shapiro's Lemma B.1.1.9 tells us that all the higher cohomologies of K vanish, as desired.

[finish G_m later] □

DEFINITION B.2.1.2. For A a functor satisfying the desired axioms over a field k, the **absolute Galois cohomology** consists of the groups

$$H^q(k^{\text{sep}}/k, A),\ q \geq 0.$$

Often we omit the notation of the extension entirely and just write $H^q(k, A)$.

PROPOSITION B.2.1.3. Let n be a positive integer coprime to the $\text{char}(k)$, and let μ_n be the group of n-th roots of unity in k^{sep}. Then

$$H^1(k, \mu_n) \cong k^\times/(k^\times)^n.$$

Proof. We have the exact sequence

$$1 \to \mu_n \to \mathbb{G}_m(k) \xrightarrow{n} \mathbb{G}_m \to 1,$$

where "n" denotes the n-th power endomorphism. This induces the exact sequence in cohomology

$$H^0(k, \mathbb{G}_m) \xrightarrow{n} H^0(k, \mathbb{G}_m) \to H^1(k, \mu_n) \to H^1(k, \mathbb{G}_m) \to \cdots$$

By Proposition B.2.1.1, we know the rightmost object in the above sequence is 0, and thus $H^1(k, \mu_n) \cong k^\times/(k^\times)^n$, as desired. □

§B.2.2 Brauer Groups

Brauer groups are objects which classify central simple algebras over a field, up to some notion of equivalence. The reason Brauer groups belong in this section is due to the isomorphism $H^2(G_k, k^\times) \simeq \text{Br}(k)$ for k a field and G_k its absolute Galois group, which we will prove later (see reference theorem once proven). However, in this book, we also see the usefulness of Brauer groups as independent objects to classify endomorphism rings of abelian varieties (see theorem 2.1.5.2). As we will see in this section, both the Galois cohomology perspective and central simple algebra are useful and have a push-and-pull relationship for giving us results about the other.

Brauer groups appear in many different texts which we will list here: compile references (FOR NOW: Colliot-Thelene, Skorobogatov's Brauer-Grothendieck group book for the most comprehensive and general setting, Stacks Project, Jacobson's Basic Algebra II, etc.).

§B.2.3 Fields of dimension ≤ 1

The main result of this subsection will be the following.

THEOREM B.2.3.1. Let k be a field, with absolute Galois group G_k. Then TFAE:

(a) $\text{cdim}(G_k) \leq 1$. If, moreover, $\text{char}(k) = p > 0$, then $\text{Br}(K)(p) = 0$ for every algebraic extension K/k.

(b) $\mathrm{Br}(K) = 0$ for every algebraic K/k.

(c) For any algebraic K/k and finite Galois L/K, the $\mathrm{Gal}(L/K)$-module $L^\times$ is cohomologically trivial.

(d) For any algebraic K/k and finite Galois L/K, the norm $N_{L/K} : L^\times \to K^\times$ is surjective.

(e) All of the above statements, but restricting to just the finite separable extensions K/k.

DEFINITION B.2.3.2. We say a field k has **dimension** ≤ 1 if any of the equivalent conditions in Theorem B.2.3.1 are satisfied. We will write $\dim(k) \leq 1$.

SANITY CHECK. Verify that if $\dim(k) \leq 1$, then the same is true for any algebraic extension K/k.

EXERCISE B.2.3.A. If k is perfect, show that $\dim(k) \leq 1$ if and only if $\mathrm{cdim}(G_k) \leq 1$. *[Hint: Use bijectivity of the Frobenius endomorphism to deduce vanishing of certain cohomology groups.]*

To begin with we discuss various characterizations and computations of cohomological dimension of absolute Galois groups G_k.

DEFINITION B.2.3.3. For G a profinite group, we denote by G_p the largest pro-p quotient of G. In the case that $G = G_k$ is the absolute Galois group of k, the group $G_{k,p}$ is the Galois group of the extension $k^{\mathrm{sep}}(p)/k$, which is called the **maximal p-extension** of k.

LEMMA B.2.3.4. For G a profinite group and p a prime, suppose $G_p = G/N$. Suppose $\mathrm{cdim}_p(N) \leq 1$. Then the canonical maps

$$H^q(G_p, \mathbb{Z}/p\mathbb{Z}) \to H^q(G, \mathbb{Z}/p\mathbb{Z})$$

are isomorphisms for all $q \geq 0$. In particular $\mathrm{cdim}(G_p) \leq \mathrm{cdim}_p(G)$.

Proof. Let N/M be the largest pro-p quotient of N. It is not too hard to verify that M is normal in G and G/M is a pro-p group. But note by definition of G_p this implies $M = N$, i.e. every morphism of N into a pro-p group is trivial. In particular, we have

$$H^1(N, \mathbb{Z}/p\mathbb{Z}) = 0.$$

Moreover $\mathrm{cdim}_p(N) \leq 1 \implies H^i(N, \mathbb{Z}/p\mathbb{Z}) = 0$ for all $i \geq 2$. From the spectral sequence of group extensions, [expand later] then, it follows the desired maps must in fact be an isomorphism. □

EXERCISE B.2.3.B. Let k be a field of dimension ≤ 1, and p a prime. Show that $\mathrm{cdim}(G_{k,p}) \leq 1$.

We will use this lemma to understand the p-cohomological dimension of G_k in two separate cases: in the equal and unequal characeristic.

PROPOSITION B.2.3.5 (Equal Characteristic Case). If $\mathrm{char}(k) = p$, then $\mathrm{cdim}(G_{k,p}) \leq \mathrm{cdim}_p(G_k) \leq 1$.

Proof. Set $f(x) := x^p - x$. Then we have the exact sequence

$$0 \to \mathbb{Z}/p\mathbb{Z} \to \mathbb{G}_a \xrightarrow{f} \mathbb{G}_a \to 0.$$

Passing to cohomology, we get the short exact subsequence:

$$H^1(G_k, \mathbb{G}_a) \to H^2(G_k, \mathbb{Z}/p\mathbb{Z}) \to H^2(G_k, \mathbb{G}_a)$$

By Proposition B.2.1.1, we know the left and right ends above are 0, and it thus follows that $H^2(G_k, \mathbb{Z}/p\mathbb{Z}) = 0$. The same equality can be deduced for any closed subgroup of G_k (since these are all Galois groups), and thus in particular to its Sylow p-subgroups. For any such subgroup H, we see that $H^2(H, \mathbb{Z}/p\mathbb{Z}) = 0$ implies that $\mathrm{cdim}(H) \leq 1$, by Proposition B.1.2.8. Then by part (a) of Corollary B.1.2.6, it follows then that $\mathrm{cdim}_p(G_k) \leq 1$.

If $N := \ker(G_k \to G_{k,p})$, we can apply identical logic to N to deduce that $\mathrm{cdim}_p(N) \leq 1$. Thus we can apply Lemma B.2.3.4 to deduce that

$$\mathrm{cdim}(G_{k,p}) \leq \mathrm{cdim}_p(G_k) \leq 1,$$

as desired. □

COROLLARY B.2.3.6 (Albert-Hochschild). If k'/k is a purely inseparable extension of a characteristic p field, then the canonical map $\mathrm{Br}(k) \to \mathrm{Br}(k')$ is surjective.

Proof. [later]. □

We now tackle the case of unequal characteristic.

PROPOSITION B.2.3.7 (Unequal Characteristic Case). Let k be a field of characterisic $\neq p$, and n and integer ≥ 1. Then TFAE:

(a) $\mathrm{cdim}_p(G_k) \leq n$.

(b) For any algebraic K/k, we have $H^{n+1}(G_K, \mathbb{G}_m)(p) = 0$ and $H^n(G_K, \mathbb{G}_m)$ is p-divisble.

(c) Same as the above assertion, but restricted to extensions K/k that are finite separable and of degree coprime to p.

Proof. [later] □

Proof of B.2.3.1. The equivalence of assertion (e) to the rest of them is either trivial or handled by Corollary B.2.3.6. □

To show a class of fields of dimension ≤ 1, we introduce the following definition.

DEFINITION B.2.3.8. We say a field k is **quasi-algebraically closed** (also often called C_1) if every $f \in k[x_1, \ldots, x_n]$ homogeneous of degree $d < n$ has a nontrivial zero in k^n.

EXERCISE B.2.3.C. Show every algebraic extension of a C_1 field is C_1.

PROPOSITION B.2.3.9. Every quasi-algebraically closed field k has dimension ≤ 1. Moreover if $\text{char}(k) = p > 0$, then $[k : k^p] = 1$ or p.

Proof. We will show that k satisfies assertion (d) of Theorem B.2.3.1. Let K be algebraic over k, and L/K Galois. Take some arbitrary $a \in K^\times$, we wish to find an element $b \in L^\times$ such that $N_{L/K}(b) = a$.

Set $d := [L : K]$. Fixing a K-basis of L, we may write any $z \in L$ with respect to said basis as a unique d-tuple $(z_1, \ldots, z_d) \in K^d$. Then, recalling we can view the norm as a polynomial in these d variables, the map

$$f(x_0, \ldots, x_d) := N_{L/K}((x_1, \ldots, x_d)) - ax_0^d$$

is an element of $K[x_0, \ldots, x_d]$. By Exercise B.2.3.C, we know K is quasi-algebraically closed, and thus this polynomial must have a non-trivial zero $(b_0, \ldots, b_d) \in K^{d+1}$. If $b_0 = 0$, this would imply that $N((b_1, \ldots, b_d)) = 0$ as well, which would only be true if all b_i were zero. Thus we must have $b_0 \neq 0$. Writing $b \in L^\times$ for the element represented by $(b_1, \ldots, b_d)$, we have

$$f(b_0, \ldots, b_d) = 0 \implies N_{L/K}(b/b_0) = a,$$

proving surjectivity of the norm, and therefore proving $\dim(k) \leq 1$.

It remains to show that $[k : k^p] = 1$ or p if $\text{char}(k) = p > 0$. Suppose $k \neq k^p$, and let K/k be a purely inseparable degree p extension. We know that $N_{K/k}(K) = k$, but $N_{K/k}(K) = K^p$. Thus $K^p = k$ and therefore $K^{p^2} = k^p$. It follows then that

$$[k : k^p] = [K : K^p] = p,$$

as desired. □

Thus to give examples of fields of dimension ≤ 1 it suffices to give examples of C_1 fields.

REMARK B.2.3.10. The converse of the above is not necessarily true; there exist perfect fields of dimension ≤ 1 that are not quasi-algebraically closed. [maybe give example]

EXAMPLE B.2.3.11. Any finite field, and any extension of transcendence degree 1 of an algebraically closed field is quasi-algebraically closed. The former statement is known as **Chevalley-Warning's theorem**, and the latter is **Tsen's theorem** [Stacks, § 59.67.10].

EXERCISE B.2.3.D. Prove (a special case of) Chevalley-Warning's theorem: any $f \in \mathbb{F}_q[x_1, \ldots, x_n]$ of total degree $d < n$ a number of solutions in $\mathbb{F}_q^n$ divisible by p. Conclude the C_1 property of $\mathbb{F}_q$ from this.

§B.2.4 Cohomology of p-adic Fields

Section B.3

Non-abelian Galois Cohomology

Appendix C

Miscellaneous Useful Theories

Section C.1

Descent

Let $p : S' \to S$ be a morphism of schemes and consider the pullback functor on the category of quasi-coherent modules

$$\begin{aligned} p^* : \underline{\mathsf{QCoh}}(S) &\to \underline{\mathsf{QCoh}}(S') \\ \mathcal{F} &\mapsto p^*\mathcal{F}. \end{aligned} \tag{C.1.0-I}$$

Question. If p has "nice enough" properties, what can we say about the image of the pullback functor p^*?

This is the starting point for descent theory. In general we may consider other categories associated to objects and other functors induced by a morphism of objects, but the fundamental question remains the same.

It would be far too much to hope that p^* is an equivalence of categories in anything but the most trivial of cases. Grothendieck's insight was that

§C.1.1 Definitions and Main Theorem

Given a morphism $p : S' \to S$, define

$$S'' := S' \times_S S',$$

and let $p_i : S'' \to S'$ denote the two projections to factors of the fiber product (note these are the same map!).

Definition C.1.1.1. With the above notation, a **covering datum** for $\mathcal{F}' \in \underline{\mathsf{QCoh}}(S')$, is an S''-isomorphism

$$\phi : p_1^*\mathcal{F}' \xrightarrow{\sim} p_2^*\mathcal{F}'.$$

The category of such pairs $(\mathcal{F}', \phi)$ will be denote $\underline{\mathsf{QCoh}}(S')^{\mathrm{cov}}$.

EXAMPLE C.1.1.2. In the affine case, if $S = \mathrm{Spec}(R), S' = \mathrm{Spec}(R')$, and pick some $M' \in \underline{R'\text{-Mod}}$, the the covering datum on $\mathcal{F}' := \mathcal{O}_{M'} \in \underline{\mathsf{QCoh}}(S')$ is given by a choice of $R' \otimes_R R'$-isomorphism

$$\phi : M' \otimes_R R' \xrightarrow{\sim} R' \otimes_R M'.$$

For example, we may use the canonical one $m' \otimes r' \mapsto r' \otimes m'$.

A morphism between elements $(\mathcal{F}', \phi)$ and $(\mathcal{G}', \psi)$ of $\underline{\mathsf{QCoh}}(S')^{\mathrm{cov}}$ is an S'-morphism $f : \mathcal{F}' \to \mathcal{G}'$ compatible with the covering data, i.e. making the following diagram commute:

$$\begin{array}{ccc} p_1^*\mathcal{F}' & \xrightarrow{\phi} & p_2^*\mathcal{F}' \\ \downarrow{\scriptstyle p_1^* f} & & \downarrow{\scriptstyle p_2^* f} \\ p_1^*\mathcal{G}' & \xrightarrow{\psi} & p_2^*\mathcal{G}' \end{array}$$

Given $\mathcal{F} \in \underline{\mathsf{QCoh}}(S)$, there is a very natural way to assign a covering datum to $p^*\mathcal{F}$, using the fact that p_1, p_2 are literally the same map, giving us the canonical isomorphism $p_1^*(p^*\mathcal{F}) = p_2^*(p^*\mathcal{F})$. Thus the functor $\mathcal{F} \mapsto p^*\mathcal{F}$ can be interpreted as a functor into $\underline{\mathsf{QCoh}}(S')^{\mathrm{cov}}$. The preliminary result is the following

PROPOSITION C.1.1.3. Take $p : S' \to S$ faithfully flat and quasi-compact, then the functor

$$\begin{aligned} p^* : \underline{\mathsf{QCoh}}(S) &\to \underline{\mathsf{QCoh}}(S)^{\mathrm{cov}} \\ \mathcal{F} &\mapsto (p^*\mathcal{F}, p_1^*(p^*\mathcal{F}) \xrightarrow{\sim} p_2^*(p^*\mathcal{F})) \end{aligned}$$

is fully faithful.

Proof. This is equivalent to showing that, for any two $\mathcal{F}, \mathcal{G} \in \underline{\mathsf{QCoh}}(S)$, setting $q := p \circ p_1 = p \circ p_2$, the diagram

$$\mathrm{Hom}_S(\mathcal{F}, \mathcal{G}) \xrightarrow{p^*} \mathrm{Hom}_{S'}(p^*\mathcal{F}, p^*\mathcal{G}) \underset{p_2^*}{\overset{p_1^*}{\rightrightarrows}} \mathrm{Hom}_{S' \times_S S'}(q^*\mathcal{F}, q^*\mathcal{G})$$

is exact. This is what we will show.

The assertion is local on S, so it suffices to prove it in the case where S is affine. Then S' is quasi-compact, covered by finitely many affines $\{S'_i\}_{i \in [n]}$. Let

$$\overline{S'} := \bigsqcup_{i \in [n]} S'_i$$

be the disjoint union. Let $u : \overline{S'} \to S'$ denote the canonical map, and let $\overline{p} := p \circ u : \overline{S'} \to S$, and let $\overline{q} : \overline{S'} \times_S \overline{S'} \to \overline{S}$ be the composition of $\overline{p}$ with the projection map(s) onto its factors. We then get the diagram

$$\begin{array}{ccccc} \mathrm{Hom}_S(\mathcal{F}, \mathcal{G}) & \xrightarrow{p^*} & \mathrm{Hom}_{S'}(p^*\mathcal{F}, p^*\mathcal{G}) & \rightrightarrows & \mathrm{Hom}_{S' \times_S S'}(q^*\mathcal{F}, q^*\mathcal{G}) \\ \| & & \downarrow{\scriptstyle u^*} & & \downarrow{\scriptstyle (u \times u)^*} \\ \mathrm{Hom}_S(\mathcal{F}, \mathcal{G}) & \xrightarrow{\overline{p}^*} & \mathrm{Hom}_{\overline{S'}}(\overline{p}^*\mathcal{F}, \overline{p}^*\mathcal{G}) & \rightrightarrows & \mathrm{Hom}_{\overline{S'} \times_S \overline{S'}}(\overline{q}^*\mathcal{F}, \overline{q}^*\mathcal{G}) \end{array}$$

Note u is faithfully flat, so the vertical maps are all injective. Thus exactness of the top row can be deduced from exactness of the bottom row. Thus we may replace $p : S' \to S$ with $\overline{p} : \overline{S'} \to S$, and thus we reduce to the case where S' is also affine.

What remains to show is given $S = \mathrm{Spec}(R), S' = \mathrm{Spec}(R')$, and a morphism of schemes corresponding to a faithfully flat morphism of rings $R \to R'$, that the diagram

$$\mathrm{Hom}_R(M,N) \xrightarrow{p^*} \mathrm{Hom}_{R'}(M \otimes_R R', N \otimes_R R') \underset{p_2^*}{\overset{p_1^*}{\rightrightarrows}} \mathrm{Hom}_{R' \times_R R'}(M \otimes_R R' \otimes_R R', N \otimes_R R' \otimes_R R').$$

is exact for any $M, N \in \underline{\mathrm{R\text{-}Mod}}$. Note $N \to N \otimes_R R'$ is injective (by faithful flatness of R'/R) and thus the map p^* in the diagram is also injective. Applying similar reasoning to the other maps in the diagram we see that showing exactness reduces to exactness of

$$M \to M \otimes_R R' \rightrightarrows M \otimes_R R' \otimes_R R', \tag{C.1.1-I}$$

which can be deduced via a faithfully flat base change as a standard exercise in commutative algebra. □

EXERCISE C.1.1.A. Show that the sequence in Eqn. C.1.1-I is exact for any faithfully flat $R \to R'$, and any R-module M, thereby finishing the proof of the last proposition.

Next we introduce the notion of descent data. Define

$$S''' := S' \times_S S' \times_S S',$$

and let $p_{i,j} : S''' \to S''$ denote the projection onto the fiber product of the ith and jth factors. Then, in order for a given $\mathcal{F}' \in \underline{\mathsf{QCoh}}(S')$ with covering datum $\phi : p_1^*\mathcal{F}' \xrightarrow{\sim} p_2^*\mathcal{F}'$ to belong to the essential image of the pullback functor, we need the following diagram to commute:

$$\begin{array}{ccccc} p_{1,2}^*p_1^*\mathcal{F}' & \xrightarrow{p_{1,2}^*\circ\phi} & p_{1,2}^*p_2^*\mathcal{F}' = p_{2,3}^*p_1^*\mathcal{F}' & \xrightarrow{p_{2,3}^*\circ\phi} & p_{2,3}^*p_2^*\mathcal{F}' \\ \| & & & & \| \\ p_{1,3}p_1^*\mathcal{F}' & & \xrightarrow{\qquad\qquad p_{1,3}^*\circ\phi \qquad\qquad} & & p_{1,3}^*p_2^*\mathcal{F} \end{array}$$

SANITY CHECK. If $\mathcal{F}' = p^*\mathcal{F}$ for some $\mathcal{F} \in \underline{\mathsf{QCoh}}(S)$, and ϕ is chosen to be the canonical isomorphism $p_1^*p^*\mathcal{F} \xrightarrow{\sim} p_2^*p^*\mathcal{F}$, then verify the above diagram commutes.

Commutivity of the above diagram is also often written more concisely as the condition that

$$p_{1,3}^* \circ \phi = p_{2,3}^* \circ \phi \circ p_{1,2}^* \circ \phi,$$

and is commonly referred to as the **cocycle condition**.

DEFINITION C.1.1.4. For $\mathcal{F}' \in \underline{\mathsf{QCoh}}(S')$, a covering datum ϕ is said to be a **descent datum** (on $\mathcal{F}'$) if it satisfies the cocycle condition mentioned above. We denote by $\underline{\mathsf{QCoh}}(S')^{\mathrm{des}}$ the category of all quasicoherent S'-modules with descent data.

The descent datum is **effective** if $(\mathcal{F}', \phi)$ is isomorphic to $(p^*\mathcal{F}, \varphi)$ for some $\mathcal{F} \in \underline{\mathsf{QCoh}}(S)$, where φ denotes the canonical isomorphism

$$p_1^*p^*\mathcal{F} \xrightarrow{\sim} p_2^*p^*\mathcal{F}.$$

We say in this case that $\mathcal{F}'$ **descends to** S.

EXAMPLE C.1.1.5. In the notation of example C.1.1.2, it is not too hard to verify that a descent datum ϕ on $\mathcal{F}' = \mathcal{O}_{M'} \in \underline{\mathsf{QCoh}}(\mathrm{Spec}(R'))$ is effective if and only if the associated diagram $M' \rightrightarrows M' \otimes_R R'$ can be enlarged into a cocartesian diagram

$$\begin{array}{ccc} M' & \longrightarrow & M' \otimes_R R' \\ \uparrow & & \uparrow \\ M & \longrightarrow & M' \end{array}$$

over the canonical $R \to R' \rightrightarrows R' \otimes_R R'$.

THEOREM C.1.1.6 (fpqc descent on quasicoherent sheaves). Let $p : S' \to S$ be a faithfully flat and quasi-compact (fpqc) morphism of schemes. Then the functor

$$\begin{aligned} p^* : \underline{\mathsf{QCoh}}(S) &\to \underline{\mathsf{QCoh}}(S')^{\mathrm{des}} \\ \mathcal{F} &\mapsto p^*\mathcal{F}, \end{aligned}$$

with the image given the canonical descent data, is an equivalence of categories.

To prove the above, it can be convenient to formulate the cocycle condition within the context of T-valued points of S'. In particular, given $(\mathcal{F}', \phi) \in \underline{\mathsf{QCoh}}(S')^{\mathrm{cov}}$ and $t_1, t_2 \in S'(T)$, denote by

$$\phi_{t_1,t_2} : t_1^*\mathcal{F}' \to t_2^*\mathcal{F}'$$

the pullback of ϕ under the morphism

$$(t_1, t_2) : T \to S'' := S' \times_S S'.$$

For a third point $t_3 \in S'(T)$, we can consider

$$(t_1, t_2, t_3) : T \to S''' := S' \times_S S' \times_S S'$$

and compose it with each of the projections down to S''. With all this in mind, we see that ϕ satisfies the cocycle condition if and only if

$$\phi_{t_1,t_3} = \phi_{t_2,t_3} \circ \phi_{t_1,t_2}$$

for all $t_1, t_2, t_3 \in S'(T)$. In particular setting all the points equal (i.e $t_1 = t_2 = t_3 = t$), we get that $\phi_{t,t} = \phi_{t,t}^2$, which implies $\phi_{t,t} = \mathrm{id}_{t^*\mathcal{F}'}$.

LEMMA C.1.1.7. Assume $p : S' \to S$ admits a section. Then any $(\mathcal{F}', \phi) \in \underline{\mathsf{QCoh}}(S')^{\mathrm{des}}$ is effective.

Proof. In particular we claim the choice of section $s : S \to S'$ determines an S-module $\mathcal{F} := s^*\mathcal{F}'$ such that $p^*\mathcal{F}$ with the canonical descent datum is isomorphic to the pair $(\mathcal{F}', \phi)$.

Consider the points $u := \mathrm{id}_{S'}$ and $v := s \circ p$ in $S'(S')$. As usual we denote $S'' := S' \times_S S'$. Then $u^*\mathcal{F}' = \mathcal{F}'$ and $v^*\mathcal{F}' = p^*\mathcal{F}$, and we can consider the isomorphism

$$f := \phi_{u,v} : \mathcal{F}' \xrightarrow{\sim} p^*\mathcal{F}.$$

It suffices to show the diagram

$$\begin{array}{ccc} p_1^*\mathcal{F}' & \xrightarrow{\phi} & p_2^*\mathcal{F}' \\ \downarrow{\scriptstyle p_1^*\circ f} & & \downarrow{\scriptstyle p_2^*\circ f} \\ p_1^*p^*\mathcal{F} & = & p_2^*p^*\mathcal{F} \end{array}$$

commutes. Consider the three S''-valued points of S' given by p_1, p_2 and

$$p_3 := s \circ p \circ p_1 = s \circ p \circ p_2.$$

Then $\phi = \phi_{p_1,p_2}$, since $(p_1, p_2) = \mathrm{id}_{S''}$. Moreover, we have by construction that

$$(u, v) \circ p_i = (p_i, p_3) : S'' \to S'' \implies p_i^* \circ f = \phi_{p_i,p_3}, \quad i = 1, 2.$$

Applying the cocycle condition for ϕ, we also have

$$\phi_{p_1,p_3} = \phi_{p_2,p_3} \circ \phi_{p_1,p_2}.$$

Substituting, into the previous equation with $i = 1$ we get

$$p_1^* \circ f = \phi_{p_2,p_3} \circ \phi_{p_1,p_2} = (p_2^* \circ f) \circ \phi,$$

thus proving commutativity of the aforementioned diagram. □

We're now ready to prove the main result.

Proof of C.1.1.6. By Proposition C.1.1.3, we know the pullback functor is fully faithful, so it remains only to show that each descent datum is effective. By Lemma C.1.1.7, we know this is true if $p : S' \to S$ admits a section. Our goal will be to reduce to this case.

Firstly, since the functor is fully faithful and descent data on $\mathcal{F}'$ can be pulled back via any morphism, we may replace the morphism $p : S' \to S$ with $p \circ u$, where $u : T \to S'$ is any faithfully flat quasicompact map. Thus via the same logic as that in the proof of Proposition C.1.1.3, we may reduce to the affine case, where $S = \mathrm{Spec}(R)$, $S' = \mathrm{Spec}(R')$ and $\mathcal{F}' = \mathcal{O}_{M'}$ for some $M' \in \underline{R'\text{-Mod}}$.

Let $\phi : M' \otimes_R R' \xrightarrow{\sim} R' \otimes_R M'$ denote our descent datum. Then ϕ determines a diagram $M \rightrightarrows M' \otimes_R R'$ over $R' \rightrightarrows R' \otimes_R R'$. As mentioned at the end of the proof of Proposition C.1.1.3, recall the diagram

$$M \to M \otimes_R R' \rightrightarrows M \otimes_R R' \otimes_R R'$$

is exact for any R-module M and faithfully flat $R \to R'$. Thus if M' does descend to R, (i.e is indeed of the form $M \otimes_R R'$), it must descend to

$$K := \ker\left(M' \rightrightarrows M' \otimes_R R'\right).$$

To show effectivity of ϕ, we need to show that the diagram

$$\begin{array}{ccc} M' & \longrightarrow & M' \otimes_R R' \\ \uparrow & & \\ K & \longrightarrow & M' \end{array}$$

is commutative and co-cartesian over $R \to R' \rightrightarrows R' \otimes_R R'$. We can apply a faithfully flat base change and therefore assume $R = R'$, and thus that $R \to R'$ admits a section. Then we can apply logic identical to that used in the proof of Lemma C.1.1.7 to deduce the desired commutative and co-cartesian properties. This completes our proof. □

§C.1.2 The Theorem of the Square

§C.1.3 Torsors

The main result we will show here is the following

> THEOREM C.1.3.1. Let $R \to R'$ be a faithfully flat extension of fields or DVRs. Let G' be an smooth, separated, finite-type R'-group scheme, and X' an R'-torsor under G. Suppose we have descent data on G' and X' (wrt $R \to R'$) compatible with the group structure on G' and the G'-action on X'. Then G' (resp. X') descends to a smooth, separated, finite-tyep R-group scheme G (resp. R-torsor under G).

Section C.2

Valuative Spaces

We frame this section around proving the following result of Raynaud and Gruson, a special case of which can be a useful technical hammer to "separate" certain non-separated schemes, as done in §8.2.1.

Theorem C.2.0.1. Let Y be a finitely-presented scheme over a quasicompact, quasiseparated scheme X, let $U \subseteq$ be a quasicompact open, and let $\mathcal{F}$ be a quasicoherent sheaf on Y of finite type. Suppose the restriction of $\mathcal{F}$ to $Y \times_X U$ is a finitely presented $\mathcal{O}_{Y \times_X U}$-module which is flat over U. Then there exists a blowup $f : X' \to X$ such that

(a) The center of the blowup of f is a finitely presented closed subscheme of X disjoint from U.

(b) The strict transform $\mathcal{F}'$ of $\mathcal{F}$ is flat over X'.

The pedagogical motivation behind proving this fact is to use a novel approach, courtesy of Guignard's [Gui19], that makes use of the theory of valuative spaces, which the authors find of interest to expound upon. Pretty much all of this section is a retelling of Guignard's paper.

§C.2.1 Constructibly Supported Rings

Definition C.2.1.1. A **family of supports** on a topological space X is a set Φ_X of closed subsets of X satisfying:

(a) $\emptyset \in \Phi_X$.

(b) Φ_X is closed under finite unions.

(c) If $Z \in \Phi_X$ then $W \in \Phi_X$ for any closed $W \subseteq Z$.

Example C.2.1.2. Given an arbitrary set of closed subsets $\{Z_i\}_{i \in I}$ of X, the set of closed subsets contained in a finite union of Z_i's is a family of supports on X. This is called the family of supports **generated** by the $\{Z_i\}_{i \in I}$.

We want to connect this geometric notion to an algebraic one via affine schemes, which motivates the following definition.

Definition C.2.1.3. Let R be a ring. A **family of constructible supports** on R is a set Φ_R of finitely generated ideals of R satisfying the following properties:

(a) Φ_R contains the unit ideal.

(b) Φ_R is closed under taking products of ideals.

(c) If $J \in \Phi_R$ then $I \in \Phi_R$ for any finitely generated $I \supseteq J$.

We call ideals $I \in \Phi$ **admissible**.

The category CSuppRing of rings with constructible supports has objects that are pairs (R, Φ_R) and morphisms $f : (R, \Phi_R) \to (S, \Phi_S)$ are ring homomorphism such that $f(I)S \in \Phi_S$ for any $I \in \Phi_R$. For ease of notation we will also refer to objects in this category as **Φ-Rings**.

The motivation behind the naming of the above definition is the following.

EXAMPLE C.2.1.4. Let R be a commutative ring and Φ_R a family of constructible supports on it. Then the set

$$\{Z \subseteq_{\text{closed}} \operatorname{Spec}(R) : Z \subseteq V(I) \text{ for some } I \in \Phi_R\}$$

is a family of supports on $\operatorname{Spec}(R)$ generated by its constructible elements.

WARNING (Abuse of Notation). We will often just drop the notation Φ_R of the family of constructible supports when talking about an object of CSuppRing, and just refer to it as R. Note of course that any such R is the underlying ring of several distinct Φ-rings, but this is no different to omitting the group operation notation when referring to a group.

PROPOSITION C.2.1.5. The category CSuppRing has small limits and colimits. Moreover the forgetful functor to Ring preserves limits and colimits, and has both a left and right adjoint.

Proof. [later] □

Let X be a topological space, and Φ a family of supports on it. Consider the functor on the category of abelian sheaves on X given by

$$\begin{aligned} \Gamma_\Phi(X, \cdot) : \underline{\operatorname{AbSh}}_X &\to \underline{\operatorname{Ab}} \\ \mathcal{F} &\mapsto \Gamma_\Phi(X, \mathcal{F}) := \{s \in \Gamma(X, \mathcal{F}) : \operatorname{Supp}(s) \in \Phi\}. \end{aligned}$$

This is a left exact additive functor, and its domain as enough injectives, thus we can define its right-derived functors, which we denote by $H^q_\Phi(X, \cdot)$. If Φ consists of the set of closed subsets of a given closed $Z \subseteq X$, we also denote this by $H^q_Z(X, \cdot)$. By using some standard cohomological arguments we can in fact obtain a Mayer-Vietoris-type result.

EXERCISE C.2.1.A. Let X be a topological space with family of supports Φ. Let Z, Z' be arbitrary closed subsets of X. Finally take any $\mathcal{F} \in \underline{\operatorname{AbSh}}_X$. Then we have:

(a) $H^q_\Phi(X, \mathcal{F}) \cong \varinjlim_{W \in \Phi} H^q_W(X, \mathcal{F})$.

(b) There exists a long exact sequence

$$\cdots \to H^q_{Z \cap Z'}(X, \mathcal{F}) \to H^q_Z(X, \mathcal{F}) \oplus H^q_{Z'}(X, \mathcal{F}) \to H^q_{Z \cup Z'}(X, \mathcal{F}) \to H^{q+1}_{Z \cap Z'}(X, \mathcal{F}) \to \cdots$$

(c) For $U := X - Z$, there exists a long exact sequence

$$\cdots \to H^q_{Z \cap Z'}(X, \mathcal{F}) \to H^q_{Z'}(X, \mathcal{F}) \to H^q_{Z' \cap U}(U, \mathcal{F}_U) \to H^{q+1}_{Z \cap Z'}(X, \mathcal{F}) \to \cdots$$

[Hint: for (a), take an injective resolution of $\mathcal{F}$, and for (b) and (c), take a flasque resolution of $\mathcal{F}$.]

DEFINITION C.2.1.6. Let R be a Φ-ring, and we also denote by Φ_R the corresponding family of supports on $\mathrm{Spec}(R)$ as described in Example C.2.1.4. Let M be an R-module corresponding to the quasicoherent sheaf $\widetilde{M}$. We then define the **Φ-cohomology groups** to be

$$H^q_I(M) := H^q_{V(I)}(X, \widetilde{M}), \quad H^q_{\Phi_R}(M) := H^q_{\Phi_R}(X, \widetilde{M}).$$

We say M is **d-deep** if $H^q_{\Phi_R}(M)$ vanishes for all $q < d$. We denote by $\underline{\mathrm{CSuppRing}}^{\geq d}$ the full subcategory of Φ-rings which are d-deep as modules over themselves.

PROPOSITION C.2.1.7. Let (R, Φ_R) be a Φ-ring, and let Φ_0 be a generating set of ideals for Φ_R. Then an R-module M is d-deep if and only if $H^q_I(M) = 0$ for all $q < d$ and $I \in \Phi_0$.

Proof. For one direction, suppose M is d-deep. We proceed by induction on d with the case $d = 0$ is trivial. Thus assume the result holds for all $(d-1)$-deep modules. Using part (a) of Exercise C.2.1.A, we have the isomorphism

$$0 = H^d_{\Phi_R}(M) \cong \varinjlim_{I \in \Phi_R} H^d_I(M).$$

If we show all the morphisms between objects in the colimit diagram are monomorphisms, it will follow they must all be 0 in order for their colimit to be 0 as well.

Thus suppose we have an inclusion of admissible ideals $I \subseteq J \trianglelefteq R$. By inducting on the number of generators we may WLOG assume $J = \langle I, g \rangle$ for some $g \in A$. Then if we set $Z := V(g), Z' := V(I)$ and $U := D(g)$, part (c) of Exercise C.2.1.A gives us the long exact sequence

$$\cdots \to H^q_J(M) \to H^q_I(M) \to H^q_I(M[g^{-1}]) \to H^{q+1}_J(M) \to \cdots.$$

By the inductive hypothesis on $(d-1)$-deepness of M, we know

$$H^{d-1}_I(M[g^{-1}]) \cong (H^{d-1}_I(M))[g^{-1}] = 0,$$

where the first isomorphism follows from a standard cohomology and base change result (see [Stacks, § 30.7]. The long-exact sequence then tells us that $H^d_J(M) \to H^d_I(M)$ is injective, which was what we wanted. This completes the proof in one direction.

For the other direction, suppose $H^q_I(M) = 0$ for any $I \in \Phi_0$ and any $q < d$. For any $J \in \Phi_0$, part (b) of Exercise C.2.1.A gives us an exact sequence

$$0 \to H^q_{IJ}(M) \to H^{q+1}_{I+J}(M) \to H^{q+1}_I(M) \oplus H^{q+1}_J(M)$$

for any $q < d$. We know from our proof of the prior direction that the rightmost map above is an injection, and thus it follows $H^q_{IJ}(M) = 0$ as well. Thus $H^q_I(M) = 0$ for any finite product of elements of Φ_0 and for any $q < d$. Moreover by definition any $I \in \Phi_R$ must contain an ideal J which is a finite product of elements of Φ_0, and again our reasoning from the prior direction tells us

$$H^q_J(M) = 0 \implies H^q_I(M) = 0.$$

It follows by definition then that the q-th Φ-cohomology vanishes for every admissible ideal, and thus by definition M is d-deep. This completes the proof of the other direction. □

There is a notion of purified and closed modules over noetherian rings. We describe the constructibly supported analogues of those.

DEFINITION C.2.1.8. Let R be a Φ-ring and M an R-module. The **purification** of M is the quotient

$$M^{pur} := M/H^0_{\Phi_R}(M).$$

If M is 1-deep, the **closure** of M is the colimit

$$M^{\triangleleft} := \varinjlim_{I \in \Phi_R} \mathrm{Hom}_R(I, M).$$

If M is not 1-deep we define $M^{\triangleleft} := (M^{pur})^{\triangleleft}$.

§C.2.2 Sheaves of Φ-Rings

DEFINITION C.2.2.1. Let $\mathcal{C}$ be a category. A **presheaf** of Φ-rings (or constructibly supported rings) on $\mathcal{C}$ is a contravariant functor $A : \mathcal{C} \to \underline{\mathrm{CSuppRing}}$. We say A is a **sheaf** of Φ-rings if the underlying presheaf of sets is a sheaf and for each integer r the contravariant functor

$$\mathrm{Adm}_{A,r} : U \mapsto \{(f_1, \ldots, f_r) \in A(U)^r : \langle f_1, \ldots, f_r \rangle \text{ is admissible in } A(U)\}$$

is a sheaf of sets on $\mathcal{C}$.

PROPOSITION C.2.2.2. Let A be a presheaf of Φ-rings on $\mathcal{C}$ and let $\widetilde{A}$ be the sheaf of rings associated to the underlying presheaf of rings. Then there is a unique sheaf of Φ-rings with underlying sheaf of rings $\widetilde{A}$ such that for every r the sheaf $\mathrm{Adm}_{\widetilde{A},r}$ is the sheaf associated to $\mathrm{Adm}_{A,r}$. Moreover the morphism $A \to \widetilde{A}$ is initial amongst the morphisms of presheaves of Φ-rings from A to sheaves of Φ-rings.

Proof. [later] □

COROLLARY C.2.2.3.

DEFINITION C.2.2.4. Let X be a topological space. We define a **Φ-ring of X** to be a sheaf of rings on X together with a subsheaf $\mathrm{Adm}_r(A)$ of A^r, for each positive integer r, such that

(i) For all open $U \subseteq X$ and for all positive r, s, for all $(f_1, \ldots, f_r) \in \mathrm{Adm}_{A,r}(U)$ and $(g_1, \ldots, g_s) \in \mathrm{Adm}_{A,s}(U)$, we have

$$(f_i g_j)_{i \in [r], j \in [s]} \in \mathrm{Adm}_{rs,A}(U) \subseteq A(U)^{rs}.$$

(ii) For all open $U \subseteq X$ and all $(f_1, \dots, f_r) \in \mathrm{Adm}_{A,r}(U)$, we have

$$\sum_{i=1}^{r} f_i A(U) \subseteq \sum_{j=1}^{s} g_j A(U) \implies (g_1, \dots, g_s) \in \mathrm{Adm}_{A,s}(U).$$

Section C.3

Some Useful Algebraic Number Theory

Proposition C.3.0.1 (Hermite). Let K be a number field, $d \geq 0$, and $S \subset \mathscr{P}_K$ a finite set of places. Then there are only a finite number of extensions of K of degree less than d and unramified outside of S.

Proof. finish later. Proof in C.3 of HS □

§C.3.1 Ramification

§C.3.2 Kummer Theory

References and Acknowledgements

References

[BL04a] C. Birkenhake and H. Lange. *Complex Abelian Varieties*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2004. ISBN: 9783540204886. URL: https://books.google.com/books?id=MOW2gEP7HIkC.

[BL04b] Christina Birkenhake and Herbert Lange. "Endomorphisms of Abelian Varieties". In: *Complex Abelian Varieties*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 113–144. ISBN: 978-3-662-06307-1. DOI: 10.1007/978-3-662-06307-1_7. URL: https://doi.org/10.1007/978-3-662-06307-1_7.

[BLR90a] S. Bosch, W. Lutkebohmert, and M. Raynaud. *Neron Models*. Springer, 1990. ISBN: 9783642514395.

[BLR90b] Siegfried Bosch, Werner Lütkebohmert, and Michel Raynaud. "Descent". In: *Néron Models*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1990, pp. 129–171. ISBN: 978-3-642-51438-8. DOI: 10.1007/978-3-642-51438-8_7. URL: https://doi.org/10.1007/978-3-642-51438-8_7.

[CK23] Alexandru Chirvasitu and Ryo Kanda. "Projective discrete modules over profinite groups". In: *Journal of Pure and Applied Algebra* 227.4 (2023). DOI: https://doi.org/10.1016/j.jpaa.2022.107260.

[Con15] Brian Conrad. *Abelian Varieties*. https://www.ams.org/open-math-notes/omn-view-listing?listingId=110657. Stanford University, 2015.

[CS86] Gary Cornell and Joseph H. Silverman, eds. *Arithmetic Geometry*. New York, NY: Springer, 1986. ISBN: 978-1-4613-8657-5. DOI: 10.1007/978-1-4613-8655-1.

[EGA] A. Grothendieck. "Eléments de Géométrie Algébrique". In: *Publ. Math. IHES* (4, 8, 11, 17, 20, 24, 28, 32 1960-67).

[Fan05] B. Fantechi. *Fundamental Algebraic Geometry: Grothendieck's FGA Explained*. Mathematical surveys and monographs. American Mathematical Society, 2005. ISBN: 9780821842454. URL: https://books.google.com/books?id=UNxU24JAHDoC.

[GR71] Laurent Gruson and Michel Raynaud. "Critères de platitude et de projectivité". In: *Invent. math.* 13 (1971), pp. 1–89.

[Gro57] Alexander Grothendieck. “Sur quelques points d’algèbre homologique, I”. In: *Tohoku Mathematical Journal* 9.2 (1957), pp. 119–221. DOI: 10.2748/tmj/1178244839. URL: https://doi.org/10.2748/tmj/1178244839.

[Gui19] Quentin Guignard. *A new proof of Raynaud-Gruson’s flattening theorem*. 2019. arXiv: 1801.01188 [math.AG]. URL: https://arxiv.org/abs/1801.01188.

[GW23] U. Görtz and T. Wedhorn. *Algebraic Geometry II: Cohomology of Schemes*. 1st. Springer Spektrum Wiesbaden, 2023. DOI: https://doi.org/10.1007/978-3-658-43031-3.

[Har77] R Hartshorne. *Algebraic Geometry*. 1st. Graduate Texts in Mathematics. Springer New York, NY, 1977. DOI: https://doi.org/10.1007/978-1-4757-3849-0.

[HS00] Marc Hindry and Joseph H. Silverman. *Diophantine Geometry: An Introduction*. New York, NY: Springer New York, 2000. ISBN: 978-1-4612-1210-2. DOI: 10.1007/978-1-4612-1210-2. URL: https://doi.org/10.1007/978-1-4612-1210-2.

[Hür92] Werner Hürlimann. “A short proof of the Albert-Brauer-Hasse-Noether theorem”. en. In: *Journées arithmétiques de Genève - 9-13 septembre 1991*. Ed. by Coray D. F. and Pétermann Y.-F. S. Astérisque 209. Société mathématique de France, 1992, pp. 215–220. URL: https://www.numdam.org/item/AST_1992__209__215_0/.

[Mat70] H. Matsumura. *Commutative Algebra*. Benjamin/Cummings Series in the Life Sciences. W. A. Benjamin, 1970. ISBN: 9780805370256.

[Mat87] H. Matsumura. *Commutative Ring Theory*. Ed. by MilesTranslator Reid. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1987.

[Mcc03] Kevin Mccrimmon. “A Taste of Jordan Algebras”. In: (Jan. 2003). DOI: 10.1007/b97489.

[Mil08] James S. Milne. *Abelian Varieties (v2.00)*. Available at www.jmilne.org/math/. 2008.

[Mil20] J.S. Milne. *Class Field Theory (v4.03)*. Available at www.jmilne.org/math/. 2020.

[Mil86] J. S. Milne. “Abelian Varieties”. In: *Arithmetic Geometry*. Ed. by Gary Cornell and Joseph H. Silverman. New York, NY: Springer New York, 1986, pp. 103–150. ISBN: 978-1-4613-8655-1. DOI: 10.1007/978-1-4613-8655-1_5. URL: https://doi.org/10.1007/978-1-4613-8655-1_5.

[Mum74] D. Mumford. *Abelian Varieties*. 2nd. Studies in mathematics. Oxford University Press for the Tata Institute of Fundamental Research, 1974.

[Pie82] Richard S. Pierce. “Simple Algebras”. In: *Associative Algebras*. New York, NY: Springer New York, 1982, pp. 218–233. ISBN: 978-1-4757-0163-0. DOI: 10.1007/978-1-4757-0163-0_12. URL: https://doi.org/10.1007/978-1-4757-0163-0_12.

[SGA3] Michael Artin et al. *Schémas en groupes*. Séminaire de Géométrie Algébrique de l’Institut des Hautes Études Scientifiques. Paris: Institut des Hautes Études Scientifiques, 1963/1966.

[SGA6] Pierre Berthelot, Alexander Grothendieck, and Luc Illusie. *Théorie des Intersections et Théorème de Riemann-Roch*. Vol. 225. Lecture notes in mathematics. Springer-Verlag, 1971.

[Stacks] The Stacks Project Authors. *Stacks Project*. https://stacks.math.columbia.edu. 2018.

[TO70] John Tate and Frans Oort. “Group schemes of prime order”. eng. In: *Annales scientifiques de l’École Normale Supérieure* 3.1 (1970), pp. 1–21. URL: http://eudml.org/doc/81855.