Néron models, Tamagawa factors, and Tate-Shafarevich groups

Brian Conrad *

October 14, 2015

1 Motivation

Let R be a discrete valuation ring, $F = \mathrm{Frac}(R)$, and k its residue field. Let A be an abelian variety over F. There are two questions we can ask ourselves:

(1) Does A extend to a smooth proper R-scheme?

(2) Does A extend even to an *abelian scheme* over R? (An abelian scheme over a base S is a smooth proper S-group with connected (geometric) fibers. The commutativity is automatic, but non-trivial; it rests on deformation-theoretic arguments given early in Chapter 6 of Mumford's GIT book.)

Although the second question looks like a stronger request than the first, the theory of Néron models will imply that they are equivalent.

Obstruction to (1). If the answer to (1) is affirmative, with smooth proper R-model X, then the smooth and proper base change theorems for étale cohomology imply upon choosing a place of F_s over that on F and letting R^{sh} denote the associated strict henselization inside F_s (so R^{sh} has residue field k_s that is a separable closure of k); it is the "maximal unramified extension" of the henselization R^{h} that the natural map $H^i_{\text{ét}}(X_{k_s}, \mathbf{Q}_\ell) \to H^i_{\text{ét}}(A_{F_s}, \mathbf{Q}_\ell)$ determined by $R^{\mathrm{sh}} \subset F_s$ is an isomorphism, where ℓ is a prime not equal to $\mathrm{char}(k)$.

This map is equivariant relative to the action of the decomposition group D in $\mathrm{Gal}(F_s/F)$ attached to the chosen place on F_s (via $D \twoheadrightarrow \mathrm{Gal}(k_s/k)$ as usual), so a consequence of the existence of such an X is that $H^i(A_{F_s}, \mathbf{Q}_\ell)$ is unramified as a $\mathrm{Gal}(F_s/F)$-module. Hence, *ramifiedness* of such cohomology on the geometric generic fiber is an obstruction to (1). This has nothing to do with abelian varieties.

*Notes taken by Tony Feng

Obstruction to (2). Since $H^1_{\text{ét}}(A_{F_s}, \mathbf{Q}_\ell) = V_\ell(A)^*$, we can alternatively think in terms of the Tate module. For any abelian scheme $\mathcal{A}/R$ and prime $\ell \neq \text{char}(k)$, the R-group scheme $\mathcal{A}[\ell^n]$ is finite étale and hence the Galois module of F_s-points of $\mathcal{A}[\ell^n]_F \simeq A[\ell^n]$ must be unramified. Hence, $V_\ell(A)$ is unramified in such cases, so ramifiedness of $V_\ell(A)$ as a $\text{Gal}(F_s/F)$-module is an obstruction to (2) (more elementary than the obstruction to (1) since it doesn't involve étale cohomology, though it entails a stronger hypothesis than in (1)!).

Miracle. The fundamental *Néron-Ogg-Shafarevich criterion* asserts that these are the *only* obstructions, and so (1) $\iff$ (2). The proof uses Néron models, as we will indicate later. (Of course, the deeper part is that an unramifiedness Galois hypothesis implies a structural property for a scheme.)

One of Néron's key insights was that if you forget about properness and focus on *smoothness* then there is a "best" integral model in general. (Note the contrast with the theory of curves, for which one focuses on proper flat models with possibly non-smooth fibers.) While understanding the construction of Néron models can be psychologically comforting, and the techniques involved in it are very useful for other purposes, in practice one typically only needs to know the general properties and existence of Néron models and not the details of their construction. The main reference for this lecture is the amazing book *Néron Models* by Bosch, Lütkebohmert, and Raynaud, hereafter denoted [BLR].

Here is the main existence result:

Theorem 1.1 (Néron). *Let S be a Dedekind scheme (i.e., a connected normal noetherian scheme of dimension* 1*), F its function field, and A an abelian variety over F. There exists a smooth separated finite type S-scheme $\mathcal{A}$ with generic fiber A such that for all smooth $T \to S$, any map $T_F \to A$ over F uniquely extends to a map $T \to \mathcal{A}$ over S:*

$$\begin{tikzcd} T_F \arrow[r] \arrow[d] & A \arrow[d] \\ T \arrow[r, dotted, "\exists!"'] & \mathcal{A} \end{tikzcd}$$

Example 1.2. If $T_F \to A$ is an isomorphism, then T dominates $\mathcal{A}$. Therefore, any smooth model of A over S dominates the Néron model (although it is non-trivial that any smooth models exist at all!). This alone implies that Néron models are functorially unique if they exist, and it is trivial from the mapping property that the formation of Néron models is compatible with direct products. However, beyond that they tend to have weak functorial properties (such as relative to ramified base change).

Remark 1.3. Let $j : \text{Spec}(F) \to S$ be the canonical map. For any functor H on the category of F-schemes, there is an associated "pushforward" functor $j_*(H)$ on the category of S-schemes defined by

$$j_*(H)(S' \to S) = H(S'_F).$$

Restricting attention to smooth S-schemes and smooth F-schemes, the mapping property of Néron models says exactly that $j_*(A)$ on the category of *smooth* S-schemes is represented by $\mathcal{A}$ (reproving functorial uniqueness, by Yoneda's Lemma), but the existence result gives *more*: $\mathcal{A}$ is separated and of finite type over S. Neither of these properties is a formal consequence of the mapping property (and there is a broader theory of Néron models that includes some tori, and there typically the Néron model is merely locally of finite type).

Example 1.4. Taking $T = S$, the theorem implies that $\mathcal{A}(S) = A(F)$, which looks like what one would get from the valuative criterion of properness, but we are only evaluating on S, not an arbitrary Dedekind scheme over S! The same applies for any quasi-compact étale morphism $S' \to S$, so $\mathcal{A}(S') = A(F')$ where $\mathrm{Spec}\,(F') = S'_F$.

Example 1.5. If $S = \mathrm{Spec}\ \mathbf{Z}$ and $K/\mathbf{Q}$ is a finite extension, then $A(K) = \mathcal{A}(U)$ where $U \subset \mathrm{Spec}\,(\mathcal{O}_K)$ is the maximal open subscheme unramified over S.

Remark 1.6. There is no analogue of the Néron model for general algebraic groups. Indeed, it is a general fact (see [1.3/1, BLR]) that a smooth group scheme G of finite type over the fraction field F of a discrete valuation ring R admits a Néron model over R (i.e., a smooth R-model that is separated, of finite type, and satisfies the Néron mapping property) if and only if $G(F^{\mathrm{sh}})$ is "bounded" in G, where F^{sh} is the fraction field of a strict henselization of R and "boundedness" is defined in terms of coordinates in suitable affine charts; see [1.1, BLR] for a detailed discussion of boundedness for general separated F-schemes of finite type.

Such boundedness never holds for nontrivial connected semisimple groups when R is complete with perfect residue field, as such groups always become quasi-split over F^{sh} by a theorem of Steinberg. There is a rich theory of "good" smooth affine models of connected semisimple groups, due to Bruhat and Tits, but it has nothing to do with Néron models (though was likely inspired by it); the resulting "Bruhat–Tits theory" is important in the representation theory of p-adic groups (providing algebro-geometric models for certain compact open subgroups of the group of rational points of a connected semisimple group over a local field).

2 Properties and examples

Properties of the Néron model.

1. It has an S-group structure extending from A. (This is immediate from the mapping property, building on the compatibility with direct products.)

2. It is compatible with étale base change $S' \to S$ and likewise localization at closed points $s \in S$, *and* base change to $\widehat{\mathcal{O}}_{S,s}$ as well as $\mathcal{O}^{\mathrm{h}}_{S,s}$ and $\mathcal{O}^{\mathrm{sh}}_{S,s}$.

3. If a closed fiber $\mathcal{A}_s$ is $k(s)$-proper then $\mathcal{A}_{\mathcal{O}_{S,s}}$ is necessarily $\mathcal{O}_{S,s}$-proper (this is not at all obvious; see [IV$_3$, 15.7.10, EGA]) with connected special fiber (by considerations with Stein factorization), so in such cases $\mathcal{A}$ is an abelian scheme over some open neighborhood of s in S.

Example 2.1. If $\mathcal{A} \to S$ is an abelian scheme, then it is the Néron model of its generic fiber. This rests on the valuative criterion of properness and the Weil Extension Lemma [4.4, BLR].

To be more specific, given a map $T_F \to A$, we can extend it across the generic points of the special fiber by the valuative criterion (as the local ring of T at such a point is a discrete valuation ring). Hence, the problem is to extend to T a map $U \to \mathcal{A}$ defined on an open subscheme $U \subset T$ whose complement is everywhere of codimension at least 2. Weil used translations in an artful way to make an extension to all of T. (Weil's proof was over fields, as he needed it in his work on Jacobian varieties; the version over R is technically more involved but rests on similar ideas.)

Example 2.2. Let $S = \mathrm{Spec}\,(R)$ where R is a discrete valuation ring and E is an elliptic curve with "split multiplicative reduction" over the function field F. The Néron model $\mathcal{E}$ of E has special fiber $\mathcal{E}_k \simeq \mathbf{G}_m \times (\mathbf{Z}/n\mathbf{Z})$, where $n = -\operatorname{ord}_R(j(E))$.

If $R \to R'$ has ramification degree e, then the Néron model $\mathcal{E}'$ of $E_{F'}$ satisfies $\mathcal{E}'_{k'} \simeq \mathbf{G}_m \times (\mathbf{Z}/ne\mathbf{Z})$ since the j-invariant doesn't change. So if $e > 1$ then the Néron model "grows". It is not obvious from this description, but these two identity components are really the "same" $\mathbf{G}_m$; i.e., the natural base change morphism $\mathcal{E}_{R'} \to \mathcal{E}'$ arising from the Néron mapping property and identification of generic fibers restricts to an *isomorphism* between identity components of special fibers.

Example 2.3. Let $X \to \mathrm{Spec}\ F$ be a smooth proper geometrically connected curve of genus > 0. Lipman's work on resolution of singularities for 2-dimensional excellent schemes ensures that there exists a proper (flat) S-model (i.e., a proper flat S-scheme with generic fiber X) which is regular (a trick of Hironaka with completions of local rings on S reduces the task to the case of excellent S, so Lipman's work is applicable). Further work in the theory of fibered surfaces ensures (since the genus is positive) that among these there is one such model $\mathcal{X} \to S$ that it is dominated by all others; it is called the *minimal regular proper model.* Since X is F-smooth, all but finitely many fibers $\mathcal{X}_s$ are smooth, and all fibers are geometrically connected since $\mathcal{X} \to S$ is its own Stein factorization.

If $J = \mathrm{Jac}(X)$ then one could ask: how is $\mathcal{X}$ is related to the Néron model $N(J)$? In general, for a smooth S-group G of finite type with (geometrically) connected generic fiber (so all but finitely many fibers are geometrically connected) let us denote by G^0 the open subscheme given by removing the closed complement of the identity component in the finitely many disconnected fibers; this is called the *relative identity component.* (The scheme G itself is connected.) Then in turns out that $N(J)^0$ has a direct description in terms of $\mathcal{X}$ provided that for each of the non-smooth fibers $\mathcal{X}_s$ the gcd of geometric multiplicities of its irreducible components is 1 (e.g., this holds whenever $\mathcal{X}_s$ has non-empty smooth locus): in such cases

$$N(J)^0 \simeq \mathrm{Pic}^0_{\mathcal{X}/S}.$$

Here, $\mathrm{Pic}_{\mathcal{X}/S}$ is the relative Picard functor for $\mathcal{X}$ over S, and it is a deep result of Raynaud that under the above hypotheses on $\mathcal{X}$ this is an algebraic space (usually non-separated when there some geometric fibers $\mathcal{X}_{\overline{s}}$ are not integral) and that its open subspace $\mathrm{Pic}^0_{\mathcal{X}/S}$ which is the identity component on every fiber is a *separated scheme* of finite type. Moreover, the component groups of the fibers of $N(J)$ are described completely by intersection theory on $\mathcal{X}/S$. These matters are discussed at length in [9.5–9.6, BLR].

The significance is that the fibers of $\mathrm{Pic}^0_{\mathcal{X}/S}$ are the Picard schemes $\mathrm{Pic}^0_{\mathcal{X}_s/k(s)}$ that have moduli-theoretic meaning in terms of the geometry of $\mathcal{X}_s$, so this makes it possible in such cases to read off information about the geometry of $N(J)^0_s$. Hence, for Jacobians we can use the fibered surface $\mathcal{X} \to \mathcal{S}$ to understand the Néron model. In the case of genus 1 with $X(F) \neq \emptyset$ (i.e., X an elliptic curve, so $J = X$), $N(J)$ coincides with the maximal smooth open subscheme $\mathcal{X}^{\mathrm{sm}}$ of $\mathcal{X}$.

Note that whenever there exists a section of $X \to \mathrm{Spec}\, F$ then it extends to a section of $\mathcal{X} \to S$ (valuative criterion), and $\mathcal{X}$ is smooth along this section (since any regular S-scheme of finite type equipped with a section is smooth along the section); hence, a-priori $\mathcal{X}^{\mathrm{sm}}$ meets every fiber $\mathcal{X}_s$ in such cases. (When there is no section then there could be fibers that are nowhere smooth.)

Definition 2.4. The *component group* of $\mathcal{A}$ at s is denoted $\Phi_s := \mathcal{A}_s/\mathcal{A}^0_s$.

This is a finite étale $k(s)$-group, so it can be "viewed" as the abelian group $\Phi_s(k(s)_{\mathrm{sep}}) = \mathcal{A}(k(s)_{\mathrm{sep}})/\mathcal{A}^0(k(s)_{\mathrm{sep}})$ equipped with its natural $\mathrm{Gal}(k(s)_{\mathrm{sep}}/k(s))$-action. Beware that in general the map $\mathcal{A}_s(k(s)) \to \Phi_s(k(s))$ is *not* surjective (there is an obstruction in the Galois cohomology group $\mathrm{H}^1(k(s), \mathcal{A}^0_s)$); for finite $k(s)$ this problem will not arise.

Example 2.5. Consider the elliptic curve 57C2 in Cremona's tables:

$$E\colon y^2 + y = x^3 + x^2 - 4390x - 113432.$$

Then $E(\mathbf{Q}) = \{0\}$, $j(E) = -(2^{12}\cdot 13171)/(3^2\cdot 19^5)$, E has good reduction away from 3 and 19, $\Phi_3 \simeq \mathbf{Z}/2\mathbf{Z}$, and Φ_{19} corresponds to the Galois module $\mathbf{Z}/5\mathbf{Z}$ equipped with a non-trivial $G_{\mathbf{F}_{19}}$-action (given in fact by the unique quadratic Galois character since the reduction type at 19 is multiplicative but non-split), so $\Phi_{19}(\mathbf{F}_{19}) = 0$.

Letting $\mathcal{E}$ denote the Néron model over $\mathbf{Z}$, the natural map

$$E(\mathbf{Q}) = \mathcal{E}(\mathbf{Z}) \to \bigoplus_p \Phi_p(\mathbf{F}_p) = \Phi_3(\mathbf{F}_3) \times \Phi_{19}(\mathbf{F}_{19}) = \Phi_3(\mathbf{F}_3)$$

is *not* surjective. The significance of this failure of surjectivity (for relating the Néron model to the Tate-Shafarevich group) will be addressed later.

3 Semistable reduction

We have noted above (with reference to EGA) the hard fact that if $\mathcal{A}_s$ is proper then $\mathcal{A}_{\mathcal{O}_s}$ is an abelian scheme (especially that it is proper over $\mathcal{O}_s$). Hence, the set of $s \in S$ for which $\mathcal{A}_s$ is proper coincides with those around which the Néron model is an abelian scheme. In particular, if the fiber at some s is not an abelian variety then it must be non-proper; we want to introduce a class of possibilities which are still reasonably nice despite the loss of properness of the fiber.

One sense in which a smooth commutative algebraic group could be considered to be "nice" is if we can probe its structure using torsion away from the characteristic (providing a certain degree of rigidity, as such torsion is étale). Abelian varieties can be probed in this way, as can tori. On the other hand, unipotent smooth connected commutative k-groups cannot: they have no nontrivial torsion away from $\mathrm{char}\,(k)$ and too much p-power torsion when $\mathrm{char}\,(k) = p > 0$. There are many other reasons why unipotent groups are worse than tori and abelian varieties (e.g., for deformation theory, representability of automorphism functors, etc.). This brings us to:

Definition 3.1. A *semi-abelian variety* over a field k is a smooth connected (commutative) k-group G such that there is an exact sequence

$$1 \to T \to G \to B \to 1$$

with T a torus and B an abelian variety. (Such an extension structure is unique if it exists, since there are no nontrivial homomorphisms from a torus to an abelian variety.)

Exercise 3.2. Use étale torsion to show that if G is a smooth connected k-group and $G_{\overline{k}}$ is semi-abelian then so is G. This is immediate via Galois descent if k is perfect, but requires some thought more generally.

Why might one care about imperfect ground fields, even if only interested in number theory in characteristic 0? Well, at the generic points of special fibers of arithmetic surfaces over $\mathbf{Z}$ the local rings are discrete valuation rings whose residue fields are global function fields over finite fields (function fields of irreducible components of mod-p fibers), and those are never perfect!

Definition 3.3. We say that A has *semistable reduction* at a closed point $s \in S$ if $\mathcal{A}_s^0$ is semi-abelian.

Example 3.4. The following fact [9.2/8, BLR] explains the terminology "semistable": if $\mathcal{X} \to S$ is a proper flat S-curve with $\mathcal{X}_F$ smooth and geometrically connected and each of the finitely many non-smooth fibers $\mathcal{X}_s$ is semistable (i.e. geometrically reduced with every geometric singularity a node) then $\mathrm{Pic}^0_{\mathcal{X}/S}$ has semistable reduction at all $s \in S$. The maximal torus in the fiber $\mathrm{Pic}^0_{\mathcal{X}_s/k(s)}$ at such s has geometric character group controlled by the reduction graph of the geometric fiber at s.

The significance of semistable reduction in the general theory is due to:

Theorem 3.5 (Grothendieck)**.** *For $S = \mathrm{Spec}\, R$ with R a discrete valuation ring there exists a finite separable extension F'/F such that $A_{F'}$ has semistable reduction at all closed points of* $\mathrm{Spec}\, R'$*. Explicitly, we can take F' to be the splitting field of $A[\ell]$ for any prime $\ell \neq \mathrm{char}\,(k)$, where when $\ell = 2$ we really use $A[4]$ instead.*

Proof. See the appendix to Expose I in SGA 7 for a beautiful proof due to Deligne using monodromy, Néron smoothening, and the Riemann Hypothesis for abelian varieties over finite fields. □

It follows that we always acquite semistable reduction everywhere by splitting the 15-torsion. Note in particular that if the dimension $g > 0$ of A is fixed then $[F' : F]$ can thereby be chosen to divide the number $\#\mathrm{GL}_{2g}(\mathbf{Z}/15\mathbf{Z})$ that depends only on g and otherwise not on A at all. Such uniform control on the degree of such an extension for attaining everywhere semistable reduction is used crucially in Faltings' proof of the Mordell Conjecture (to permit reducing all work to the study of everywhere-semistable abelian varieties of a fixed dimension over a number field).

Definition 3.6. A *semi-abelian scheme* $\mathcal{G} \to S$ is a commutative smooth separated S-group of finite type with semi-abelian fibers.

Theorem 3.7 (7.4, BLR)**.** *If an abelian variety A over F extends to a semi-abelian scheme $\mathcal{A}$ over S then the natural map $\mathcal{A} \to N(A_F)$ is an isomorphism onto $N(A_F)^0$.*

An interesting consequence is that if one has semistable reduction at all closed points of S then for any finite separable extension F'/F with associated finite (typically non-étale!) normalization $S' \to S$, the natural "base change morphism"

$$N(A)_{S'} \to N(A_{F'})$$

is an isomorphism between relative identity components (because $N(A)^0_{S'}$ is a semi-abelian scheme over S' with generic fiber $A_{F'}$). This explains the precise sense in which, once everywhere-semistable reduction is achieved, after any further finite separable extension on F all change in the Néron model is concentrated in the component groups of the non-proper fibers.

Remark 3.8. For the notions of "good reduction" (i.e., proper fiber) and "semistable reduction" (i.e., semi-abelian identity component for the fiber), how does one work with them in practice for a given abelian variety A over F? For instance, it is not obvious from the definitions that these notions should be invariant under isogeny. The key to such invariance is that each of these conditions on a fiber is equivalent to a Galois-theoretic condition for the inertial action on $V_\ell(A)$ for a prime $\ell \neq \mathrm{char}\,(k)$.

For "good reduction" there is the Néron-Ogg-Shafarevich criterion (equivalence to unramifiedness of $V_\ell(A)$), the proof of which treats the existence of $N(A)$ as a black box. Grothendieck *used* the semistable reduction theorem (and the black-box existence of $N(A)$) to show that semistable reduction is characterized by unipotence of the inertial action on $V_\ell(A)$; see [Exp IX, §3, SGA7].

Let us indicate the *key link* between inertial action and the structure of the special fiber. Suppose R is a henselian (e.g., complete) discrete valuation ring, and let R^{sh} denote its strict henselization; i.e., the valuation ring of the maximal unramified extension F^{sh}/F inside F_s. For a prime $\ell \neq \mathrm{char}\,(k)$ consider the unramified ℓ^n-torsion points $A[\ell^n](F_s)^I$ (where I is the inertia subgroup of the Galois group of F). This is $A[\ell^n](F^{\mathrm{sh}})$, which by the property of the Néron model coincides with $\mathcal{A}[\ell^n](R^{\mathrm{sh}})$. But $\mathcal{A}[\ell^n]$ a quasi-finite separated étale R-scheme, so by Zariski's Main Theorem (or more specifically its application to the structure of quasi-finite separated schemes over henselian local rings) the reduction map

$$\mathcal{A}[\ell^n](R^{\mathrm{sh}}) \to \mathcal{A}_k[\ell^n](k_s)$$

is bijective.

This gives control over the ℓ-power torsion in the special fiber, which tells us about the structure of $\mathcal{A}_k^0$ if we know a general structure theorem for *general* smooth connected commutative k-groups (e.g., we seek a way to show that if the size of the ℓ^n-torsion grows on the order of ℓ^{2n} then the group must be an abelian variety). Such a structure theorem was proved by Chevalley (who also went beyond the commutative case, but we will not discuss that here): the Chevalley structure theorem for *commutative* smooth connected group over a perfect field k says that any such G can be presented (necessarily uniquely!) as an extension

$$1 \to T \times U \to G \to B \to 1 \tag{3.1}$$

where B is an abelian variety, T is a torus (i.e., $T_{\overline{k}} \simeq \mathbf{G}_m^N$), and U is a smooth connected unipotent k-group (i.e., has a filtration over $\overline{k}$ with successive quotients isomorphic to $\mathbf{G}_a$).

The main work in the proof of Chevalley's result is over $\overline{k}$, and when k is perfect we can bring the result down to k via Galois descent. A bonus of perfect fields k is that smooth connected unipotent k-groups admit a composition series over k with successive quotients isomorphic to $\mathbf{G}_a$ when k is perfect. Perfectness is crucial: over *every* imperfect k with characteristic $p > 0$ there are 1-dimensional smooth connected unipotent groups not isomorphic to $\mathbf{G}_a$ (e.g., $y^p = x - ax^p$ where $a \in k - k^p$), and Chevalley's structure theorem is also false over such k (counterexamples are given by the *non-proper* Weil restriction $\mathrm{R}_{k'/k}(A')$ for any nontrivial purely inseparable finite extension k'/k and any nonzero abelian variety A' over k').

Let's apply the snake lemma for $[\ell^n]$ on (3.1), with $\ell \neq \mathrm{char}\,(k)$. Multiplication by ℓ is an automorphism on U (look at the filration by $\mathbf{G}_a$'s over $\overline{k}$), and on T and B it is surjective with finite kernel. Therefore, the snake lemma gives the exact sequence of finite étale k-groups

$$1 \to T[\ell^n] \to G[\ell^n] \to B[\ell^n] \to 1.$$

Taking the inverse limit on k_s-points and tensoring with $\mathbf{Q}_\ell$ gives an exact sequence

$$0 \to V_\ell(T) \to V_\ell(G) \to V_\ell(B) \to 0. \tag{3.2}$$

Thus the unipotent part has dropped away!

We apply the preceding to study an abelian variety A over a field $F = \mathrm{Frac}(R)$ where R is a discrete valuation ring with *perfect residue field* k (later to be finite). We take $G = \mathcal{A}_k^0$. Since the ℓ-adic Tate module of a finite étale k-group vanishes, we have $T_\ell(G) = T_\ell(\mathcal{A}_k)$; hence, likewise $V_\ell(G) = V_\ell(\mathcal{A}_k)$. We have

$$V_\ell(A)^I \simeq V_\ell(\mathcal{A}(R^{\mathrm{sh}})) \simeq V_\ell(\mathcal{A}_k(k_s)),$$

so since $V_\ell(\mathcal{A}_k(k_s)) = V_\ell(\mathcal{A}_k^0(k_s))$ we may substitute into (3.2) to get an exact sequence of $\mathrm{Gal}(F^{\mathrm{sh}}/F)/I = \mathrm{Gal}(k_s/k)$-modules

$$0 \to V_\ell(T) \to V_\ell(A)^I \to V_\ell(B) \to 0 \tag{3.3}$$

where T and B are from the Chevalley structure theorem for $\mathcal{A}_k^0$. For F a global field, this will be the key to disposing of our independence-of-ℓ problem at bad places (for the definition of the L-function), as well as motivating how to define the volume term Ω_A (in a conceptual manner, inspired by an idea of Tamagawa).

4 L-factors via point-counting

Now consider a complete discrete valuation ring R with *finite* residue field k of size q and fraction field F. Let $\phi \in \mathrm{Gal}(k_s/k)$ be the q-power Frobenius automorphism of k_s, and $I \subset \mathrm{Gal}(F_s/F)$ the inertia subgroup. Let A be an abelian variety over F, and ℓ a prime distinct from $p := \mathrm{char}\,(k)$. We are going to use Néron models to settle various open issues in earlier lectures.

Application 1. Recall our first puzzle: for the linear dual $W = V_\ell(A)^* = \mathrm{H}^1_{\text{ét}}(A_{F_s}, \mathbf{Q}_\ell)$, is the polynomial $\det(1 - \phi^{-1}t \,|\, W^I) \in \mathbf{Q}_\ell[t]$ in $\mathbf{Q}[t]$, and as such independent of ℓ? Via the Weil pairing, $W \simeq V_\ell(\widehat{A})(-1)$ as $\mathrm{Gal}(F_s/F)$-modules where $\widehat{A}$ is the dual abelian variety, and $\mathbf{Q}_\ell(-1)$ is unramified, so $W^I = V_\ell(\widehat{A})^I(-1)$. Since A is F-isogenous to $\widehat{A}$, so their V_ℓ's are $\mathrm{Gal}(F_s/F)$-equivariantly isomorphic, $W^I \simeq V_\ell(A)^I(-1)$. The action of ϕ^{-1} on $\mathbf{Q}_\ell(1)$ is multiplication by $1/q$, so on $\mathbf{Q}_\ell(-1)$ it acts through multiplication by q. Hence,

$$L_\ell(t) := \det(1 - \phi^{-1}t \,|\, W^I) = \det(1 - q\phi^{-1}t \,|\, V_\ell(A)^I).$$

We want this to lie in $\mathbf{Q}[t]$ and to be independent of ℓ. In view of the $\mathrm{Gal}(k_s/k)$-equivariant (3.3), this reduces to the analogous statement for $V_\ell(T)$ and $V_\ell(B)$ in place of $V_\ell(A)^I$.

The rationality and independence-of-ℓ for $V_\ell(B)$ is part of the theory of abelian varieties over finite fields. Indeed, the Riemann Hypothesis tells us that $\lambda \mapsto q/\lambda$ is a permutation of the roots (in an algebraic closure of $\mathbf{Q}_\ell$) of the characteristic polynomial of ϕ on $V_\ell(B)$, so

$$\det(1 - q\phi^{-1}t \,|\, V_\ell(B)) = \det(1 - \phi t \,|\, V_\ell(B)),$$

and the theory over finite fields tells us that this final polynomial has the desired properties.

Let's now pass to the torus. By definition the covariant geometric cocharacter group $\mathrm{X}_*(T) = \mathrm{Hom}_{k_s}(\mathbf{G}_m, T_{k_s})$ has mod-ℓ^m reduction $T[\ell^m](k_s) \otimes \mu_{\ell^m}^{-1}$ as a Galois module, and passage to the inverse limit over m (and then inverting ℓ) gives a $\mathrm{Gal}(k_s/k)$-equivariant isomorphism

$$V_\ell(T) = \mathrm{X}_*(T)_{\mathbf{Q}_\ell} \otimes_{\mathbf{Q}_\ell} \mathbf{Q}_\ell(1).$$

The ℓ-adic representation $\mathrm{X}_*(T)_{\mathbf{Q}_\ell}$ of $\mathrm{Gal}(k_s/k)$ is manifestly tensored up from the rational representation $\mathrm{X}_*(T)_{\mathbf{Q}}$ (on which an open subgroup of $\mathrm{Gal}(k_s/k)$ acts trivially), so the relevant characteristic polynomial is obviously rational and independent of ℓ; the extra Tate twist by $\mathbf{Q}_\ell(1)$ merely has the effect of scaling the variable in the characteristic polynomial by q. So our independence-of-ℓ problem is finally settled, thanks to the magic of Néron models!

Application 2. Let's prove a formula for $\#\mathcal{A}^0(k)$ in terms of

$$L(t) := \det(1 - \phi^{-1}t \,|\, \mathrm{H}^1_{\text{ét}}(A_{F_s}, \mathbf{Q}_\ell)^I) = \det(1 - q\phi^{-1}t \,|\, V_\ell(A)^I) \in \mathbf{Q}[t]$$

whose independence-of-ℓ properties were analyzed above:

Theorem 4.1. *We have*

$$\#\mathcal{A}^0_k(k) = q^{\dim A} L(1/q)$$

where $L(t)$ is the "local L-function" attached to A.

Proof. Note that $L(1/q) = \det(1 - \phi^{-1} \,|\, V_\ell(A)^I)$. Consider the exact sequence

$$1 \to T \times U \to \mathcal{A}^0_k \to B \to 0.$$

Lang's theorem (the vanishing of degree-1 Galois cohomology of smooth connected groups over finite fields) implies that the induced diagram of k-points is short exact, so

$$\#\mathcal{A}^0_k(k) = \#B(k) \cdot \#T(k) \cdot \#U(k).$$

To analyze the torus contribution, we use the following formula from §1.5, Chapter I of Osterlé's awesome 1984 Inventiones paper on Tamagawa numbers:

$$\#T(k) = \det(q - \phi \mid \mathrm{X}^*(T)) \in \mathbf{Q}_{>0}.$$

The $\mathrm{Gal}(k_s/k)$-equivariant perfect duality

$$\mathrm{X}^*(T) \times \mathrm{X}_*(T) \to \mathrm{End}(\mathbf{G}_m) = \mathbf{Z}$$

defined via composition of cocharacters and characters implies that $\mathrm{X}^*(T)_{\mathbf{Q}}$ is the $\mathbf{Q}$-linear dual to $\mathrm{X}_*(T)_{\mathbf{Q}}$ in a manner that identifies the action of $\gamma \in \mathrm{Gal}(k_s/k)$ on the rationalized geometric character group with the linear dual of the action of γ^{-1} (check!) on the rationalized geometric cocharacter group.

Since passage to the dual preserves the determinant, we conclude that

$$\#T(k) = \det(q - \phi^{-1} \mid \mathrm{X}_*(T)_{\mathbf{Q}}) = q^{\dim T} \cdot \det(1 - (q\phi)^{-1} \mid \mathrm{X}_*(T)_{\mathbf{Q}}).$$

We can compute the determinant over $\mathbf{Q}_\ell$ and use the Galois-equivariant isomorphism $\mathrm{X}_*(T)_{\mathbf{Q}} \otimes_{\mathbf{Q}} \mathbf{Q}_\ell \simeq V_\ell(T)(-1)$ to get that

$$\#T(k) = q^{\dim T} \det(1 - \phi^{-1} \mid V_\ell(T)).$$

To analyze the B-part, if $\{\lambda_i\}$ is the set of roots (with multiplicity) of $L(B,t) = \det(1 - \phi t \mid V_\ell(B))$ then the Riemann Hypothesis over finite fields gives that $\lambda_i \mapsto q/\lambda_i$ is a permutation of the roots, so

$$\begin{aligned}
\#B(k) &= L(B,1) \\
&= q^{\dim B} \prod (1 - \lambda_i/q) \\
&= q^{\dim B} \prod (1 - 1/\lambda_j) \\
&= q^{\dim B} \det(1 - \phi^{-1} t \mid V_\ell(B)).
\end{aligned}$$

Since $\dim B + \dim T = \dim \mathcal{A}_k^0 - \dim U = \dim A - \dim U$ we now get

$$\#T(k)\#B(k) = q^{\dim A - \dim U} \det(1 - \phi \mid V_\ell(A)^I) = q^{\dim A - \dim U} L(1/q).$$

But recall that over a *perfect* field (such as k) a unipotent group U has a composition series with successive quotients isomorphic to $\mathbf{G}_a$, so $\#U(k) = q^{\dim U}$. Multiplying the previous equation by this one yields that $\#\mathcal{A}_k^0(k) = \#U(k)\#T(k)\#B(k)$ is equal to $q^{\dim A} L(1/q)$. $\square$

So far we have only counted points of the identity component. But by Lang's theorem applied to $\mathcal{A}_k^0$ we see that the exact sequence of smooth k-groups

$$0 \to \mathcal{A}_k^0 \to \mathcal{A}_k \to \Phi \to 0$$

gives an exact sequence

$$0 \to \mathcal{A}_k^0(k) \to \mathcal{A}_k(k) \to \Phi(k) \to 0.$$

Hence,

$$\#\mathcal{A}_k(k) = \#\Phi(k) \cdot \#\mathcal{A}_k^0(k) = \#\Phi(k) q^{\dim A} L(1/q).$$

The factor $\#\Phi(k)$ is called the *Tamagawa factor*, usually denoted c_v when $F = K_v$ for a global field K and non-archimedean place v of K.

Example 4.2. For the elliptic curve 57C2 from Example 2.5, we have $\Phi_3 \simeq \mathbf{Z}/2\mathbf{Z}$ and $\Phi_{19}(\overline{\mathbf{F}}_{19}) \simeq \mathbf{Z}/5\mathbf{Z}$ with nontrivial Galois action, so $c_3 = 2$ and $c_{19} = 1$ but as a finite étale $\mathbf{F}_{19}$-group the "order" of Φ_{19} is 5. So don't mix up the Tamagawa factor with the order of the component group (as a finite étale group scheme over the residue field)!

Now we switch gears to consider volumes and measures, again using the Néron model. Since R is complete, for any R-smooth scheme $\mathcal{X}$ the natural map $\mathcal{X}(R) \twoheadrightarrow \mathcal{X}(k)$ is surjective due to the Zariski-local structure of smooth schemes (étale over an affine space). Moreover, since $\mathcal{A}^0$ is open in $\mathcal{A}$, if a section Spec $R \to \mathcal{A}$ takes the closed point into $\mathcal{A}^0$ then it must factor through $\mathcal{A}^0$. Hence, we have a fiber square

$$\begin{array}{ccc} \mathcal{A}^0(R) & \longrightarrow & \mathcal{A}^0(k) \\ \downarrow & & \downarrow \\ \mathcal{A}(R) & \twoheadrightarrow & \mathcal{A}(k) \end{array}$$

so $\Phi(k) = \mathcal{A}(R)/\mathcal{A}^0(R)$. But $\mathcal{A}(R) = A(F)$ by the Néronian property, and usually the open finite-index subgroup $\mathcal{A}(R)^0$ is denoted $A(F)^0$ (strictly meaningless since $A(F)$ is totally disconnected), so one may also write $A(F)/A(F)^0 = \Phi(k)$.

Fix the standard Haar measure μ on F (normalized so that $\mu(R) = 1$, and hence $\mu(\mathfrak{m}) = 1/q$). The Change of Variables Formula and Inverse/Implicit Function Theorems from multivariable calculus also hold over any non-archimedean local field (using the *normalized* absolute value in the Change of Variables Formula); this is an extended exercise for the reader (though for the Inverse and Implicit Function Theorems over an arbitrary non-archimedean field one can find a complete discussion in Serre's book *Lie groups and Lie algebras*). This allows us to define a measure $|\omega|\mu^d$ on $X(F)$ for any smooth F-scheme X with pure dimension d and top-degree differential form ω on X. (There is no relevance for "orientations" since we work with $|\omega|$; we do not care about additivity in ω, in contrast with integration on usual smooth manifolds.) If we scale ω by $c \in F^\times$ then this measure scales by $|c|$.

Now consider the case $X = A$. The Néron model defines a preferred R-line $\Omega^{\mathrm{top}}_{\mathcal{A}/\mathcal{R}}$ inside the F-line $\Omega^{\mathrm{top}}_{A/F}$ of top-degree differential forms on A. Choose ω to be a generator of this R-line; this choice is unique up to $R^\times$. Using this choice, the measure $|\omega|\mu^d$ on $A(F)$ is *independent* of such ω, so it is truly canonical (but with definition resting crucially on $\mathcal{A}$ to pick out a preferred R-line inside $\Omega^{\mathrm{top}}_{A/F}$). This is quite remarkable: in contrast with the situation over $\mathbf{R}$ and $\mathbf{C}$, over a non-archimedean local field the group of rational points of an abelian variety has a *canonical* measure! Moreover, this is a Haar measure because any such ω is translation-invariant (in the scheme-theoretic sense, over R).

Exercise 4.3. Using this canonical measure, what is $\mathrm{vol}(A(F))$? We claim it is equal to $\#\Phi(k)L(1/q)$, where $L(t) \in \mathbf{Q}[t]$ is the local L-function attached to A.

To prove this, the key point is the second equality in:

$$\mathrm{vol}(A(F)) = \#\Phi(k) \cdot \mathrm{vol}(\mathcal{A}^0(R)) = \#\Phi(k) q^{-\dim A} \#\mathcal{A}^0(k).$$

One establishes the second equality by rigorously proving (using the Zariski-local structure of smooth morphisms) that for any smooth R-scheme $\mathcal{X}$ (e.g., $\mathcal{A}$) with relative dimension d and any top-degree differential form ω on $\mathcal{X}$ that is nowhere-vanishing on $\mathcal{X}_k$ (if one exists, as is the case for $\mathcal{A}$!), the fibers of $\mathcal{X}(R) \twoheadrightarrow \mathcal{X}(k)$ are analytically isomorphic to $\mathfrak{m}^d$ carrying $|\omega|\mu^d$ to the *standard* Haar measure obtained from R^d.

With that second equality established, the formula $\#\mathcal{A}^0(k) = q^{\dim A} L(A, 1/q)$ from Theorem 4.1 then gives the desired result.

5 The Tate-Shafarevich group

Definition 5.1. Let K be a global field. Let $S = \mathrm{Spec}\,(\mathcal{O}_K)$ in the number field case, and let S be the associated smooth proper curve over a finite field in the function field case. Consider an abelian variety A over K and let $\mathcal{A}$ be its Néron model over S. The *Tate-Shafarevich group* $\text{Ш}(A)$ is defined to be

$$\text{Ш}(A) = \ker(H^1(K, A) \to \prod_v H^1(K_v, A)).$$

We make some remarks on deciphering this definition. By Galois descent, the cohomology group $H^1(K, A) := H^1(K_s/K, A(K_s))$ is identified with the set of isomorphism classes of A-torsors X over K. For such a torsor to be trivial in $H^1(K_v, A)$ is exactly to say that it has a K_v-point. Hence, $\text{Ш}(A)$ classifies torsors with a rational point over *all* completions of K. That is:

$$\text{Ш}(A) \simeq \{A\text{-torsors } X \text{ such that } X(K_v) \neq \emptyset \text{ for all } v\}.$$

In [6.5, BLR] it is proved that the torsors X classified by $\text{Ш}(A)$ (or more generally admitting a point over the maximal unramified extension of K_v for all v) *also* admit a (separated and finite type) Néron model $\mathcal{X}$ which is moreover a torsor for $\mathcal{A}$. The same goes if X arises from the modified (and more "algebro-geometric") subgroup

$$\text{Ш}(A)' \subset \mathrm{H}^1(K, A)$$

defined similarly to $\text{Ш}(A)$ but without local triviality conditions at archimedean (or equivalently, real) places (it agrees with $\text{Ш}(A)$ in the function field case and contains $\text{Ш}(A)$ with finite 2-power index in the number field case). The operation $X \rightsquigarrow \mathcal{X}$ defines a map of sets

$$\text{Ш}(A)' \to \mathrm{H}(S, \mathcal{A}) \tag{5.1}$$

that is easily seen to be a homomorphism by using the "contracted product" description of the group law on H^1 in terms of torsors. The kernel is trivial since if $\mathcal{X}(S) \neq \emptyset$ then certainly $X(K) \neq \emptyset$. What is $\text{Ш}(A)'$ as a subgroup of $\mathrm{H}^1(S, \mathcal{A})$?

Proposition 5.2 (Mazur)**.** *For any A as above,*

$$Ш(A)' = \mathrm{Im}\,(\mathrm{H}^1(S, \mathcal{A}^0) \to \mathrm{H}^1(S, \mathcal{A})). \tag{5.2}$$

Before we give the proof of this result, we note that it really is necessary to speak of "image" on the right side of (5.2) because the map of H^1's induced by $\mathcal{A}^0 \hookrightarrow \mathcal{A}$ can *fail* to be injective. Indeed, by consideration with $\mathcal{O}_{S,s}^{\mathrm{sh}}$-points for s in the finite locus Σ of bad fibers in S we get (exercise!) an exact sequence

$$1 \to \mathcal{A}^0 \to \mathcal{A} \to \bigoplus_{s \in \Sigma} (j_s)_*(\Phi_s) \to 0 \tag{5.3}$$

for the étale topology on S, where j_s is pushforward along $\mathrm{Spec}\,(k(s)) \hookrightarrow S$. Hence, the obstruction to injectivity of the map of H^1's on the right side of (5.2) is precisely the cokernel of

$$A(K) = \mathcal{A}(S) \to \bigoplus_{s \in \Sigma} \Phi_s(k(s)).$$

In Example 2.5 (with $S = \mathrm{Spec}\ \mathbf{Z}$) we saw an example in which this latter map fails to be surjective (i.e., has nonzero cokernel), using an explicit elliptic curve over $\mathbf{Q}$.

Now we turn to the proof of Proposition 5.2

Proof. Let $T(A)$ denote the right side of (5.2). By (5.3), we have

$$T(A) := \ker(\mathrm{H}^1(S, \mathcal{A}) \to \bigoplus_{s \in \Sigma} \mathrm{H}^1(s, \Phi_s)). \tag{5.4}$$

We will show that the right side of (5.2) is contained in $Ш(A)'$ and that the right side of (5.4) contains $Ш(A)'$. (Note that $\mathrm{H}^1(s, \Phi_s)$ is finite for each s, as is $\Phi_s(s)$, so $T(A)$ is off from $\mathrm{H}^1(S, \mathcal{A})$ by a "finite amount"; in particular, finiteness for $Ш(A)$ will be equivalent to finiteness of $\mathrm{H}^1(S, \mathcal{A})$.)

We begin with (5.2), but must first address a delicate technical question: does every class in $\mathrm{H}^1(S, \mathcal{A})$ actually arise from a $\mathcal{A}$-torsor? The Čech-cohomology interpretation in degree 1 identifies such cohomology classes with (an equivalence class of) étale descent data for such a torsor, but effectivity of étale descent is not obvious since $\mathcal{A}$ is not affine. The effectivity is a hard theorem, generally applicable with $\mathcal{A}$ replaced by any smooth separated group schemes of finite type over a Dedekind base, and is proved in [6.5/1, BLR] (the main point of which is a quasi-projectivity result for torsors over a Dedekind base); the same holds for $\mathcal{A}^0$.

To analyze the right side of (5.2), we may now consider an $\mathcal{A}^0$-torsor $\mathcal{Y}$ over S, and let $\mathcal{X}$ be its pushout along $\mathcal{A}^0 \to \mathcal{A}$ to an $\mathcal{A}$-torsor over S. We want to show that the class of the generic fiber $X = \mathcal{X}_K$ in $\mathrm{H}^1(K, A)$ lies in $Ш(A)'$. Note that $X = \mathcal{Y}_K$ since $\mathcal{A}^0$ has the same generic fiber as $\mathcal{A}$. Our problem then is to prove that $\mathcal{Y}(K_v)$ is non-empty for all v. For that purpose it is suffices to prove $\mathcal{Y}(\mathcal{O}_{K_v})$ is non-empty for all v. Since $\mathcal{Y}$ is S-smooth and $\mathcal{O}_{K_v}$ is henselian, it is sufficient to show that the special fiber $\mathcal{Y}_v$ at each closed point $v \in S$ has a rational point over the finite field $k(v)$. But $\mathcal{Y}_v$ is a torsor for the smooth finite type $k(v)$-group $\mathcal{A}_v^0$ that is *connected*. Lang's theorem gives that a torsor for a smooth connected group scheme over a *finite* field is always trivial, so we win: the right side of (5.2) is contained inside $Ш(A)'$.

Now consider an A-torsor X arising from $Ш(A)'$, and let $\mathcal{X}$ be its Néron model, which we have noted earlier exists and is an $\mathcal{A}$-torsor (due to the hard work in [6.5, BLR]). Our task is to show that the class of $\mathcal{X}$ in $\mathrm{H}^1(S, \mathcal{A})$ lies inside the right side of (5.4). This is now an entirely local problem: for each bad place s, is the image of this global class in $\mathrm{H}^1(s, \Phi_s)$ trivial? It is an instructive exercise to check that the class obtained in $\mathrm{H}^1(s, \Phi_s)$ is exactly the Φ_s-torsor $\mathcal{X}_s/\mathcal{A}_s^0$. Hence, we want this latter torsor to be trivial. Even better, $\mathcal{X}_s$ as an $\mathcal{A}_s$-torsor is trivial! This does not come for free from Lang's theorem (as $\mathcal{A}_s$ is generally disconnected), but rather is due to the magic of the Néronian mapping property: it suffices to show that $\mathcal{X}(\mathcal{O}_{S,s}^\wedge)$ is non-empty (as then passing to the special fiber gives a rational point on $\mathcal{X}_s$), but by the very construction of $\mathcal{X}$ as a Néron model it retains the Néronian mapping property after base change to the *completion* $\mathcal{O}_{S,s}^\wedge$. Hence, $\mathcal{X}(\mathcal{O}_{S,s}^\wedge) = \mathcal{X}(K_s) = X(K_s)$, and by our hypothesis involving X and $Ш(A)'$ we know that $X(K_s)$ is non-empty. □

Considering $Ш(A)'$ as a subgroup of $\mathrm{H}^1(S, \mathcal{A})$, clearly $Ш(A)$ is obtained from this subgroup by imposing the further condition on an $\mathcal{A}$-torsor that it admit a local point at all real places of K (a vacuous condition unless K is a number field with a real embedding). The ability to interpret classes in $Ш(A)'$ in terms of $\mathcal{A}$-torsors over the global base S yields an alternative description of $Ш(A)'$ as follows.

Under the restriction map $\mathrm{H}^1(S, \mathcal{A}) \to \mathrm{H}^1(K_s, A)$, the image lies inside the subgroup

$$\mathrm{H}^1(K_s^{\mathrm{sh}}/K_s, A(K_s^{\mathrm{sh}})) = \ker(\mathrm{H}^1(K_s, A) \to \mathrm{H}^1(K_s^{\mathrm{sh}}, A))$$

of "unramified torsors" (where the fraction field K_s^{sh} of the strict henselization $(\mathcal{O}_{S,s}^\wedge)^{\mathrm{sh}}$ is the maximal unramified extension of $K_s = \mathrm{Frac}(\mathcal{O}_{S,s}^\wedge)$). Indeed, if $\mathcal{X}$ is an $\mathcal{A}$-torsor over S with generic fiber A-torsor X then the set $X(K_s^{\mathrm{sh}}) = \mathcal{X}(K_s^{\mathrm{sh}})$ is non-empty because $\mathcal{X}((\mathcal{O}_{S,s}^\wedge)^{\mathrm{sh}})$ is non-empty (as for *any* smooth scheme with non-empty fiber over a strictly henselian local ring, by lifting any separable closed point in the non-empty smooth special fiber). This leads to:

Corollary 5.3. *Let $U \subset S$ be a dense open subscheme whose complement contains the set Σ of bad places for A. Then the kernel of the natural map*

$$\mathrm{H}^1(S, \mathcal{A}) \to \bigoplus_{s \in S-U} \mathrm{H}^1(K_s^{\mathrm{sh}}/K_s, A(K_s^{\mathrm{sh}}))$$

is equal to $Ш(A)'$.

Proof. Composing with the injective map from unramified classes of A-torsors over K_s into the group of (isomorphism classes of) all A-torsors over K_s, it is equivalent to show that $Ш(A)'$ consists of the classes in $\mathrm{H}^1(S,\mathcal{A})$ that have trivial restriction over K_s for all $s \in S - U$. Certainly if X is an A-torsor coming from $Ш(A)'$ and $\mathcal{X}$ is its associated $\mathcal{A}$-torsor Néron model then $\mathcal{X}_{K_s} = X_{K_s}$ has a K_s-point for *all* closed points $s \in S$. Thus, the kernel in the statement of the Corollary contains $Ш(A)'$.

Conversely, if $\mathcal{X}$ is a representative of a class in that kernel then we want to show that $\mathcal{X}$ arises from $Ш(A)'$, which is to say that its generic fiber X has a K_s-point for all closed points $s \in S$. If s is a point of good reduction for A (such as anything outside $S - U$) then $\mathcal{A}_s$ is an abelian variety, and in particular is connected, so its torsor $\mathcal{X}_s$ has a rational point by Lang's theorem. Thus, by smoothness this lifts to an $\mathcal{O}^{\wedge}_{S,s}$-point of $\mathcal{X}$. Passing to the generic fiber gives a K_s-point of X for such s. It therefore remains to consider bad s, all of which are contained in $S - U$. But by the choice of $\mathcal{X}$ we know $X(K_v)$ is non-empty for all $v \in S - U$, so we are done. $\square$

The displayed map in Corollary 5.3 has a cohomological interpretation as follows. Let $j : U \hookrightarrow S$ be the natural open immersion, so $\mathcal{A}_U$ is an abelian scheme over U, and its associated sheaf on $U_{\text{ét}}$ is precisely $j^*(\mathcal{A})$ (why?). The pushforward sheaf $j_*(\mathcal{A}_U)$ on $S_{\text{ét}}$ is precisely the functor of points of $\mathcal{A}$ by the Néronian mapping property (exercise!). Consider the resulting short exact sequence of abelian sheaves

$$0 \to j_!(\mathcal{A}_U) \to j_*(\mathcal{A}_U) = \mathcal{A} \to \mathscr{G} \to 0$$

on $S_{\text{ét}}$ where the first map is the natural inclusion (an isomorphism over U!) and $\mathscr{G}$ is defined to be the cokernel.

Note that $\mathscr{G}$ is a skyscraper sheaf supported at the finite set $S - U$ of closed points. Its stalk at each $s \in S - U$ is the discrete Galois module (at s) associated to the global sections of its pullback over $\mathcal{O}^{\mathrm{sh}}_{S,s}$ (as for computing the stalk of any étale sheaf on any scheme whatsoever). But that stalk is the same as for $j_*(\mathcal{A}_U)$ (by how $j_!$ is defined, with vanishing stalks outside U), and the formation of j_* commutes with ind-étale base change (such as pullback to $\mathrm{Spec}\,(\mathcal{O}^{\mathrm{sh}}_{S,s})$), so we conclude that $\mathscr{G}_s$ corresponds to the discrete Galois module

$$\mathcal{A}_U(K^{\mathrm{sh},s}) = A(K^{\mathrm{sh},s}),$$

where $K^{\mathrm{sh},s}$ denotes the fraction field of a strict henseliization of $\mathcal{O}_{S,s}$.

Thus, passing to the long exact cohomology sequence yields an exact sequence

$$\bigoplus_{s \in S-U} A(K^{\mathrm{h},s}) \xrightarrow{\delta} \mathrm{H}^1(S, j_!(\mathcal{A}_U)) \to \mathrm{H}^1(S,\mathcal{A}) \to \bigoplus_{s \in S-U} \mathrm{H}^1(K^{\mathrm{sh},s}/K^{\mathrm{h},s}, A(K^{\mathrm{sh},s}))$$

where $K^{\mathrm{h},s}$ is the fraction field of the henselization of $\mathcal{O}_{S,s}$. The final term in this long exact sequence classifies unramified A-torsors over $K^{\mathrm{h},s}$, and we claim that this H^1 *injects* into the analogue over the *completion* K_s (of K at s, or equivalently of $K^{\mathrm{h},s}$). Unramified torsors admit Néron models with non-empty special fiber (again, by the hard work in [6.5, BLR]), and by the Néronian property triviality of such a torsor is equivalent to the special fiber of its Néron model having a rational point, so we have established injectivity of the natural map

$$\mathrm{H}^1(K^{\mathrm{sh},s}/K^{\mathrm{h},s}, A(K^{\mathrm{sh},s})) \to \mathrm{H}^1(K_s^{\mathrm{sh}}/K_s, A(K_s^{\mathrm{sh}})).$$

Consequently, we get an exact sequence

$$\bigoplus_{s\in S-U} A(K^{\mathrm{h},s}) \xrightarrow{\delta} \mathrm{H}^1(S, j_!(\mathcal{A}_U)) \to \mathrm{H}^1(S, \mathcal{A}) \to \bigoplus_{s\in S-U} \mathrm{H}^1(K_s^{\mathrm{sh}}/K_s A(K_s^{\mathrm{sh}}))$$

where now the final map is exactly what arose in Corollary 5.3. So applying that corollary now gives two descriptions of $Ш(A)'$: it is a cokernel

$$\bigoplus_{s\in S-U} A(K^{\mathrm{h},s}) \xrightarrow{\delta} \mathrm{H}^1(S, j_!(\mathcal{A}_U)) \to Ш(A)' \to 0 \tag{5.5}$$

as well as a kernel:

$$Ш(A)' = \ker(\mathrm{H}^1(S, \mathcal{A}) \to \prod_{s\in S-U} \mathrm{H}^1(K_s, A)).$$

The cokernel presentation of $Ш(A)'$ in (5.5) underlies a very conceptual definition of the Cassels–Tate pairing

$$Ш(A) \times Ш(\widehat{A}) \to \mathbf{Q}/\mathbf{Z}$$

by means of étale cohomology over open subschemes of S (though other definitions will be useful too). There is yet another description of $Ш(A)'$ is useful when constructing this conceptual construction of the Cassels–Tate pairing and when relating Tate–Shafarevich groups to Brauer groups in the setting of the Artin–Tate conjecture (as we will see in a later lecture):

Proposition 5.4. *The natural restriction map* $\mathrm{H}^1(U, \mathcal{A}) \to \mathrm{H}^1(K, A)$ *is injective and its image contains* $Ш(A)'$*. Under this identification of* $Ш(A)'$ *as a subgroup of* $\mathrm{H}^1(U, \mathcal{A})$*, we have*

$$Ш(A)' = \ker(\mathrm{H}^1(U, \mathcal{A}) \to \prod_{s\in S-U} \mathrm{H}^1(K_s, A)).$$

Proof. We have seen that every class in $\mathrm{H}^1(U, \mathcal{A})$ is represented by an $\mathcal{A}_U$-torsor $\mathcal{X}$ (this was discussed over S, but the same reasoning applies over any non-empty open subscheme of S), and this inherits properness from $\mathcal{A}_U$. Hence, by the valuative criterion for properness applied over the Dedekind base U we have $\mathcal{X}(U) = \mathcal{X}(K)$, so injectivity is clear.

We have identified $Ш(A)'$ with the image of $\mathrm{H}^1(S, \mathcal{A}^0)$ in $\mathrm{H}^1(S, \mathcal{A})$ via the formation of Néron models of A-torsors unramified at all finite places, so composition with the restriction map $\mathrm{H}^1(S, \mathcal{A}) \to \mathrm{H}^1(K, A)$ amounts to forming the generic fibers of such torsors; that recovers the definition of $\text{Ш}(A)'$ as a subgroup of $\mathrm{H}^1(K, A)$. This shows that $\mathrm{H}^1(U, \mathcal{A})$ as a subgroup of $\mathrm{H}^1(K, A)$ contains $\text{Ш}(A)'$ as a subgroup killed by the restriction map to $\mathrm{H}^1(K_s, A)$ for all $s \in S - U$.

To show that the local triviality condition at all $s \in S - U$ already cuts $\mathrm{H}^1(U, \mathcal{A})$ down to $\text{Ш}(A)'$, it suffices to prove that for all $s \in U$ the restriction map $\mathrm{H}^1(U, \mathcal{A}) \to \mathrm{H}^1(K_s, A)$ vanishes. This factors through $\mathrm{H}^1(\mathcal{O}^{\wedge}_{S,s}, \mathcal{A})$, so it is enough to prove that this latter H^1 vanishes. More generally, if R is a complete (or henselian) discrete valuation ring with finite residue field κ and fraction field F, and if B is an abelian scheme over R, then we claim that $\mathrm{H}^1(R, B) = 0$.

Since every étale cover of R is dominated by a *finite* étale cover (as R is local henselian), and that in turn is dominated by one which is Galois, it suffices to fix an unramified finite Galois extension F'/F with associated valuation ring R' over R and show that the Čech cohomology group $\mathrm{H}^1(R'/R, B)$ vanishes. This classifies Galois descent datum for a B-torsor, and such descent is effective since B is quasi-projective over R. Hence, such cohomology classes correspond to B-torsors X over R that split over R'. But then X_κ is a B_κ-torsor, so by connectedness of B_κ and Lang's theorem we see that the set $X(\kappa) = X_\kappa(\kappa)$ is non-empty. By smoothness of X and the henselian property of R, such a κ-point lifts to an R-point of X, so we are done. □

Remark 5.5. The preceding considerations identify $\text{Ш}(A)'$ as the image of an excision map $\mathrm{H}^1_c(U, \mathcal{A}_U) \to \mathrm{H}^1(S, \mathcal{A})$ for any dense open $U \subset S$ avoiding the bad points, up to the caveat that a good definition of $\mathrm{H}^\bullet_c(U, \cdot)$ when K admits real places requires some modification of $\mathrm{H}^\bullet(S, j_!(\cdot))$. We will return to this matter after setting up Tate global duality by means of étale cohomology on open subschemes of S.

6 The volume factor

Finally, for an abelian variety A of dimension $d > 0$ over a global field K, we wish to discuss Ω_A appearing in the BSD conjecture. This is most elegantly defined to be $\mathrm{vol}(A(\mathbf{A}_K))$ where the volume is with respect to the "Tamagawa measure", a God-given Haar measure on the locally compact group $A(\mathbf{A}_K)$. (See my expository paper on Weil and Grothendieck approaches to topologies of adelic points beyond the affine case for a discussion of putting a topology on the $\mathbf{A}_K$-points of any separated K-scheme of finite type, especially going beyond the affine case.)

We will build this measure using a global top-degree differential form. There is a subtlely, only apparent when $\mathrm{Pic}(S)$ is non-trivial: the line bundle $\Omega^{\mathrm{top}}_{\mathcal{A}/S}$ on S of top-degree differential forms on $\mathcal{A}$ (this line bundle is the top exterior power of the cotangent space along the identity section, using relative translation arguments) may not be globally free, so we may not be able to choose a global generator ω of the sheaf of top-degree differential forms on $\mathcal{A}$.

Pick an arbitrary nonzero $\omega \in \Omega^{\mathrm{top}}_{A/K}$. For all but finitely many closed points $s \in S$, this is a generator of the $\mathcal{O}_s$-line $\Omega^{\mathrm{top}}_{\mathcal{A}_{\mathcal{O}_s}/\mathcal{O}_s}$ inside the K-line $\Omega^{\mathrm{top}}_{A/K}$. In general, for a general closed point $v \in S$ (viewed as a non-archimedean valuation on K), ω defines a global nowhere-vanishing top-degree differential form on $A(K_v)$, If η_v is a global generator of top-degree differential forms on $\mathcal{A}_{\mathcal{O}_{K_v}}$ then $\omega = a_v \eta_v$ for some $a_v \in K_v^\times$.

For the normalized measure μ_v on K_v, the Haar measure $|\omega|_v \mu_v^d$ is equal to $|a_v|_v |\eta_v|_v \mu_v^d$, which is to say it is equal to $|a_v|_v$ times the *canonical* Haar measure of $A(K_v)$ considered earlier. For *archimedean* places v on K there is no canonical measure on $A(K_v)$, but we do get an associated Haar measure $|\omega|_v \mu_v^d$ where μ_v is taken to be the Lebesgue measure on K_v in the real case and twice the Lebesgue measure on K_v in the complex case, and likewise $|\cdot|_v$ is the standard absolute value on K_v at real places and the *square* of the standard one at complex places (so the product formula holds when using these and the normalized valuation at the finite places, and the scaling effect on all additive Haar measures under multiplication by $c \in K_v^\times$ is $|c|_v$). The factor of 2 in the normalized Haar measure at complex places will emerge in an adelic calculation later on. Note that for complex v, our convention for the meaning of $|\cdot|_v$ implies that $|\omega|_v$ coincides with what is usually denoted $|\omega \wedge \overline{\omega}|$ on a complex manifold.

What happens if we change the choice of ω? The possible choices are $c\omega$ for $c \in K^\times$, so the associated Haar measure on $A(K_v)$ for any place v of K changes by $|c|_v$. Hence, the *product* of such discrepancy factors over all places is equal to 1. Thus, if we could make sense of a (restricted) product measure $\prod_v |\omega|_v \mu_v^d$ on $A(\mathbf{A}_K)$ then it would be *independent* of ω!

However, such a product measure does not make sense because the volume of $A(K_v) = \mathcal{A}(\mathcal{O}_{K_v})$ under such measures is generally not equal to 1 for all but finitely many v. More specifically, for all but finitely many v (namely, for the non-archimedean v such that ω generates the $\mathcal{O}_{K_v}$-line of top-degree v-integral differential forms) this volume has already been computed (in Exercise 4.3) to be $\#\Phi_v(k(v))L_v(1/q_v)$ where $q_v = \#k(v)$ and $L_v(t)$ is the local L-function at v. For such v which are good-reduction places, this is equal to

$$L_v(1/q_v) = q_v^{-d} L_v(1) = \#\mathcal{A}(k(v))/q_v^d,$$

and by the Riemann Hypothesis this is on the order of $1 + O(1/\sqrt{q_v})$ with an O-constant that is uniform across all such v. The product of such terms is generally *divergent*, so we make an adjustment as follows.

For each non-archimedean place v of K, define the "convergence factor"

$$\lambda_v = 1/L_v(1/q_v) = q_v^d/\#\mathcal{A}(k(v)).$$

For archimedean v, define $\lambda_v = 1$. Upon choosing ω, we get the associated Haar measures

$$m_v = \lambda_v |\omega|_v \mu_v^d$$

on $A(K_v)$ that assign total volume 1 for all but finitely many v (but which "all but finitely many v" now depends on ω!). Thus, it makes sense to define the product measure

$$m_A = \prod m_v \tag{6.1}$$

on the (topological!) product $A(\mathbf{A}_K) = \prod A(K_v)$ (not just restricted product), and this is *independent* of the choice of ω due to the product formula. Hence, it is truly canonical.

Remark 6.1. For finite places v of K, the convergence factor λ_v is the value at $s = 1$ of the reciprocal $1/L_v(q_v^{-s})$ of the local factor at v in the L-function $L(A/K, s)$.

Remark 6.2. Consider a finite separable extension of global fields K'/K and assume $A = \mathrm{R}_{K'/K}(A')$ for an abelian variety A' over K'. For each finite place v of K we claim that $\lambda_v = \prod_{v'|v} \lambda_{v'}$. Néron models are compatible with Weil restriction due to the mapping property. That is, the respective Néron models over valuation rings satisfy the relation

$$\mathcal{A}_v \simeq \prod_{v'|v} \mathrm{R}_{\mathcal{O}_{v'}/\mathcal{O}_v}(\mathcal{A}'_{v'}),$$

so

$$\mathcal{A}_v(k(v)) = \prod_{v'|v} \mathcal{A}'_{v'}(\mathcal{O}_{v'}/\mathfrak{m}_v\mathcal{O}_{v'}).$$

Letting $e_{v'}$ and $f_{v'}$ denote the ramification index and residual degrees for v' over v, so $q_{v'} = q_v^{f_v}$, we have by smoothness of the Néron model that

$$\#\mathcal{A}'_{v'}(\mathcal{O}_{v'}/\mathfrak{m}_v\mathcal{O}_{v'}) = \#\mathcal{A}_{v'}(k(v'))q_{v'}^{d'(e_{v'}-1)} = q_{v'}^{d'e_{v'}} \frac{\#\mathcal{A}'_{v'}(k(v'))}{q_{v'}^{d'}} = q_{v'}^{d'e_{v'}}/\lambda_{v'}$$

where $d' = \dim A'$. Thus,

$$\frac{q_v^d}{\lambda_v} = \prod_{v'|v} \frac{q_{v'}^{d'e_{v'}}}{\lambda_{v'}}.$$

But $q_{v'}^{d'e_{v'}} = q_v^{d'f_{v'}e_{v'}}$, so the product of these numerators over all such $v'|v$ is equal to $q_v^{d'[K':K]} = q_v^d$ (as $\mathrm{R}_{K'/K}$ multiplies the dimension of a smooth K'-scheme by $[K' : K]$). Comparing denominators now gives the desired result.

It is convenient to make one final adjustment in this construction, essential to satisfy compatibility of the associated volume relative to Weil restriction through a finite separable extension of global fields. Observe that in the preceding we have always worked with the normalized Haar measures μ_v on the local fields K_v. Such normalized measures interact poorly with finite separable Weil restriction in the presence of ramification, so it is better to make a construction that uses an *arbitrary* choice of such local Haar measures μ'_v on each K_v provided that it is the normalized choice for all but finitely many v (without needing to specify which ones).

In other words, we choose an *arbitrary* Haar measure μ' on the adele ring $\mathbf{A}_K$, which in turn always decomposes as a restricted product measure of such a collection of local Haar measures $\{\mu'_v\}$; this local collection has scaling ambiguity for any particular local measure, but the overall scaling ambiguities (all but finitely many of which are 1) multiply to 1. Also choose a nonzero top-degree differential form ω on A. Consider the associated Haar measure $m'_v = \lambda_v|\omega|_v{\mu'_v}^d$ on $A(K_v)$ for all v. These assign volume 1 to $A(K_v)$ for all but finitely many v, so it makes sense to form the product measure

$$m_{A,\mu'} = \prod m'_v.$$

This is independent of ω due to the product formula, and is indeoendent of the collection of measures $\{\mu'_v\}$ as above whose restricted product is the chosen μ', but $m_{A,\mu'}$ depends on μ'. (This refines (6.1) because it rests on an arbitrary Haar measure μ' on the adele ring.)

We make the dependence on μ' cancel out with the following scaling trick. The Haar measure μ' induces a quotient Haar measure $\overline{\mu'}$ on the compact quotient $\mathbf{A}_K/K$ compatibly with μ' on $\mathbf{A}_K$ and with counting measure on the discrete closed subgroup $K \subset \mathbf{A}_K$. The resulting volume $\overline{\mu'}(\mathbf{A}_K/K)$ is a finite positive real number, and $\overline{\mu'}(\mathbf{A}_K/K)^d$ scales under a change in μ' *exactly the same way* that $m_{A,\mu'}$ does under a change in μ'. Voila, so we define the *Tamagawa measure* on $A(\mathbf{A}_K)$ to be the Haar measure

$$m_A := \overline{\mu'}(\mathbf{A}_K/K)^{-d} m_{A,\mu'} \tag{6.2}$$

on $A(\mathbf{A}_K)$. This Haar measure is also independent of μ', so it will provide more robustness with respect to change in K below.

In the special case of our initial construction using the normalized local measures μ_v as the choice of measures μ'_v for all places v of K, the effect of this new scaling factor in (6.2) is to multiply by $\overline{\mu}(\mathbf{A}_K/K)^{-d}$, where $\overline{\mu}$ is the measure on $\mathbf{A}_K/K$ induced by counting measure on K and the restricted product μ of the measures μ_v as a measure on $\mathbf{A}_K$. It is well-known that $\overline{\mu}(\mathbf{A}_K/K)$ is equal to $\mathrm{disc}(K)^{1/2}$ in the number field case (here using our convention to insert the factor of 2 at complex places for the definition of μ_v) and q^{g-1} in the function field case (with K the function field of a geometrically connected genus-g smooth proper curve over a finite field of size q).

Definition 6.3. The volume factor Ω_A in the BSD Conjecture is

$$\Omega_A = m_A(A(\mathbf{A}_K)).$$

Exercise 6.4. Suppose K is a number field. A given choice of ω satisfies $\mathfrak{a}_\omega \cdot \omega = \Omega^d_{\mathcal{A}/\mathcal{O}_K}$ for a (generally non-principal!) fractional ideal $\mathfrak{a}_\omega$ of K that depends on ω. Then

$$\Omega_A = \mathbf{N}(\mathfrak{a}_\omega)\mathrm{disc}(K/\mathbf{Q})^{-d/2} \int_{A(K_\infty)} |\omega|_\infty \cdot \prod_{v\,\mathrm{bad}} c_v \tag{6.3}$$

where $c_v := \#\Phi(k(v))$ for the finitely many bad places v and integration on $A(K_\infty) = \prod_{v|\infty} A(K_v)$ is defined using the product of the measures $|\omega|_v$ built from ω on each $A(K_v)$ and the *standard* Haar measure on K_v for each $v|\infty$. Note that for complex v, the measure $|\omega|_v$ is precisely the traditional $|\omega \wedge \overline{\omega}|$ on $A(K_v)$ with $K_v \simeq \mathbf{C}$. This is especially classical for $K = \mathbf{Q}$ and ω a choice of "Néron differential"; i.e. generator of the $\mathbf{Z}$-line of top-degree global differential forms on the Néron model (so $\mathfrak{a}_\omega = (1)$).

The formula in (6.3) is given as a *definition* out of thin air at the end of §5 in Ch. III in Lang's book *Number Theory III*.

Recall that since L-functions are invariant under induction, the invariance of the BSD-coefficient under finite separable extension of global fields is a (weak) necessary test of the well-posedness of the conjecture. We have already seen that all pieces of the coefficient individually satisfy such invariance except for possibly $\#Ш(A)$ (whose finiteness is not known in general) and Ω_A (which we have only now finally defined). We conclude our discussion by addressing the invariance for both of these terms.

For $Ш(A)$ one can say something with content about invariance even though we do not know it to be finite: if K'/K is a finite separable extension of global fields and $A = \mathrm{R}_{K'/K}(A')$ for an abelian variety A' over K' then we claim that naturally $Ш(A') \simeq Ш(A)$. Shapiro's Lemma gives

$$\mathrm{H}^1(K', A') \simeq \mathrm{H}^1(K, A);$$

this can be defined on the level of cocycles, but it is described more conceptually in the language of torsors: $X' \rightsquigarrow \mathrm{R}_{K'/K}(X')$! (Exercise: prove these two definitions coincide.) The advantage of the torsor description is that

$$\mathrm{R}_{K'/K}(X')(K_v) = X'(K' \otimes_K K_v) = \prod_{v'|v} X'(K'_{v'})$$

for any place v of K. Hence, it is then immediate that X comes from $Ш(A)$ if and only if X' comes from $Ш(A')$, which is to say that the above Shapiro isomorphism restricts to an isomorphism $Ш(A') \simeq Ш(A)$ (which can of course also be proved in cocycle language via more notation and bookkeeping of places).

The equality between $\Omega_{A'}$ equal Ω_A is much harder, as we now explain. First, note that this is certainly sensitive to the fact that we introduced the factor $\overline{m}_{\mu'}(\mathbf{A}_K/K)^{-d}$ into the definition of the Tamagawa measure on $A(\mathbf{A}_K)$! Moreover, the equality is really *not* obvious at all, and if one tries to attack it using the explicit formula in Exercise 6.4 in the number field case then probably one gets mired in a mess.

The essential difficulty is that tracking the behavior of local measures through Weil restriction gets caught up in the Exercise at the end of last time concerning how determinant of vector bundles and norm of line bundles interact with pushforward through a finite flat map (such as an extension of valuation rings of non-archimedean local fields). This matter is addressed in elegant detail in §4–§5 of Ch. II of Oesterlé's paper, and the technique in §6 of Ch. II of that paper with abelian varieties in place of smooth connected affine groups (using Remark 6.2!) gives that the volumes of

$$A(\mathbf{A}_K) = A'(K' \otimes_K \mathbf{A}_K) = A'(\mathbf{A}_{K'})$$

associated to the respective Tamagawa measures arising from A and A' satisfy

$$\Omega_A = c^d \cdot \Omega_{A'}$$

for some constant $c > 0$ determined *solely* by considerations with local extensions arising from K'/K and has *nothing* to do with the abelian variety A' over K'.

The way that c emerges from calculations with local Haar measures on the completions of K and K' shows that the *same* c arises in the analogous considerations with smooth connected affine groups in place of abelian varieties (using a theory of convergence factors $\{\lambda_v\}$ for that case too, all trivial unless the group has nontrivial geometric character group). As in §6, Ch. II of Oesterlé's paper, we can then focus on the K'-group $\mathbf{G}_a$ (!), whereupon we see that $c = 1$ due to the scaling factor in (6.2). Hence, amusingly the invariance of Ω_A relative to finite separable Weil restriction comes down to a universal identity that we check by a calculation with smooth connected *affine* groups.