The Cassels–Tate Pairing and work of Poonen–Stoll

David Sherman

11 November 2015

Let A be an abelian variety over a number field K. For any abelian group M, let M_{div} denote the subgroup of divisible elements (i.e., $M_{\text{div}} := \bigcap_{n>0} nM$) and let $M_{\text{nd}} := M/M_{\text{div}}$ (the maximal quotient of M with no nonzero divisible elements). We're going to define a pairing

$$\langle \cdot, \cdot \rangle_A : Ш(A) \times Ш(A^\vee) \to \mathbf{Q}/\mathbf{Z}$$

such that $Ш(A)_{\text{div}}$ annhilates $Ш(A^\vee)$ and vice-versa with the roles of A and $A^\vee$ swapped, and the induced pairing $Ш(A)_{\text{nd}} \times Ш(A^\vee)_{\text{nd}} \to \mathbf{Q}/\mathbf{Z}$ will be perfect in the sense that the annihilator of each side of the pairing vanishes in the other side. (It is known that for each prime ℓ, the ℓ-primary part of $Ш(A)_{\text{nd}}$ is finite.) Moreover, for any symmetric isogeny $\varphi : A \to A^\vee$ (e.g. a polarization or the negative of such) the resulting self-pairing $\langle \cdot, \cdot \rangle_\varphi = \langle \cdot, \varphi(\cdot) \rangle_A$ on $Ш(A)$ will be skew-symmetric.

Tate introduced $\langle \cdot, \cdot \rangle_A$ in his 1962 ICM talk as a generalization of Cassels' version in the case that A is an elliptic curve. In Theorem 3.2 of his ICM talk, Tate announced that for each prime ℓ the induced pairing between ℓ-primary components $Ш(A)[\ell^\infty]$ and $Ш(A^\vee)[\ell^\infty]$ (each of which is an extension of a finite abelian ℓ-group by a divisible group of finite corank) makes the divisible part of each side annihilate the other side, and moreover that the resulting pairing between the finite groups $Ш(A)_{\text{nd}}[\ell^\infty]$ and $Ш(A^\vee)_{\text{nd}}[\ell^\infty]$ is perfect. Going further, in Theorem 3.3 of his ICM talk Tate announced that for a principal polarization $f : A \simeq A^\vee$, the self-pairing $\langle \cdot, \cdot \rangle_f$ on $Ш(A)$ is alternating provided that $f = \phi_{\mathscr{L}}$ for some ample line bundle $\mathscr{L}$ on A. (In §1 we will review the notion of polarization and its relation to the "Mumford construction" $\phi_{\mathscr{L}}$; any polarization of A_{K_s} will arise from the Mumford construction associated to a line bundle on A_{K_s} but this is generally not true on A.)

Near the end of §1 of Tate's 1966 Bourbaki talk on a geometric analogue of BSD he asserted (without reference or proof) the alternating property for $\langle \cdot, \cdot \rangle_f$ for any principal polarization f, dropping the hypothesis involving a line bundle $\mathscr{L}$ on A. The omission of that hypothesis was likely a typo by Tate, and 30 years of mistaken folklore followed until in 1998 it was shown definitively by Poonen and Stoll that dropping that hypothesis is erroneous: they built an obstruction inside $Ш(A)$ to alternation relative to a principal polarization and they built explicit examples in which their obstruction doesn't vanish. As a teaser, one such an example is the genus-g hyperelliptic curve over $\mathbf{Q}$ given by $y^2 = -(x^{2g+2} + x + 1)$ for even g.

Note that if A admits a principal polarization φ then non-degeneracy of the skew-symmetric $\langle \cdot, \cdot \rangle_\varphi$ between ℓ-primary parts of $Ш(A)_{\text{nd}}$ implies that the ℓ-part of $Ш(A)_{\text{nd}}$ has square order for odd ℓ (whereas for $\ell = 2$ it tells us nothing of the sort, so if $Ш(A)$ is finite then its overall order is a square or twice a square, depending on the parity of the multiplicity of 2 as a factor of the order). Alternation implies the 2-part has size 2^e for some even e, so in such cases the order (if finite) would be a square. But Poonen and Stoll show that alternation can fail, even when $Ш(A)$ has square order (they give an involved example where this holds).

1 Reminders: Picard schemes and polarizations

Before we talk about the pairing itself and its properties, we have to discuss some set-up related to abelian varieties and their polarizations. Though much of what we say below works in greater generality (especially, over a more general base), we'll just work over a field for simplicity. In particular, the definition(s) we shall use for the Cassels–Tate pairing will not involve integral structures such as Néron models (or torsors over them), in contrast with what was discussed in Lecture 3.

1.1 Picard Schemes

See Chapter 9 of the book "FGA Explained" for a good exposition of the material below. Notational note: in these notes we'll use Roman font to denote functors and boldface to denote representing schemes. For example, $\mathrm{Hilb}_{X/k}$ denotes a Hilbert functor (classifying flat families of closed subschemes of a projective k-scheme X) and $\mathbf{Hilb}_{X/k}$ is the representing Hilbert scheme.

Let X be a geometrically integral projective k-scheme; note that for any k-scheme T, the structure map $f_T : X_T \to T$ satisfies $\mathscr{O}_T \simeq (f_T)_*(\mathscr{O}_{X_T})$ and hence likewise $\mathscr{O}_T^\times \simeq (f_T)_*(\mathscr{O}_{X_T}^\times)$. In other words, letting $f : X \to \mathrm{Spec}(k)$ denote the structure map, we can say "$f_*(\mathbf{G}_{m,X}) = \mathbf{G}_m$ universally".

We want to parametrize families of line bundles on X via a moduli scheme. A naïve try would be to represent the functor $T \mapsto \mathrm{Pic}(X_T)$ for varying k-schemes T. But this will fail to be a sheaf for two reaons: the image of $\mathrm{Pic}(T)$ in $\mathrm{Pic}(X_T)$ is generally nontrivial but becomes trivial Zariski-locally on T (so the functor $T \mapsto \mathrm{Pic}(X_T)$ is not a Zarsiki sheaf), and line bundles on X_T are not sufficiently rigid to expect a global "universal line bundle" to exist that can be classified using morphisms on the base.

To fix the first of these problems it is tempting to Zariski-sheafify, or as an intermediate stage to consider the functor $T \mapsto \mathrm{Pic}(X_T)/\mathrm{Pic}(T)$. (The map $\mathrm{Pic}(T) \to \mathrm{Pic}(X_T)$ is injective because the natural map $\mathscr{L} \to (f_T)_* f_T^*(\mathscr{L})$ for any line bundle $\mathscr{L}$ on T, as we may check Zariski-locally on T to reduce to the case $\mathscr{L} = \mathscr{O}_T$.) It is not obvious if this quotient construction should be expected to have any reasonable sheaf properties, as is necessary for representability (any representable functor is an fpqc sheaf). Define the *relative Picard functor* $\mathrm{Pic}_{X/k}$ to be the étale-sheafification of $T \mapsto \mathrm{Pic}(X_T)$, so this is also the étale-sheafication of $T \mapsto \mathrm{Pic}(X_T)/\mathrm{Pic}(T)$. This is the same as the higher direct image sheaf $\mathrm{R}^1(f_*)(\mathbf{G}_{m,X})$ for the étale topology (since the Zariski-cohomology group $\mathrm{H}^1(X_T, \mathscr{O}_{X_T}^\times)$ is unaffected by passing to the étale topology, due to the link to Čech cohomology in degree 1 and étale descent theory for line bundles).

Theorem 1.1 (Grothendieck)**.** *The functor* $\mathrm{Pic}_{X/k}$ *is represented by a locally of finite type k-group scheme* $\mathbf{Pic}_{X/k}$ *that is a countable disjoint union of quasi-projective k-schemes. The identity component* $\mathbf{Pic}^0_{X/k}$ *is projective when X is smooth.*

The representability of $\mathrm{Pic}_{X/k}$ is deduced from the existence of Hilbert schemes. The properness of $\mathbf{Pic}^0_{X/k}$ for smooth X is a simple application of the valuative criterion for properness (using that X_R is regular for a discrete valuation ring R, as it is R-smooth; the regularity ensures that irreducible closed subschemes of pure codimension 1 are Cartier; i.e., have an *invertible* ideal sheaf). This properness generally fails otherwise. For example, if X is a nodal plane cubic then $\mathbf{Pic}^0_{X/k}$ is a 1-dimensional torus.

Example 1.2. Suppose X is a smooth curve with genus g. In this case $J := \mathbf{Pic}^0_{X/k}$ is a proper k-group scheme whose tangent space coincides with

$$\mathrm{T}_0(\mathbf{Pic}_{X/k}) = \ker(\mathbf{Pic}_{X/k}(k[\epsilon]) \to \mathbf{Pic}_{X/k}(k)) = \mathrm{H}^1(X, \mathbf{O}_X),$$

so $\dim J \leq g$ with equality provided that J is smooth. The infinitesimal deformation theory of line bundles on a proper k-scheme is governed by coherent H^2's, and those vanish on curves. Thus, by the infinitesimal smoothness criterion, $\mathbf{Pic}_{X/k}$ is smooth and hence J is smooth.

One problem when working with the relative Picard functor $\mathrm{Pic}_{X/k}$ is that the sheafification involved in its definition makes it rather unclear what "concrete meaning" can be ascribed to its value on a k-scheme T. Namely, there is a natural map $\mathrm{Pic}(X_T)/\mathrm{Pic}(T) \to \mathbf{Pic}_{X/k}(T)$ and we will see shortly that it is always injective but surjectivity is unclear. Bijectivity will turn out to hold when $X(k) \neq \emptyset$, but if $X(k)$ is empty then there will be an obstruction in $\mathrm{Br}(k)$ to a k-rational point of the Picard scheme $\mathbf{Pic}_{X/k}$ to come from a line bundle on X. In terms of divisors, the issue is that one might have a divisor class on X_{k_s} that is $\mathrm{Gal}(k_s/k)$-invariant as a divisor class but contains no divisor that arises from one on X (due to a 2-cocycle in $\mathrm{Z}^2(\mathrm{Gal}(k_s/k), k_s^\times)$ that may fail to be a coboundary). In particular, if X is a smooth projective curve with no rational point (as will be of much interest below) then it is not evident what concrete meaning can be given to elements of $J(k)$ for the Jacobian $J = \mathbf{Pic}^0_{X/k}$.

To make this issue more explicit, note that (as for any k-scheme) $\mathrm{Pic}_{X/k}(k) = \mathrm{Pic}_{X/k}(k_s)^{\mathrm{Gal}(k_s/k)}$. Moreover, since $\mathrm{Pic}_{X/k}$ is a sheafification for the étale topology we know that $\mathrm{Pic}(X_{k_s}) \simeq \mathrm{Pic}_{X/k}(k_s)$. Thus, an element of $\mathrm{Pic}_{X/k}(k)$ is a line bundle $\mathscr{L}$ on X_{k_s} that is isomorphic to all of its $\mathrm{Gal}(k_s/k)$-twists. In contrast, to descend $\mathscr{L}$ to a line bundle on X we have to be able to choose these isomorphisms so that they satisfy a cocycle condition expressing compatibility with the multiplication in the Galois group (hence induce a Galois descent datum on $\mathscr{L}$).

Because X is geometrically integral and projective, the units on X_{k_s} we'd use to adjust an initial "random" choice of isomorphisms from $\mathscr{L}$ to its Galois twists all come from units in k_s. The obstruction to adjusting to get a 1-cocycle is a Galois 2-cocycle valued in $k_s^\times$, and defines an element in $\mathrm{Br}(k) = \mathrm{H}^2(k_s/k, k_s^\times)$ whose vanishing is necessary and sufficient for $\mathscr{L}$ to descend to X. Similarly, for a general k-scheme T the low-degree part of the Leray spectral sequence

$$\mathrm{H}^i(T, \mathrm{R}^j(f_T)_*(\mathbf{G}_{m,X_T})) \Rightarrow \mathrm{H}^{i+j}(X_T, \mathbf{G}_{m,X_T})$$

and the equality $\mathbf{G}_{m,T} = (f_T)_*(\mathbf{G}_{m,X_T})$ give an exact sequence:

$$0 \to \mathrm{Pic}(X_T)/\mathrm{Pic}(T) \to \mathbf{Pic}_{X/k}(T) \to \mathrm{H}^2(T, \mathbf{G}_{m,T}) =: \mathrm{Br}(T)$$

This "Brauer obstruction" is often nontrivial when k is a number field or a local field (except $\mathbf{C}$), whereas if $k = k_s$ or k is finite then the Brauer obstruction is trivial. Here's a concrete example:

Example 1.3. Let X be the conic $x^2 + y^2 + z^2 = 0$ in $\mathbf{P}^2_{\mathbf{R}}$, so $X_{\mathbf{C}}$ is isomorphic to $\mathbf{P}^1_{\mathbf{C}}$. For a general ring R a line bundle on $\mathbf{P}^n_R$ is given *Zariski-locally* on R by $\mathscr{O}(d)$ for some d, and the resulting fiberwise degree map $\deg : \mathbf{Pic}_{\mathbf{P}^n_k/k} \to \underline{\mathbf{Z}}_k$ is an isomorphism (as we can see by using the theorem on formal functions). Especially, the degree function is invariant under complex conjugation, so the point $\mathscr{O}(1) \in \mathbf{Pic}_{X/\mathbf{R}}(\mathbf{C})$ is an $\mathbf{R}$-rational point of the Picard scheme. On the other hand, X has no degree-1 line bundles, as any such would arises from an $\mathbf{R}$-point due to Riemann–Roch for the genus-0 curve X yet $X(\mathbf{R})$ is empty. The map $\mathrm{Pic}(X) \to \mathbf{Pic}_{X/\mathbf{R}}(\mathbf{R}) = \mathbf{Z}$ is the inclusion of $2\mathbf{Z}$ into $\mathbf{Z}$ (the line bundle $\Omega^1_{X/\mathbf{R}}$ has degree -2). Thus, the $\mathbf{R}$-point arising from a degree-1 line bundle on $X_{\mathbf{C}}$ must have nontrivial image in $\mathrm{Br}(\mathbf{R})$.

Now, if $f : X \to \mathrm{Spec}(k)$ has a section e (i.e. a k-rational point) then for any k-scheme T and line bundle $\mathscr{L}$ on X_T the line bundle $\mathscr{N} := \mathscr{L} \otimes f_T^*(e_T^*(\mathscr{L}))^{-1}$ has the same image in $\mathrm{Pic}(X_T)/\mathrm{Pic}(T)$, so also the same image in $\mathrm{Pic}_{X/k}(T)$, but it has an additional property: it admits a *trivialization along* e (i.e., an isomorphism $i : \mathscr{O}_T \simeq e_T^*(\mathscr{N})$). Such pairs $(\mathscr{M}, j)$ consisting of a line bundle $\mathscr{M}$ on X_T and a trivialization i of $\mathscr{M}$ along e are called *rigidified* line bundles and have a crucial advantage: for the evident notion of isomorphism among such pairs there are *no nontrivial automorphisms* (a simple exercise, using that $(f_T)_*(\mathbf{G}_{m,X}) = \mathbf{G}_{m,T}$). Consequently, in contrast with the set of isomorphism classes of line bundles, the functor $\mathrm{Pic}_{X/k,e}$ of *rigidified line bundles* (with respect to e) is a Zariski sheaf and even an fpqc sheaf.

The above construction of $(\mathscr{N}, i)$ from $\mathscr{L}$ readily shows (exercise!) that the natural forgetful map

$$\mathrm{Pic}_{X/k,e}(T) \to \mathrm{Pic}(X_T)/\mathrm{Pic}(T)$$

is an isomorphism for every T. Thus, the right side is already an étale (and even fpqc) sheaf, so the forgetful map $\mathrm{Pic}_{X/k,e} \to \mathrm{Pic}_{X/k}$ is an isomorphism of functors. Hence, in such cases we achieve the goal of giving a concrete meaning to the elements of $\mathrm{Pic}_{X/k}(T)$ (and in particular for any local k-scheme R we have $\mathrm{Pic}(X_R) = \mathrm{Pic}_{X/k}(R)$, the case $R = k$ being of much interest).

See Theorem 9.2.5 of "FGA Explained" for further discussion of this rigidification construction, which is applicable in particular to the case when $X = A$ is an abelian variety (in which case it is standard to take e to be the identity section). In such cases $A^\vee := \mathbf{Pic}^0_{A/k}$ is not only proper (since A is smooth) and connected but also *smooth*; this is the *dual abelian variety.* Going beyond the 1-dimensional case in Example 1.2, the smoothness of $A^\vee$ is remarkable in positive characteristic, as already for smooth projective surfaces X the Picard scheme $\mathbf{Pic}_{X/k}$ can be non-smooth when $\mathrm{char}\, k > 0$. Mumford's book "Lectures on curves on an algebraic surface" is largely about this phenomenon (which cannot occur in characteristic 0, by Cartier's theorem).

1.2 Polarizations

For this subsection, a good reference is Mumford's "Abelian Varieties," or notes and homework from Brian's course (see Tony's webpage).

For any line bundle $\mathscr{L}$ on an abelian variety A over a field k, define the map of k-schemes $\phi_{\mathscr{L}} : A \to A^\vee$ functorially via Mumford's construction: $x \mapsto t_x^* \mathscr{L} \otimes \mathscr{L}^{-1}$ on T-valued points for any k-scheme T. Here are some properties of this construction:

(a) The map $\phi_{\mathscr{L}}$ is a homomorphism (by the Theorem of the Square) and is symmetric with respect to double duality.

(b) For the Poincaré bundle $\mathscr{P}_A$, $(1, \phi_{\mathscr{L}})^*(\mathscr{P}_A) \simeq \mathscr{L}^{\otimes 2} \otimes \mathscr{M}$ with $\mathscr{M} \in \mathbf{Pic}^0_{A/k}(k) \subset \mathrm{Pic}(A)$.

(c) If $\mathscr{L}$ is ample then $\phi_{\mathscr{L}}$ is an isogeny with square degree.

(d) If $k = \overline{k}$ then every symmetric homomorphism $A \to A^\vee$ arises via the Mumford construction. (See §23 of Mumford's book for a proof. This uses crucially the structure of simple finite commutative group schemes over an algebraically closed field: $\mathbf{Z}/\ell\mathbf{Z}$ for an arbitrary prime ℓ, and also μ_p and α_p when $\mathrm{char}(k) = p > 0$.)

If $\ell \neq \mathrm{char}(k)$ and $f : A \to A^\vee$ is a homomorphism then the $\mathbf{Z}_\ell(1)$-valued bilinear form $e_{f,\ell}(x, y) = e_\ell(x, T_\ell(f)(y))$ on $T_\ell(A)$ is skew-symmetric if and only if f is symmetric (due to symmetry properties of the ℓ-adic Weil pairing with respect to swapping the role of A and $A^\vee$ via double-duality). This underlies the proof that $\phi_{\mathscr{L}}$ is always symmetric.

It is a general (nontrivial) fact that ampleness or not of a line bundle on A depends only on its geometric component in the Picard scheme (this is a special feature of abelian varieties, not true for a general geometrically integral projective k-scheme), so property (b) above shows that ampleness of $\mathscr{L}$ is equivalent to that of $(1, \phi_{\mathscr{L}})^*(\mathscr{P}_A)$. In view of property (d), we are thereby led to make the following definition over k that does not explicitly mention the Mumford construction and is in the spirit of considerations with positive definite quadratic forms over $\mathbf{Q}$ (when studied in the language of symmetric bilinear forms):

Definition 1.4. A homomorphism $\phi : A \to A^\vee$ is called a *polarization* when it is symmetric with respect to double duality and $(1, \phi)^*(\mathscr{P}_A)$ is ample on A. A polarization of degree 1 is called a *principal* polarization (these are special symmetric isomorphisms $A \simeq A^\vee$).

Over $\overline{k}$ we see via property (d) above that the polarizations are precisely the maps $\phi_{\mathscr{L}}$ for line bundles $\mathscr{L}$ on $A_{\overline{k}}$, so in general any polarization of A is an isogeny of square degree (as we may check over $\overline{k}$). By working a bit harder, we can improve a bit on $\overline{k}$: any polarization in the form $\phi_{\mathscr{L}}$ over k_s (a very useful improvement when k is not perfect).

The Mumford construction defines a map of k-group schemes

$$\mathbf{Pic}_{A/k} \to \mathbf{Hom}^{\mathrm{sym}}(A, A^\vee)$$

from the Picard scheme into the group scheme of symmetric homomorphisms (built into the Hom-scheme of the underlying projective schemes); this Hom-scheme is étale since we may check via the vanishing of its tangent space at the zero map via the rigidity of homomorphisms between abelian schemes. Consequently, this group scheme homomorphism kills the identity component. In Mumford's book it is shown that the condition $\phi_{\mathscr{L}} = 0$ says exactly that $\mathscr{L}$ arises from the dual abelian variety (as may be checked over $\overline{k}$), so we get a map

$$\mathrm{NS}(A) := \mathbf{Pic}_{A/k}/\mathbf{Pic}^0_{A/k} \to \mathbf{Hom}^{\mathrm{sym}}(A, A^\vee)$$

from the étale component group of $\mathbf{Pic}_{A/k}$ (called the *Néron–Severi group* of A) into the étale k-group of symmetric homomorphisms. But this is bijective on $\overline{k}$-points (injectivity being much simpler; surjectivity is property (d) above that ultimately rests on Proposition 1.7 below) and hence is an isomorphism of étale k-schemes. In particular, $\mathrm{NS}(A)$ corresponds to the "Galois lattice" of symmetric homomorphisms from A_{k_s} to $A^\vee_{k_s}$.

Remark 1.5. It is a general result of Lang and Néron that the étale Néron–Severi group $\mathbf{Pic}_{X/k}/\mathbf{Pic}^0_{X/k}$ is finitely generated on geometric points for any projective geometrically integral k-scheme X. This is called the Theorem of the Base; the case of abelian varieties is very special since we can "interpret" the finite generation via injection into a lattice of homomorphisms. Over $\mathbf{C}$ there is a proof for smooth X via the exponential sequence to inject the Néron–Severi group into $\mathrm{H}^2(X(\mathbf{C}), \mathbf{Z}(1))$. The Theorem of the Base was proved more generally by Grothendieck and Kleiman for any proper k-scheme; see Theorem 5.1 in Exposé XIII of SGA6 (whose proof uses étale cohomology and resolution of singularities for projective surfaces).

Consider the resulting exact sequence of smooth k-groups

$$0 \to A^\vee \to \mathbf{Pic}_{A/k} \to \mathrm{NS}(A) \to 0.$$

Polarizations of A are points $\phi \in \mathrm{NS}(A)(k) = \mathrm{Hom}^{\mathrm{sym}}(A, A^\vee)$ that satisfy an additional "positivity" condition (encoding ampleness of $(1,\phi)^*(\mathscr{P}_A)$), and those which arise from a (necessarily!) ample line bundle are the ones in the image of $\mathbf{Pic}_{A/k}(k) = \mathrm{Pic}(A)$. Hence, the obstruction to ϕ arising from a line bundle on A is the connecting map in the cohomology sequence

$$0 \to A^\vee(k) \to \mathrm{Pic}(A) \to \mathrm{Hom}^{\mathrm{sym}}(A, A^\vee) \xrightarrow{\delta} \mathrm{H}^1(k, A^\vee) \to \mathrm{H}^1(k, \mathbf{Pic}_{A/k}) \to \mathrm{H}^1(k, \mathrm{NS}(A)). \tag{1}$$

(Note that the second term really is $\mathrm{Pic}(A)$ because $A(k) \neq \emptyset$.) By Lang's theorem it follows that for finite fields k there is no obstruction: every ϕ arises from a line bundle over such k. Also, if k is *separably closed* then there is certainly no obstruction. Since k_s is exhausted by finite Galois extensions of k, we conclude:

Proposition 1.6. *For any polarization ϕ on A there exists a finite Galois extension K/k such that the polarization ϕ_K on A_K has the form $\phi_{\mathscr{L}}$ for some ample line bundle $\mathscr{L}$ on A_K.*

Now we can see the Galois-theoretic obstacle to realizing a general polarization as arising from a line bundle: for a finite Galois extension K/k and (ample) line bundle $\mathscr{L}$ on A_K, a Galois descent datum on $\phi_{\mathscr{L}}$ can define a polarization ϕ on A via Galois descent for homomorphisms, but such a descent datum is *much weaker* than a Galois descent datum on $\mathscr{L}$ itself (due to the $A^\vee(K)$-translation ambiguity in $\mathscr{L}$ when only $\phi_{\mathscr{L}}$ is given).

Remarkably, the obstruction class $\delta(\phi) \in \mathrm{H}^1(k, A^\vee)$ for a symmetric homomorphism ϕ to arise from a line bundle on A is always 2-torsion. To explain this, we recall that symmetry of a homomorphism $f : A \to A^\vee$ is equivalent to skew-symmetry of the associated self-pairing $e_{f,\ell} = e_\ell(\cdot, T_\ell(f)(\cdot))$ on $T_\ell(A)$ for a prime $\ell \neq \mathrm{char}(k)$. Hence, the 2-torsion property for the obstruction class is immediate from the following result (Theorem 2 in §20 of Mumford's book), as we will discuss at the start of §3.1:

Proposition 1.7. *If a homomorphism $f : A \to A^\vee$ is skew-symmetric with respect to the ℓ-adic Weil pairing then $2f = \phi_{\mathscr{L}_f}$ for $\mathscr{L}_f := (1, f)^*(\mathscr{P}_A)$.*

This result is analogous to an elementary fact in linear algebra: if $q : V \to F$ is a quadratic form on a vector space V over a field F and $B_q : V \times V \to F$ is the associated symmetric bilinear form $(v, v') \mapsto q(v+v') - q(v) - q(v')$ then for the associated linear map $L_q : V \to V^*$ defined by $v \mapsto B_q(v, \cdot) = B_q(\cdot, v)$ the composition of the evaluation map $V \times V^* \to F$ with $(\mathrm{id}, L_q) : V \to V \times V^*$ is $v \mapsto B_q(v, v) = 2q(v)$. Since Mumford's book is written over an algebraically closed field, we note that since Proposition 1.7 concerns an equality of maps given over k it is sufficient to check the equality after scalar extension to $\overline{k}$.

Remark 1.8. Though we speak of "polarizations" throughout this discussion. most of what we do will not use the ampleness property of a polarization; any symmetric isogeny ϕ (for example, one for which $(1,\phi)^*(\mathscr{P}_A)$ is anti-ample; i.e., the additive opposite of a polarization) would do, or sometimes even just a symmetric homomorphism.

Also, the way we've explained the relationship between symmetric isogenies and the Mumford construction is backwards: Proposition 1.7 is proved in §20 of Mumford's book and is *used* in §23 along with a study of possibility for the structure of the finite commutative group scheme $\ker f$ over an algebraically closed field to remove the factor 2 over such fields (i.e., to show that $\mathscr{L}_f$ is the square of another line bundle $\mathscr{N}$, so then $f = \phi_{\mathscr{N}}$).

The last topic we discuss in our background review is how to get a canonical principal polarization of the Jacobian J of a smooth proper and geometrically connected curve X of genus $g > 0$ over k. This will rest on an algebraic incarnation of a map defined by Abel for compact Riemann surfaces using integration. The starting point is the observation that we have a natural map

$$j : X \to \mathbf{Pic}^1_{X/k}$$

into the "degree 1" component of the Picard scheme via $x \mapsto \mathscr{O}(x)$ (associating to any $x \in X(T)$ the inverse of the invertible ideal sheaf of the section to $X_T \to T$ defined by the graph of x, for any k-scheme T). By the Riemann–Roch Theorem this is injective on geometric points (using that $g > 0$), and a mild refinement (exercise!) gives injectivity on artinian points, so j is a closed immersion.

In general $X(k)$ may be non-empty, but since X is smooth there is a finite Galois extension K/k such that $X(K)$ is non-empty. For such K/k and $x_0 \in X(K)$ we get a pointed map

$$i_{x_0} : X_K \to J_K = \mathbf{Pic}^0_{X_K/K}$$

defined functorially by $x \mapsto \mathscr{O}((x_0)_T) \otimes \mathscr{O}(x)^{-1} =: \mathscr{O}(x_0 - x)$ for any K-scheme T. If we replace x_0 with $x_0' \in X(K)$ then the map changes via translation against a point in $J(K)$ (namely, $i_{x_0'} = \xi + i_{x_0}$ for $\xi = \mathscr{O}(x_0' - x_0)$). But such translation induces the identity on $J^\vee$ since (on geometric points) the dual abelian variety consists of "translation-invariant line bundles". Hence, the induced map $\mathrm{Pic}^0(i_{x_0}) : J_K^\vee \to \mathbf{Pic}^0_{X_K/K} = J_K$ is *independent* of the choice of x_0! However, it is clear (why?) that if $\sigma \in \mathrm{Gal}(K/k)$ then $(\sigma^{-1})^*(i_{x_0}) = i_{\sigma(x_0)}$, so

$$(\sigma^{-1})^*(\mathrm{Pic}^0(i_{x_0})) = \mathrm{Pic}^0(i_{\sigma(x_0)}) = \mathrm{Pic}^0(i_{x_0});$$

Galois descent then implies that the K-homomorphism $\mathrm{Pic}^0(i_{x_0})$ descends to a *canonical* k-homomorphism $J^\vee \to J$ independent of all choices.

In the theory of Jacobian varieties it is shown that this final homomorphism is an *isomorphism* whose inverse $\phi_J : J \simeq J^\vee$ is a *polarization*. Explicitly, for K/k as above for which there exists $x_0 \in X(K)$, $(\phi_J)_K = \phi_{\mathscr{L}(\Theta_{x_0})}$ where $\Theta_{x_0} \subset J$ is the so-called "theta divisor" that is the image of $X_K^{g-1} \to J_K$ defined by $(x_1, \dots, x_{g-1}) \mapsto \sum_{i=1}^{g}(x_i - x_0)$, and $\mathscr{L}(\Theta_{x_0})$ is *ample*. (Beware that if we had used $x \mapsto \mathscr{O}(x - x_0)$ in the definition of i_{x_0}, which is to say we negated the definition of i_{x_0}, then the effect on ϕ_J would have been to introduce an overall negation, thereby getting the *negative* of a polarization. This is not just a matter of conventions, but is a real difference!) More elegantly, the theta divisor canonically lies in $\mathbf{Pic}^{g-1}_{X/k}$ as the image of X^{g-1} without any need for a base point over k, and when $\mathbf{Pic}^{g-1}_{X/k}$ contains a K-point (e.g., $\mathscr{O}(x_0)^{\otimes(g-1)}$ for $x_0 \in X(K)$) then translating by that brings the theta divisor into $\mathbf{Pic}^0_{X/k}$ as a well-defined divisor class.

In particular, if $X(k)$ is empty then there is no evident reason why ϕ_J should arises from a line bundle on J. Only in the latter case will it be generally true that the self-pairing on Ш(J) arising from the Cassels–Tate pairing via ϕ_J is alternating.

Remark 1.9. Here is a more geometric way to see that the above construction of ϕ_J is independent of x_0. For any k-scheme S and $x_0 \in X(S)$, we get a map $X_S \to J_S$ defined on points valued in an S-scheme T by $x \mapsto \mathscr{O}((x_0)_T - x)$. Applying $\mathbf{Pic}^0$ then gives a map of k-schemes

$$X \to \mathbf{Hom}(J^\vee, J)$$

where the target is the k-scheme of group homomorphisms (rather than general scheme morphisms). This Hom-scheme over k is étale by rigidity for abelian schemes yet X is geometrically connected over k, so this map must factor through a single k-point of the target. This proves the independence of x_0 and that it yields a (canonical) homomorphism $J^\vee \to J$ over k.

Remark 1.10. There is yet another geometric method to present the principal polarization, as follows. We built a canonical map $X \to \mathbf{Pic}^1_{X/k} =: P$ where P is a J-torsor. Applying $\mathbf{Pic}^0$ then gives a canonical map $\mathbf{Pic}^0_{P/k} \to J$. This latter map is what is shown to be an isomorphism whose inverse is a polarization via considerations over $\overline{k}$, but at a basic level how can we *canonically* identify $\mathbf{Pic}^0_{P/k}$ with $J^\vee$?

Rather generally, for *any* abelian variety A and A-torsor P we claim that *canonically* $A^\vee \simeq \mathbf{Pic}^0_{P/k}$. First observe that the group scheme $P' := \mathbf{Pic}^0_{P/k}$ is an abelian variety, as it suffices to check this over $\overline{k}$ (where the A-torsor P becomes trivial, so we can appeal to the theory of the dual abelian variety). For any k-scheme T and $\xi \in P(T)$ we get an isomorphism of T-schemes $f_\xi : A_T \simeq P_T$ defined functorially by $a \mapsto a.\xi$ (on points valued in any T-scheme), and hence an isomorphism $\mathbf{Pic}^0(f_\xi) : P'_T \to A_T^\vee$ between abelian schemes over T. This construction is compatible with base change on T and so defines a map of k-schemes

$$P \to \mathbf{Isom}(P', A^\vee).$$

But the target is étale and the source is geometrically connected over k, so it factors through a single k-point of the target; this is a distinguished k-isomorphism $f : P' \simeq A^\vee$.

For any K/k such that $P'(K)$ contains a point ξ we see that the canonical f_K is exactly $\mathrm{Pic}^0(f_\xi)$. The independence of ξ is explained more concretely by the fact that geometric points of $A^\vee$ are translation-invariant line bundles, and this concrete description shows (check!) that f is an isomorphism of $A^\vee$-torsors.

2 Definitions and Formal Properties of the Cassels–Tate Pairing

For an abelian variety A over a global field k we'd like to define a bilinear pairing (the Cassels–Tate pairing)

$$\langle\ ,\ \rangle : Ш(A) \times Ш(A^\vee) \to \mathbf{Q}/\mathbf{Z}$$

with the following properties:

- The kernels on left and right are the respective maximal divisible subgroups (especially, if $\#Ш < \infty$ then $\langle\cdot,\cdot\rangle$ is non-degenerate).
- If we pull back $\langle\cdot,\cdot\rangle$ to a self-pairing on $Ш(A)$ via a polarization (or any symmetric homomorphism) then the result is skew-symmetric (not necessarily alternating!).
- Functoriality holds with respect to any k-homomorphism $f : A \to B$ and its dual (as described and used in Lecture 6): $\langle f(a), b\rangle = \langle a, f^\vee(b)\rangle$.

There are maybe a half-dozen ways to define this pairing. Each of them has its own appeal and purpose, and to obtain all desired properties one needs to prove equality among (some of) the various definitions.

Remark 2.1. One can find a "fancy" definition of the pairing in Chapter II, Section 5 of Milne's book "Arithmetic Duality Theorems" (ADT). As Milne notes, we should view that as an étale cohomology version of the definition described in Tate's 1962 ICM talk, which is in turn analogous to Tate's global duality result (*loc cit.* or Theorem 4.10 in Chapter I of ADT) on finite discrete Galois modules. In particular, the ICM definition is essentially the "Weil pairing" definition of Poonen and Stoll, namely the one used in the proof of Proposition 6.9 in ADT. We don't give a proof here (yet?) of these compatibilities.

Poonen and Stoll use four definitions: the "homogeneous space" definition, the "Weil pairing" definition, the "Albanese–Picard" definition, and the "Albanese–Albanese" definition. Both of the latter two are quite painful, involving pushing zero-cycles around and evaluating rational functions at them, so we don't give them here. Poonen and Stoll use the Albanese–Picard definition crucially to compute the value of the pairing in the case that A is the Jacobian of a curve X. We'll black-box some of that computation, referring to Poonen–Stoll for full details. The Appendix of Poonen–Stoll proves that these four definitions coincide.

The Weil pairing definition is slightly less complicated, but we won't give it either.[1] An advantage of the Weil pairing definition is that the Cassels–Tate pairing defined in this way inherits most of the basic desired properties (bilinearity, anti-symmetry with respect to double-duality, and functoriality) directly from the known analogues for the Weil pairing. The non-degeneracy based on this definition is a bit more complicated to prove (see Chapter II, Section 6 of ADT, which uses Tate's 9-term exact sequence from Lecture 4).

Soon we'll give the homogeneous space definition, but first we want to see that the obstruction to alternation of a skew-symmetric pairing on $Ш(A)_{\mathrm{nd}}$ is a formal consequence of the desired properties when $Ш(A)_{\mathrm{nd}}$ is finite. The real miracle then will be that Poonen and Stoll succeed in describing such an obstruction directly without assuming finiteness of $Ш(A)_{\mathrm{nd}}$, and they use their description of it to perform calculations.

Consider quite generally a $\mathbf{Q}/\mathbf{Z}$-valued non-degenerate skew-symmetric pairing $\langle\cdot,\cdot\rangle$ on a *finite* abelian group G. The map $a \mapsto \langle a, a\rangle$ is easily checked to be a homomorphism via the skew-symmetry. Hence, by perfectness, there exists a unique $c \in G$ such that this homomorphism has the form $a \mapsto \langle a, c\rangle$. Also due to perfectness, this is the zero homomorphism (i.e. $\langle\cdot,\cdot\rangle$ is alternating) if and only if $c = 0$.

One can use the structure theorem of finite abelian groups and work separately on each primary factor to show that if $\langle\cdot,\cdot\rangle$ is alternating then $\#G$ is a square. On odd-primary parts skew-symmetry and alternation

[1] Perhaps it will be given in a later version of these notes.

are the same thing, so in general $\#G$ is either a square or twice a square. A slightly more involved argument (see Theorem 8 in Poonen–Stoll) using just elementary algebra shows that $\langle c, c\rangle = 0$ if and only if $\#G$ is actually a square. The argument shows that if $c \neq 0$ but $\langle c, c\rangle = 0$ then we can find square-sized subgroup V of G containing c such that $\langle\cdot,\cdot\rangle$ is alternating on $V^{\perp}$. The phenomenon that $c \neq 0$ but $\langle c, c\rangle = 0$ does sometimes happen for the Cassels–Tate self-pairing $\langle\cdot,\cdot\rangle_\lambda$ on $\text{Ш}(A)$ where c is the Poonen–Stoll obstruction c_λ to alternation for a principally polarized (A, λ).

3 Homogeneous Space Definition and Counterexamples

At last, here is the homogeneous space definition of $\langle\cdot,\cdot\rangle_A$, which is the most palatable among the various definitions. Choose $a \in \text{Ш}(A)$ and $a' \in \text{Ш}(A^\vee)$. Let P be an A-torsor representing a, so P has a rational point over every completion k_v. By Remark 1.10, we have canonically $\mathbf{Pic}^0_{P/k} \simeq A$. This does *not* identify $A(k)$ with a subgroup of $\text{Pic}(P)$ since we have no concrete description of what points in $\mathbf{Pic}^0_{P/k}(k)$ mean! Nonetheless,

$$a' \in \text{Ш}(\mathbf{Pic}^0_{P/k}) \subset \mathrm{H}^1(k, \mathbf{Pic}^0_{P/k}).$$

In general $\mathbf{Pic}_{P/k}(k)$ receives a homomorphism from $\text{Pic}(P)$ but we cannot say much more.

Let $j : \eta = \text{Spec}(k(P)) \hookrightarrow P$ be the inclusion of the generic point, so there is a short exact sequence of étale sheaves on P

$$1 \to \mathbf{G}_{m,P} \to j_*(\mathbf{G}_{m,P}) \to \mathscr{D}iv \to 0$$

where the cokernel term is the sheaf of Weil divisors (surjectivity expressing that they are locally principal, even for the Zariski topology, as P is regular!); the value of $\mathscr{D}iv$ on an étale P-scheme U is the group of *principal* Weil divisors on U. Applying higher direct images relative to $f : P \to \text{Spec}(k)$ for the étale topology gives an exact sequence of étale sheaves on k:

$$1 \to \mathbf{G}_{m,k} \to (f \circ j)_*(\mathbf{G}_{m,\eta}) \to f_*(\mathscr{D}iv) \to \mathrm{R}^1 f_*(\mathbf{G}_{m,P}) = \mathbf{Pic}_{P/k}. \tag{2}$$

We now make this more explicit for the corresponding exact sequence of k_s-points as discrete $\text{Gal}(k_s/k)$-modules. Note that the natural map $\text{Pic}(P_{k_s}) \to \mathbf{Pic}_{P/k}(k_s)$ is an isomorphism since $P(k_s) \neq \emptyset$ (as P inherits smoothness from A), and $f_*(\mathscr{D}iv)$ corresponds to the group $\text{Div}(P_{k_s})$ of Weil divisors on P_{k_s}. But every line bundle on P_{k_s} does arise from a global Weil divisor on P_{k_s} (as for any regular scheme), so (2) is actually surjective at the end. In other words, we recover the classical 4-term exact sequence of discrete $\text{Gal}(k_s/k)$-modules

$$1 \to k_s^\times \to k_s(P)^\times \overset{\text{div}}{\to} \text{Div}(P_{k_s}) \to \text{Pic}(P_{k_s}) \to 1.$$

Inside $\text{Pic}(P_{k_s})$ it makes sense to *define*

$$\text{Pic}^0(P_{k_s}) := \mathbf{Pic}^0_{P/k}(k_s).$$

(This has no "degree 0" interpretation when $\dim P > 1$; it is defined via the topology of the Picard scheme.) Inside the group $\text{Div}(P_{k_s})$ of Weil divisors on P_{k_s}, we define $\text{Div}^0(P_{k_s})$ to be those divisors whose associated line bundle lies in $\text{Pic}^0(P_{k_s})$. Thus, replacing $\text{Pic}(P_{k_s})$ with $\text{Pic}^0(P_{k_s})$ and replacing $\text{Div}(P_{k_s})$ with its subgroup $\text{Div}^0(P_{k_s})$ gives a 4-term exact sequence

$$1 \to k_s^\times \to k_s(P)^\times \overset{\text{div}}{\to} \text{Div}^0(P_{k_s}) \to \text{Pic}^0(P_{k_s}) \to 1$$

that we chop into two short exact sequences of discrete $\text{Gal}(k_s/k)$-modules

$$0 \to k_s^\times \to k_s(P)^\times \to k_s(P)^\times/k_s^\times \to 0, \;\; 0 \to k_s(P)^\times/k_s^\times \to \text{Div}^0(P_{k_s}) \to \text{Pic}^0(P_{k_s}) \to 0.$$

This provides a connecting map

$$\delta : \mathrm{H}^1(k, \mathbf{Pic}^0_{P/k}) \to \mathrm{H}^2(k, k_s(P)^\times/k_s^\times)$$

and an exact sequence

$$\mathrm{H}^2(k, k_s(P)^\times) \to \mathrm{H}^2(k, k_s(P)^\times/k_s^\times) \overset{\delta'}{\to} \mathrm{H}^3(k, \mathbf{G}_m).$$

Up to here there has been no number theory. Now class field theory enters: the cohomological approach to global class field theory shows that $\mathrm{H}^3(k, \mathbf{G}_m) = 1$ for the global field k, so $\delta(a')$ lifts to an element $f' \in \mathrm{H}^2(k, k_s(P)^\times)$ well-defined up to the image of $\mathrm{H}^2(k, \mathbf{G}_m) = \mathrm{Br}(k)$. The Galois groups $\mathrm{Gal}(k_s/k)$ and $\mathrm{Gal}(k_s(P)/k(P))$ are naturally identified, and in this manner (using Hilbert 90) $\mathrm{H}^2(k, k_s(P)^\times)$ is identified with the subgroup of $\mathrm{Br}(k(P))$ consisting of classes that split over $k_s(P)$. We now use Grothendieck's work on Brauer groups of schemes to analyze when the natural map $\mathrm{Br}(k) \to \mathrm{H}^2(k, k_s(P)^\times) \subset \mathrm{Br}(k(P))$ is injective.

Grothendieck proved (see Corollary 2.6 in Chapter IV of Milne's book on étale cohomology) that for any regular connected scheme Z with generic point η_Z, the natural map of étale cohomology groups

$$\mathrm{Br}(Z) := \mathrm{H}^2(Z, \mathbf{G}_m) \to \mathrm{H}^2(\eta_Z, \mathbf{G}_m) = \mathrm{Br}(\kappa(\eta_Z))$$

is injective. Applying this with $Z = P$ and using functoriality of cohomology relative to the factorization

$$\eta \to P \to \mathrm{Spec}(k)$$

shows that injectivity of $\mathrm{Br}(k) \to \mathrm{Br}(k(P))$ is *equivalent* to that of $\mathrm{Br}(k) \to \mathrm{Br}(P)$. If there were a k-point of P, which is to say a section to $P \to \mathrm{Spec}(k)$, then the latter map of Brauer groups would obviously be injective. Of course, generally $P(k)$ is empty (after all, the whole point of the present discussion is to grapple with the possibility that $a' \neq 0$), but by hypothesis $P(k_v)$ is non-empty for all v. Hence, by the same considerations applied over k_v rather than over k we conclude that the local analogue $\mathrm{Br}(k_v) \to \mathrm{Br}(k_v(P))$ *is* injective for all v.

To see why the "local injectivity" is interesting, note that the local restrictions a'_v vanish for all v (as $a' \in Ш(A^\vee)$). Hence, by the functoriality of our purely algebraic considerations relative to extension of the ground field, the local restriction $f'_v \in \mathrm{H}^2(k_v, k_{v,s}(P)^\times) \subset \mathrm{Br}(k_v(P))$ maps to zero in $\mathrm{H}^2(k_v, k_{v,s}(P)^\times/k_{v,s}^\times)$ for all v. It follows that f'_v comes from some $c_v \in \mathrm{Br}(k_v)$ for each v, and such c_v is *unique* for each v. Consequently, we get a collection

$$(c_v)_v \in \prod_v \mathrm{Br}(k_v)$$

that is actually *well-defined* up to precisely the image of $\mathrm{Br}(k)$ (the ambiguity in the global class f'). We need $(c_v)_v$ to lie inside the direct sum rather than just the direct product, so now we prove:

Lemma 3.1. *Any element $\xi \in \mathrm{H}^2(k, k_s(P)^\times) \subset \mathrm{Br}(k(P))$ has trivial restriction $\xi_v \in \mathrm{Br}(k_v(P))$ for all but finitely many v. In particular, for all but finitely many v the local restriction $f'_v \in \mathrm{Br}(k_v(P))$ vanishes and so $c_v = 0$.*

Proof. [2] □

Local class field theory provides natural isomorphisms $\mathrm{inv}_v : \mathrm{Br}(k_v) \simeq \mathbf{Q}/\mathbf{Z}$ for all non-archimedean v and a unique obvious inclusion for archimedean v with the property that the resulting map

$$\mathrm{inv} : \bigoplus_v \mathrm{Br}(k_v) \to \mathbf{Q}/\mathbf{Z}$$

defined by $\sum_v \mathrm{inv}_v$ kills $\mathrm{Br}(k)$. Hence, the $\mathrm{Br}(k)$-ambiguity in the definition of (c_v) is wiped out by applying inv, so it is *well-posed* to define

$$\langle a, a' \rangle_A = \sum_v \mathrm{inv}(c_v) \in \mathbf{Q}/\mathbf{Z}.$$

That is the "homogenous space" definition of the Cassels–Tate pairing.

As with most definitions of the pairing, this definition is rather asymmetric in how it treats A and $A^\vee$, so the anti-symmetry of the pairing relative to swapping the roles of A and $A^\vee$ and using double duality isn't obvious. The dependence on a via the torsor P relates additivity in a to "addition" of torsors, so with some care one can verify additivity in a. Additivity in a' is immediate, so overall we see that bilinearity holds. Functoriality in A is a bit tricky, but can be shown just via careful work with the definition. The behavior on divisible elements and non-degeneracy of the resulting pairing between $Ш(A)_{\mathrm{nd}}$ and $Ш(A^\vee)_{\mathrm{nd}}$ are completely opaque via this definition.

[2] Needs to be justified. Does not appear to be addressed in ADT or Poonen–Stoll.

3.1 Explicit Description of c

Assuming now the properties of the Cassels–Tate pairing and agreement among multiple definitions, our goal is to reduce the study of the alternation of the pairing and the size of $\text{Ш}(A)_{\text{nd}}$ (if finite!) for principally polarized A to the study of a single element of $\text{Ш}(A)$ (depending on the principal polarization), which in turn can be computed explicitly in some cases. In particular, we want to find the c that, thanks to those properties, we know must exist if $\text{Ш}(A)_{\text{nd}}$ is finite (and A admits a principal polarization).

The starting point is the cohomology sequence (1). For any polarization (or merely symmetric homomorphism) $\lambda : A \to A^\vee$, define c_λ to be the image of λ under the connecting homomorphism

$$\delta : \text{Hom}^{\text{sym}}(A, A^\vee) = \text{NS}(A)(k) \to \text{H}^1(k, A^\vee).$$

By exactness, $c_\lambda = 0$ if and only if λ has the form $\phi_{\mathscr{L}}$ for some line bundle $\mathscr{L}$ on A. By Proposition 1.7, 2λ does arise via the Mumford construction. Thus, $2c_\lambda = c_{2\lambda} = 0$. Up to here, the construction of $c_\lambda \in \text{H}^1(k, A^\vee)[2]$ is pure algebraic geometry, valid over any field. Note that the formation of c_λ is compatible with extension of the ground field. Number theory enters in the following result:

Lemma 3.2. *If k is a (possibly archimedean) local field then $c_\lambda = 0$. In particular, if k is a global field then $c_\lambda \in \text{Ш}(A^\vee)$.*

Proof. The class $c_\lambda \in \text{H}^1(k, A^\vee)$ maps to zero in $\text{H}^1(k, \text{Pic}_{A/k})$ due to how it was constructed, Thus, by the commutativity of the diagram of long exact ohomology sequences arising from the commutative diagram of short exact sequences

$$\begin{array}{ccccccccc} 0 & \longrightarrow & k_s(A)^\times/k_s^\times & \longrightarrow & \text{Div}^0(A_{k_s}) & \longrightarrow & A^\vee(k_s) & \longrightarrow & 0 \\ & & \| & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & k_s(A)^\times/k_s^\times & \longrightarrow & \text{Div}(A_{k_s}) & \longrightarrow & \text{Pic}(A_{k_s}) & \longrightarrow & 0 \end{array}$$

it follows that c_λ is killed by the connecting map $\text{H}^1(k, A^\vee) \to \text{H}^2(k, k_s(A)^\times/k_s^\times)$. Hence, suffices to show that this connecting map is injective, or equivalently (by long exactness for the cohomology sequence arising from the top row) that the natural map $\text{H}^1(k, \text{Div}^0(A_{k_s})) \to \text{H}^1(k, A^\vee)$ vanishes.[3] □

To go further we want to compute c_λ using homogeneous spaces. The argument is typical of those in Poonen–Stoll in that it involves pushing around divisors and cocycles, so we give this proof and omit the others. Before we state the result, we need to extend the Mumford construction to the setting of line bundles on A-torsors. That is, for a line bundle $\mathscr{L}$ on an A-torsor P, we explain how to make sense of $\phi_{\mathscr{L}}$ as a symmetric homomorphism $A \to A^\vee$.

Note that the Mumford construction $a \mapsto t_a^*(\mathscr{L}) \otimes \mathscr{L}^{-1}$ gives a meaningful scheme map $A \to \mathbf{Pic}_{P/k}$ since the formula only involves A throught its action on P, and this map carries $a = 0$ to the trivial line bundle, so for topological reasons it carries A into $\mathbf{Pic}^0_{P/k}$ respecting identity sections. Thus, this is a map of abelian varieties $A \to \mathbf{Pic}^0_{P/k}$ over k. But the target is canonically identified with $A^\vee$ as we saw in Remark 1.10! This map $\phi_{\mathscr{L}}$ is symmetric because we can check it over $\overline{k}$ (where P becomes a trivial torsor upon choosing a base point in $P(\overline{k})$, in terms of which this construction becomes the usual Mumford construction that is always a symmetric homomorphism).

Proposition 3.3. *Let $\lambda : A \to A^\vee$ be a symmetric homomorphism and P an A-torsor such that $\lambda = \phi_{\mathscr{L}}$ for a line bundle $\mathscr{L} \in \text{Pic}(P)$. Then c_λ is the image of $[P]$ under $\text{H}^1(\lambda) : \text{H}^1(k, A) \to \text{H}^1(k, A^\vee)$.*

We shall express the proof in the language of line bundles and the Picard group $\text{Pic}(A_{k_s})$ of isomorphism classes of line bundles rather than the more explicit language of divisors and the group $\text{Div}(A_{k_s})/k_s(A)^\times$ of divisor classes (as used by Poonen–Stoll). Much as the definition of c_λ as an element of $\text{H}^1(k, A^\vee)$ (without reference to $\text{Ш}(A^\vee)$) involved no number theory and made sense over any field, the statement of Proposition 3.3 also makes sense over any field and the proof will work in that generality (i.e., there is no number theory in the proof, just basic algebraic geometry).

[3] Needs to be finished, in a later revision.

Proof. Since P is an A-torsor, for any two points $z_1, z_2 \in P(k_s)$ the "difference" $z_1 - z_2 \in A(k_s)$ makes sense as the unique $a \in A(k_s)$ such that a carries z_2 to z_1. In the same manner, if $\mathscr{N}$ is a line bundle on P and $z \in P(k_s)$ then we can make sense of the line bundle $\mathscr{N}_z = f_z^*(\mathscr{N}_{k_s})$ on A_{k_s} where $f_z : A_{k_s} \simeq P_{k_s}$ is the torsor isomorphism via the base point z. For any $a \in A(k_s)$ we have $f_{a+z} = f_z \circ t_a$, so $\mathscr{N}_{a+z} \simeq t_a^*(\mathscr{N}_z)$. For the calculation we are about to carry out, the significance of $\mathscr{N}$ arising on P rather than just on P_{k_s} is that $(\sigma^{-1})^*(\mathscr{N}_z) = \mathscr{N}_{\sigma(z)}$ for $\sigma \in \mathrm{Gal}(k_s/k)$.

[As a sanity check, the reason for σ^{-1} appearing on the left side is forced by both functoriality and geometric reasons: (i) it cancels out the contravariance of pullback and so corresponds to making a *left* action of $\mathrm{Gal}(k_s/k)$ on $\mathrm{Pic}(A_{k_s})$ in accordance with the habitual convention in the theory of modules for a group, or equivalently it is the pushforward σ_* along the automorphism action of σ on the scheme A_{k_s}, (ii) it corresponds to making $\mathrm{Div}(A_{k_s}) \twoheadrightarrow \mathrm{Pic}(A_{k_s})$ Galois-equivariant since σ^* applied to the line bundle for a divisor is the line bundle for the σ-*preimage* of that divisor and σ-preimage is the same as image under σ^{-1}.]

Now comes the cocycle calculation. Choose $z \in P(k_s)$. Following the conventions as in §5.3 of Chapter I of Serre's book on Galois cohomology (look at the proof of Proposition 33), the class $[P] \in \mathrm{H}^1(k, A)$ is the class of the cocycle $\sigma \mapsto z^\sigma - z$. (Warning: this may look like it's a 1-coboundary for $A(k_s)$ but that is erroneous because z and z^σ are not points in $A(k_s)$, whereas "$z^\sigma - z$" is.) The consistency of $\phi_{\mathscr{L}}$ and the Mumford construction over k_s upon trivializing P_{k_s} via the base point z identifies λ_{k_s} with $\phi_{\mathscr{L}_z}$. Thus, c_λ is represented by the 1-cocycle

$$\sigma \mapsto \delta(\mathscr{L}_z)(\sigma) = (\sigma^{-1})^*(\mathscr{L}_z) \otimes \mathscr{L}_z^{-1} \simeq \mathscr{L}_{\sigma(z)} \otimes \mathscr{L}_z^{-1} \simeq t^*_{\sigma(z)-z}(\mathscr{L}_z) \otimes \mathscr{L}_z^{-1}$$

valued in $A^\vee(k_s)$. But this is $\phi_{\mathscr{L}_z}(z^\sigma - z) = \lambda_{k_s}(z^\sigma - z)$, so we have the class of $\mathrm{H}^1(\lambda)([P]) \in \mathrm{H}^1(k, A^\vee)$. □

Corollary 3.4. *If $A = J$ is the Jacobian of a curve X with genus $g > 0$ and canonical principal polarization λ then $c_\lambda = \mathrm{H}^1(\lambda)([\mathbf{Pic}^{g-1}_{X/k}])$. In particular, $\mathrm{H}^1(\lambda)([\mathbf{Pic}^{g-1}_{X/k}]) \in Ш(J)$.*

The statement and proof of the formula for c_λ will work over any field; only the membership in $Ш(J)$ has arithmetic content.

Proof. The theta divisor exists over k as a canonical divisor D on the J-torsor $P = \mathbf{Pic}^{g-1}_{X/k}$, and if $x_0 \in X(k_s)$ then for $z = \mathscr{O}(x_0)^{\otimes(g-1)} \in P(k_s)$ we see that the isomorphism $f_z : J_{k_s} \simeq P_{k_s}$ via the base point z carries D back to the classical theta divisor $\Theta_{x_0} \subset A_{k_s}$ associated to x_0. The definition of λ in terms of Θ_{x_0} thereby shows that λ coincides with the homomorphism $\phi_{\mathscr{L}(D)}$ is associated to the line bundle $\mathscr{L}(D)$ arising from the divisor D on P. Proposition 3.3 may now be applied to obtain the formula for c_λ, and the assertion concerning $Ш(J)$ is a consequence of Lemma 3.2. □

Using similar argumentation (i.e., cocycle manipulation) and the homogeneous space definition, one proves the following result (Theorem 5 in Poonen–Stoll):

Theorem 3.5. *For every $a \in Ш(A)$ and a symmetric k-homomorphism $\lambda : A \to A^\vee$,*

$$\langle a, \lambda a + c_\lambda \rangle_A = \langle a, \lambda a - c_\lambda \rangle_A = 0.$$

Hence, c_λ as we've constructed it satisfies $\langle a, \lambda a \rangle_A = \langle a, c_\lambda \rangle_A$. For λ of degree 1 it follows that the class $c = \lambda^{-1} c_\lambda \in Ш(A)$ (which depends on λ) makes sense and its image in $Ш(A)_{\mathrm{nd}}$ is the obstruction to alternation of the self-pairing on $Ш(A)_{\mathrm{nd}}$ induced by the Cassels–Tate construction when $Ш(A)_{\mathrm{nd}}$ is finite.

3.2 Counterexamples to alternation

In this subsection, we assume $A = J = \mathbf{Pic}^0_{X/k}$ for a genus g curve X, and we fix λ to be the canonical principal polarization for turning the Cassels–Tate pairing into a skew-symmetric form

$$\langle \cdot, \cdot \rangle_X : Ш(J) \times Ш(J) \to \mathbf{Q}/\mathbf{Z}.$$

Using Corollary 3.4, we know that the class $[\mathbf{Pic}^{g-1}_{X/k}] \in \mathrm{H}^1(k, J)$ lies in $Ш(J)$ and is the obstruction to the alternating property for $\langle \cdot, \cdot \rangle_X$; likewise, the pairing of this class against itself is the obstruction to $\#Ш(J)_{\mathrm{nd}}$ being a square (when finite).

By computing using the Albanese–Picard definition, Poonen and Stoll show that for any $n \in \mathbf{Z}$ such that $[\mathbf{Pic}^n_{X/k}] \in Ш(J)$ (as always happens for $n = g-1$),

$$\langle [\mathbf{Pic}^n_{X/k}], [\mathbf{Pic}^n_{X/k}] \rangle_X = \frac{N}{2} \bmod \mathbf{Z} \in \mathbf{Q}/\mathbf{Z} \tag{3}$$

where N is the number of places of k for which the J-torsor $\mathbf{Pic}^n_{X/k}$ does *not* have a local point.

Remark 3.6. To see that N is finite, it suffices to show that $\mathrm{Pic}^1_{X/k}(k_v)$ is non-empty for all but finitely many v, or even better that $X(k_v)$ is non-empty for all but finitely many v. Consider any non-archimedean place v of k for which X_{k_v} admits a smooth proper model $\mathfrak{X}_v$ over O_{k_v}. By the Riemann Hypothesis for curves,

$$|\mathfrak{X}_v(\kappa(v)) - q_v - 1| \le 2g q_v^{1/2}$$

where q_v is the size of the residue field at v. Hence, for $q_v \ge 4g^2$ (which rules out only finitely many v), it follows that $\mathfrak{X}_v(\kappa(v))$ must be non-empty. Any such $\kappa(v)$-point lifts to an O_{k_v}-point of $\mathfrak{X}_v$, due to the Zariski-local structure theorem for smooth and étale morphisms [EGA, IV$_4$, 17.12.2(d), 18.4.6(ii)] (usually incorrectly called "Hensel's Lemma"). Passing to the generic fiber gives a point in $X(k_v)$ for such v.

Combining Corollary 3.4 with (3) for $n = g-1$, it follows that $\#Ш(J)_{\mathrm{nd}}$ is a square (assuming it is finite) if and only if the number N of places v for which X_{k_v} has no Weil divisor of degree $g-1$ is even. Likewise, $\langle \cdot, \cdot \rangle_X$ is alternating if and only if X has a divisor of degree $g-1$ (i.e., the J-torsor $\mathbf{Pic}^{g-1}_{X/k}$ is trivial).

As applications of these conclusions, Poonen and Stoll do the following:

(i) They give some density computations to show that a positive proportion of Jacobians (of a specific type of hyperelliptic curve) have Ш with non-square size if finite.

(ii) They build a genus-2 curve whose Jacobian has non-alternating Cassels–Tate pairing but whose Tate–Shafarevich group has provably finite size that is a square!

We finish by discussing the family in (i).

For even $g > 0$ and any $t \in \mathbf{Q}$, let X_t be the genus-g hyperelliptic curve over $\mathbf{Q}$ given by

$$y^2 = -(x^{2g+2} + x + t). \tag{4}$$

(Note that X_t is *not* the Zariski closure in $\mathbf{P}^2$ of this smooth affine plane curve; such a closure is not smooth!) Let $J_t = \mathbf{Pic}^0_{X_t/\mathbf{Q}}$ be its Jacobian. Using computations resting on Hensel's Lemma, for every prime p one finds that $(X_t)_{\mathbf{Q}_p}$ has a divisor of $\mathbf{Q}_p$-degree $g-1$. The map $x : X_t \to \mathbf{P}^1_{\mathbf{Q}}$ has degree 2, so its fiber over any $\mathbf{Q}$-point is a $\mathbf{Q}$-point in $\mathbf{Pic}^2_{X_t/\mathbf{Q}}$. Thus, whether or not the J_t-torsor $\mathbf{Pic}^n_{X_t/\mathbf{Q}}$ has a point over a given extension of $\mathbf{Q}$ (such as a completion of $\mathbf{Q}$) depends only on the parity of n. In particular, since $g-1$ is odd, the existence of a divisor on $(X_t)_{\mathbf{Q}_p}$ with $\mathbf{Q}_p$-degree $g-1$ is equivalent to $(X_t)_{\mathbf{Q}_p}$ admitting a (possibly non-effective!) divisor of $\mathbf{Q}_p$-degree equal to 1.

We now show that $\mathbf{Pic}^{g-1}_{X_t/\mathbf{Q}}(\mathbf{R})$ is empty for suitable t, in which case the J_t-torsor $\mathbf{Pic}^{g-1}_{X_t/\mathbf{Q}}$ fails to have a local point at exactly one place of $\mathbf{Q}$ (namely, the archimedean place). The key point is to use special features of $\mathbf{R}$ to relate the existence of an odd-degree divisor to the existence of an *effective* divisor of degree 1 (i.e., an $\mathbf{R}$-point of the curve X_t).

Assume $\mathbf{Pic}^n_{X_t/\mathbf{Q}}(\mathbf{R})$ is non-empty for some odd n, so there exist line bundles on $(X_t)_{\mathbf{R}}$ with arbitrarily large odd $\mathbf{R}$-degree. Choosing one such $\mathscr{L}$ that is very ample yields a closed immersion

$$(X_t)_{\mathbf{R}} \hookrightarrow \mathbf{P}(\Gamma((X_t)_{\mathbf{R}}, \mathscr{L})) = \mathbf{P}^N_{\mathbf{R}}$$

as a curve with odd degree $d > 0$ in some projective space over $\mathbf{R}$. A "generic" projection to a line in $\mathbf{P}^N_{\mathbf{R}}$ is a finite map $f : (X_t)_{\mathbf{R}} \to \mathbf{P}^1_{\mathbf{R}}$ whose fibers are scheme-theoretic intersections of $(X_t)_{\mathbf{R}}$ with hyperplanes in $\mathbf{P}^N_{\mathbf{R}}$, so $\deg(f) = d$. The fiber of f over a point in $\mathbf{P}^1(\mathbf{R})$ outside the finite branch locus of f is an étale $\mathbf{R}$-scheme in $(X_t)_{\mathbf{R}}$ with $\mathbf{R}$-degree d that is odd. This étale fiber cannot be a union of copies of $\mathrm{Spec}(\mathbf{C})$ (as otherwise its $\mathbf{R}$-degree would be even), so it must contain an $\mathbf{R}$-point. In other words, the existence of some divisor on $(X_t)_{\mathbf{R}}$ with odd $\mathbf{R}$-degree implies that there exists an $\mathbf{R}$-point! (This is specific to the field $\mathbf{R}$.)

Since $(X_t)_{\mathbf{R}}$ is $\mathbf{R}$-smooth with pure dimension 1, by the analytic inverse function theorem over $\mathbf{R}$ it follows that $X_t(\mathbf{R})$ is infinite when it is non-empty. In particular, the affine open subscheme of X_t given by (4) has a solution $(x_0, y_0) \in \mathbf{R}^2$ in such cases. This is impossible if we choose the rational t to be positive.