Overview

Brian Conrad and Akshay Venkatesh *

October 7, 2015

1 Introduction to Bloch-Kato Conjecture (continued)

Recall some notation from last time. Let X be a smooth proper $\mathbf{Q}$-scheme and $L^i(X, s)$ the L-function attached to $H^i_{\text{ét}}(X)$. Set $p = i + 1$ and $q^* = p - q$ where $q \geq p/2$ (swap q and q^* if necessary). There is a (conjectural) functional equation relating

$$L^i(X, s) \leftrightarrow L^i(X, p - s)$$

so information about $L^i(X, q^*)$ is equivalent to information about $L^i(X, q)$. However, the conjectures are nicer to state for $L^i(X, q^*)$, as in the classical story of the zeta function.

1.1 Deligne's conjecture

Deligne's conjecture (made in the late '70s, in the Corvallis volume) predicts that

$$\boxed{L^i(X, q^*) \in \mathbf{Q}(2\pi i)^{(1-q)\dim H^i} \det(\langle \omega_i, \gamma_j \rangle)}$$

where $\{\omega_i\}$ is a $\mathbf{Q}$-basis for $F^q H^2_{\text{dR}}(H, \mathbf{Q})$ and $\{\gamma_j\}$ is a $\mathbf{Q}$-basis for $H_{\text{sing}}(X(\mathbf{C}), \mathbf{Q}))^{\pm}$ with $\pm$ denoting the $(-1)^{q-1}$-eigenspace for complex conjugation.

(The fact that the two vector spaces have the same dimension is not obvious, and is a consequence of the criticality. You can think of the $\pm$ space as roughly picking out "half" of the cohomology.)

Deligne's conjecture was based on computations of Shimura, in which the conjectures are clear almost by definition.

Remark 1.1.1. There is a "dual" formulation for $L^i(X, q)$ by taking the "other halves" of the cohomology. For instance, the determinant that appears is that of the map $H^i_{\text{sing}}(X, \mathbf{Q})^{\mp} \to H^i_{\text{dR}}/F^q H^i_{\text{dR}}$. This formulation is better for certain purposes.

*Notes taken by Tony Feng

1.2 K-theory and class groups

There is another thread of thought, which seems to have started with Tate's discovering of an interesting generalization of the class number formula.

Suppose F is a totally real field. Then it was known by work of Siegel that

$$\zeta_F(-1) \in \mathbf{Q}$$

Since we have a rational number, you can ask if it measures the size of something.

Conjecture 1.2.1 (Birch-Tate, ≈1970)**.** *If F is totally real then*

$$\zeta_F(-1) = \frac{\# K_2(\mathcal{O}_F)}{w_2}$$

Here w_2 is an analogue of the number of roots of unity. Tate probably arrived at this because he was interested in K_2 (not because he was looking at zeta values). He was computing its size by hand for several number fields and function fields, which is how he arrived at this observation.

Introduction to K_2. For a field F, a *symbol* in F is a bilinear map $F^* \times F^* \to A$ to an abelian group A with the property that $(x, 1-x) = 0$.

This seems like a weird definition, but it is motivated by the many examples of symbols in nature, such as :

- The Hilbert symbol of a local field (via local class field theory).
- The *tame symbol* of a local field F with residue field k, which is defined by

$$(x, y) = \frac{x^{v(y)}}{y^{v(x)}}(-1)^{v(x)v(y)} \pmod{\mathfrak{m}} \in k^{\times}.$$

- The *differential symbol* of a local field F, which is defined by

$$\frac{dx}{x} \wedge \frac{dy}{y} \in \Omega^2_F.$$

With all these different symbols, it is natural to ask for the universal one.

Definition 1.2.2. The group $K_2(F)$ is the target of the universal symbol, so we have an explicit presentation

$$K_2(F) = F^{\times} \otimes_{\mathbf{Z}} F^{\times} / \langle x \otimes (1-x) \rangle.$$

Example 1.2.3. What is $K_2(\mathbf{Q})$? It's hard to get our hands on it, but we can define maps out of it. For each p, we get a tame symbol $K_2(\mathbf{Q}) \to (\mathbf{Z}/p\mathbf{Z})^{\times}$. We also have a Hilbert symbol $K_2(\mathbf{Q}) \to \{\pm 1\}$, which could be viewed as the tame symbol at ∞, which sends $a \otimes b \mapsto -1$ if and only if a, b are both negative.

Theorem 1.2.4 (Tate). *The product of the Tame symbol maps induces an isomorphism*

$$K_2(\mathbf{Q}) \simeq \prod_p (\mathbf{Z}/p\mathbf{Z})^\times \times \{\pm 1\}.$$

(Note that the factor at $p = 2$ is trivial, but for larger number fields the contribution at 2 is generally not trivial.)

Now let's discuss K_2 of a ring of integers.

Definition 1.2.5. For F a number field, we define

$$K_2(\mathcal{O}_F) = \ker\left(K_2(F) \to \prod_{\mathfrak{p}} (\mathcal{O}_F/\mathfrak{p})^\times\right).$$

Example 1.2.6. Therefore $K_2(\mathbf{Z}) = \mathbf{Z}/2\mathbf{Z}$.

The definition looks terribly complicated - how do you compute this in practice? Suppose we want to compute $K_2(\mathbf{Q})$. Start with a small group of units, say $U = \langle -1, 2, 3, 5, 7\rangle \subset \mathbf{Q}^\times$. This has rank 16. We could then try to understand the image of $U \otimes U$ in $K_2(\mathbf{Q})$.

How do you go about this? The image is a quotient of $U \otimes U$ by $x \otimes (1-x)$ when $x, 1-x$ both belong to U, i.e. when x and $1-x$ are both S-units for $S = \{2, 3, 5, 7\}$. The free part of U has rank 16. If you want to know how big this quotient is then you want to know how many such x there are. They are in correspondence with solutions to the equation

$$a + b = c \quad a, b, c \text{ are integers divisible only by } 2, 3, 5, 7$$

because we can then set $x = \frac{a}{c}, 1 - x = \frac{b}{c}$. It is an interesting exercise to convince yourself that there are many, many solutions to this equation. The point is that the number of relations is vastly greater than the rank, so you should view K_2 as being something giant modulo something even more giant.

In practice it turns out to be the case that if you slowly increase U, the image stabilizes very quickly. Of course, actually proving that this stabilizes is very difficult, but in practice this works well.

Remark 1.2.7. You might find it weird that after dividing a rank 16 group by some enormous set of relations, there's still something left (we know that $K_2(\mathbf{Z}) \simeq \mathbf{Z}/2\mathbf{Z}$). That suggests that there are "relations among the relations". Indeed, there is a sequence which describes these relations among relations, and the kernel is K_3 (the cokernel is K_2). That gives an effective method for computing K_3; higher K-groups are not really computable.

Somewhat later it was noticed that w_2 was, up to powers of 2, the size of K_3. (Of course, this was unavailable to Tate because he didn't have a definition of K_3.)

Therefore, a reformulation in these terms of the Birch-Tate conjecture (up to powers of 2) is

$$\zeta_F(-1) \sim_{\mathbf{Q}^\times} \frac{\#K_2(\mathcal{O}_F)}{\#K_3(\mathcal{O}_F)}.$$

This is a tantalizing generalization of the class number formula, which can be interpreted as

$$\zeta_F(0) \sim_{\mathbf{Q}^\times} \frac{\#K_0(F)}{\#K_1(F)}.$$

Why the restriction to totally real field? For imaginary quadratic, K_3 is infinite and the zeta function vanishes. One wants to look at the first non-zero term, and Tate speculated that there should be a regulator, but he didn't know what.

1.3 Towards Bloch-Kato

Borel computed $\operatorname{rank} K_i(\mathcal{O}_F)$ and found that $\operatorname{rank} K_3(\mathcal{O}_F)$ is the order of vanishing of ζ_F at -1. The rank of $K_5(\mathcal{O}_F)$ is the order of vanishing of ζ_F at -2. He also constructed a "regulator map" for $K_1(\mathcal{O}_F), K_3(\mathcal{O}_F), \ldots$, meaning maps to some real vector space of the right dimension. (Recall that $K_1(\mathcal{O}_F) = \mathcal{O}_F^\times$, whose rank $r_1 + r_2 - 1$ is the order of vanishing of ζ_F at $s = 0$, as we noted early in the first lecture.) He showed that the first nonvanishing derivative of ζ_F at $-i$ is equal, up to $\mathbf{Q}^\times$, to the regulator of K_{2i}. (The regulator is defined by taking a rational basis and taking the regulator.)

Remark 1.3.1. The rational number is something like the order of K_{2i}.

We haven't defined the higher K-groups. Whatever they are, they have a map

$$K_i\mathcal{O}_F \to H_i(\mathrm{GL}_n\,\mathcal{O}_F).$$

Borel's argument is by writing down cycles in this homology group and computing their volume.

Lichtenbaum then made a general conjecture combining the Birch-Tate conjecture and Borel.

That story was for number fields. Number fields are very special; for instance, there's basically no other class of fields for which we know finite generation of the K-theory.

Bloch conjectured a relation between $L(E, 2)$ (for E an elliptic curve over a number field) and $K_2(E)$. (We can define K_i of any scheme.) He conjectured that it is a rational multiple of the regulator of $K_2(E)$.

Shortly thereafter, Beilinson made a general conjecture (up to $\mathbf{Q}^\times$). Beilinson constructs a regulator on $K^*(X)$, which is a map - let's call it r - from a "piece" (i.e. direct summand) of $K_{2q-p}(X) \otimes \mathbf{Q}$ to the cokernel of the map

$$\int : (H^i)_{\mathrm{sing}}(X, \mathbf{Q})^\mp \to H^i_{\mathrm{dR}}/F^q H^i_{\mathrm{dR}}$$

That means that we can try to "put r and $\int$ together into a square matrix" and thus define "$\det(\int \oplus r)$."

This is for X smooth and proper over $\mathbf{Q}$; to generalize it to number fields you restrict scalars down to $\mathbf{Q}$.

Conjecture 1.3.2 (Beilinson)**.** *We have*

$$\det(\int \oplus r) \sim_{\mathbf{Q}^\times} L^i(X, q).$$

1.4 Bloch-Kato

Finally, let's say something about Bloch-Kato. (See also Fontaine's exposition.) Bloch and Kato made a conjecture without the $\mathbf{Q}^\times$-ambiguity. (Strictly speaking one would say that they made a conjecture up to $\mathbf{Z}^\times$, but it's always easy to figure out the sign.)

To highlight one challenge in generalizing BSD, recall that BSD predicts

$$L(E, 1) = \frac{Ш_E \Omega_E R_E}{E(\mathbf{Q})^2_{\text{tors}}}.$$

The $\Omega_E = \int_{E(\mathbf{R})} |\omega|$ is a bit weird: you have to use a Néron model to normalize the form ω. The Néron model is a miracle of abelian varieties, and definitely has no analogue in general.

The key input in Bloch-Kato is that they generalize the logarithm map. If you have $E/\mathbf{Q}_p$ then there is a logarithm map

$$\log\colon E(\mathbf{Q}_p) \to T_0(E) = H^0(E, \Omega^1)^\vee.$$

Moreover, this is an isomorphism after tensoring with $\mathbf{Q}_p$. The Néron form ω lives in $H^0(E, \Omega^1)^\vee$, and thusu defines an integral structure on $H^0(E, \Omega^1)^\vee \otimes \mathbf{Q}_p$. There is also an integral structure on $E(\mathbf{Q}_p) \otimes \mathbf{Q}_p$, namely the image of $E(\mathbf{Q}_p)$. These two integral structures *don't match up*, but they are related by something concrete.

The generalization of log is the "Bloch-Kato exponential"

$$H^i_{\text{dR}}(X)/F^q H^i_{\text{dR}} \otimes \mathbf{Q}_\ell \to H^q_{\text{ét}}(X \times \overline{\mathbf{Q}_\ell}, \mathbf{Q}_\ell(q))$$

the latter being regarded as a "piece" of $H^p_{\text{ét}}(X, \mathbf{Z}(q))$. The right side has an ℓ-adically integral structure, so it gives a way of attaching such an integral structure to the left hand side. (This is a little loose. Really the point is that the choices made in defining integral structures cancels out.)

Example 1.4.1. Suppose you have a modular form f, with coefficients in some totally real field K. This predicts $L(f, q)$ up to $\mathcal{O}_K^\times$. A similar thing came up before for $\mathbf{Z}$, but in that case it was no problem because the units were small; for general K the units are a lot bigger. Nobody really knows how to fix this.

2 Recap of issues in the formulation of BSD

Let A be an abelian variety over a global field K. We are now going to discuss some of the parts of the BSD formula that didn't quite make sense last time. In particular, we defined the L-function to be

$$L(A/K, s) = \prod_v \det(1 - \mathrm{Frob}_v\, q_v^{-s} \mid V_\ell(A)^{I_v})^{-1}.$$

There were several problems with this. One was that q_v^{-s} is a complex number, and inserting it into such a determinant doesn't literally make sense as written. Also, this definition is for $\ell \neq \mathrm{char}\, \mathbf{F}_v$, so we cannot use the same ℓ for *every* finite place v (for K a number field). Therefore, we must prove independence of ℓ. More precisely:

1. Is the characteristic polynomial of Frob_v on $V_\ell(A)^{I_v}$, which is a priori in $\mathbf{Q}_\ell[T]$, actually in $\mathbf{Q}[T]$, and as such independent of ℓ?

   The answer turns out to be *yes*, which we will see by using the geometry of the v-reduction of the Néron model. (In the number field case the Néron model is a scheme over the ring of integers; in the function field case it is a scheme over the proper curve defining the function field.)

   (A bonus that comes out of this is that we'll get the right estimates to know that the product converges, by combining it with the Riemann hypothesis for abelian varieties over $\mathbf{F}_v$ when v is good. The estimates on the magnitude of the Frobenius eigenvalues ensures that we get convergence for $\mathrm{Re}(s) > \frac{3}{2}$.)

2. The BSD conjecture predicts that we can analytically continue to $\mathbf{C}$, and that

   $$L(s) \sim C_A (s-1)^{\mathrm{rank}\, A(K)} \text{ as } s \to 1$$

   where

   $$C_A = \frac{(\#\text{Ш}_A) R_A \Omega_A}{\#A(K)_{\mathrm{tor}} \widehat{A}(K)_{\mathrm{tor}}}.$$

   Here

   - the factor R_A is defined via height pairings (discussed below), which *don't* rely on the theory of Néron models.
   - Ω_A and Ш_A involve the Néron model.

   So the next order of business is to discuss the Néron model, height pairings, and Tate-Shafarevich group. Today we'll only have time to discuss height pairings.

Just as an aside, note that for an elliptic curve the quantity $R_E/\#E(K)_{\mathrm{tor}}^2$ can be interpreted as the leading coefficient of an asymptotic:

$$\#\{p \in E(K) \colon \text{naive height}(p) \leq x\} \sim \frac{R_E}{\#E(K)_{\mathrm{tor}}^2} (\log x)^{\mathrm{rank}\, E(K)}.$$

You might wonder if there is such an interpretation for abelian varieties. The answer is yes, in some sense, but it's not as natural because we need to choose a polarization.

3 Height Pairings

We shall begin with a review of canonical height functions on $A(\overline{K})$ associated to *any* line bundle $\mathcal{L}$ on A, not just ample symmetric ones as usually done in a first course on abelian varieties: we will need to work with line bundles that come from the dual abelian variety, and are those are *never* ample and are always *anti-symmetric* in the sense that $[-1]^*(\mathcal{L}) \simeq \mathcal{L}^{-1}$ (for reasons stemming from the Theorem of the Square, to be recalled below). **Warning**: in terms of the correspondence between line bundles and Weil divisors, if $\mathcal{L} \leftrightarrow D$ then $[-1]^*\mathcal{L}$ has nothing to do with $-D$ (think about even just elliptic curves).

[Reference for everything that follows: §9 of B. Conrad's "Chow trace" article and references therein, as well as the discussion of heights near the end of his course on abelian varieties (at `http://web.stanford.edu/~tonyfeng/249C.pdf`) for some of the basics.]

3.1 Height functions

Let $\mathcal{L}$ be a line bundle on A. Recall that we say that $\mathcal{L}$ is *symmetric* if $\mathcal{L} \simeq [-1]^*\mathcal{L}$; e.g., a source of such examples is $\mathcal{N} \otimes [-1]^*\mathcal{N}$ for line bundles $\mathcal{N}$. (You can imagine the analogous definition for asymmetric bundles.) If also $\mathcal{N}$ is ample then since the pullback under an automorphism is also ample it follows that get many symmetric ample line bundles.

Choose $\mathcal{L}$ to be very ample, so we get

$$A \hookrightarrow \mathbf{P}(\Gamma(A, \mathcal{L})) \overset{\theta}{\simeq} \mathbf{P}^N_K$$

where θ is induced by choosing a basis for $\Gamma(A, \mathcal{L})$. Then we get a "naïve height" $h_{K,\mathcal{L}} : A(\overline{K}) \hookrightarrow \mathbf{P}^N(\overline{K}) \xrightarrow{h_{N,K}} \mathbf{R}$ where $h_{N,K}$ is the standard height function relative to K. The function $h_{K,\mathcal{L}}$ depends on $\mathcal{L}$ *and on* θ (but in only a minor way on the latter).

Some Facts:

- The function $h_{K,\mathcal{L}}$ is independent of θ modulo $O(1)$ (i.e. modulo bounded functions on $A(\overline{K})$. This is elementary.

- (*Additivity in* $\mathcal{L}$) We have

  $$h_{K,\mathcal{L}_1 \otimes \mathcal{L}_2} = h_{K,\mathcal{L}_1} + h_{K,\mathcal{L}_2} \mod O(1).$$

  So in particular, $h_{K,\mathcal{L}^{\otimes n}} \sim n h_{K,\mathcal{L}}$, allowing us to extend to ample $\mathcal{L}$.

  Using this additivity we can define $h_{K,\mathcal{L}}$ modulo $O(1)$ for *any* $\mathcal{L}$ since $\mathcal{L} \simeq \mathcal{L}_1 \otimes \mathcal{L}_2^{-1}$ for very ample $\mathcal{L}_1, \mathcal{L}_2$. (With a bit of work one can show that if $\mathcal{L}$ is generated by global sections then $h_{K,\mathcal{L}}$ agrees modulo $O(1)$ with the function obtained by composing $A \to \mathbf{P}(\Gamma(A, \mathcal{L})) \simeq \mathbf{P}^N_K$ on $\overline{K}$-points with $h_{N,K}$.)

- If also $\mathcal{L}$ is symmetric then $h_{K,\mathcal{L}}$ is a "quadratic form mod $O(1)$". Tate showed that by a clever limit trick that in the $O(1)$-equivalence class there is a unique genuine *quadratic form* $\widehat{h}_{K,\mathcal{L}}$, given by

$$a \mapsto \lim_{n\to\infty} \frac{h_{K,\mathcal{L}}(na)}{n^2}.$$

- If $\mathcal{L}$ is *anti-symmetric* then one has a similar story replacing "quadratic form" with "additive function". That is, $h_{K,\mathcal{L}}$ is "additive modulo $O(1)$" and its $O(1)$-class contains a unique additive function, defined by

$$\widehat{h}_{K,\mathcal{L}} \colon a \mapsto \lim_{n\to\infty} \frac{h_{K,\mathcal{L}}(na)}{n}.$$

Recall that if A is an abelian variety then its *dual abelian variety* is

$$\widehat{A} = \mathrm{Pic}^0_{A/K} \subset \mathrm{Pic}_{A/K}.$$

By the Theorem of the Square, one characterization of line bundles $\mathcal{L}$ coming from Pic^0 is that

$$m^*\mathcal{L} \simeq p_1^*\mathcal{L} \otimes p_2^*\mathcal{L}$$

on $A \times A$. Now consider pulling back along the antidiagonal map $(a,a) \mapsto (a,-a)$. That kills the left side, so we get

$$\mathcal{O}_A \simeq \mathcal{L} \otimes [-1]^*\mathcal{L}.$$

This shows that line bundles coming from $\widehat{A}$ are always anti-symmetric.

Just as we can express a homomorphism from an abelian group into a $\mathbf{Z}[1/2]$-module uniquely as a sum of an even function and an odd function, for any line bundle $\mathcal{L}$ on A we define the line bundles

$$\begin{aligned} \mathcal{L}^+ &= \mathcal{L} \otimes [-1]^*\mathcal{L}, \\ \mathcal{L}^- &= \mathcal{L} \otimes [-1]^*\mathcal{L}^{-1} \end{aligned}$$

that are respectively symmetric and anti-symmetric. Then the function

$$\boxed{\widehat{h}_{K,\mathcal{L}} := \frac{1}{2}\left(\widehat{h}_{K,\mathcal{L}^+} + \widehat{h}_{K,\mathcal{L}^-}\right)}$$

is in the $O(1)$-class of $h_{K,\mathcal{L}}$. This is the unique function in the $O(1)$-class of $h_{K,\mathcal{L}}$ that is "polynomial of degree ≤ 2": the first piece $\widehat{h}^+_{K,\mathcal{L}}$ "extracts" the quadratic part and the second piece $\widehat{h}_{K,\mathcal{L}^-}$ "extracts" the additive part. This all depends additively on $\mathcal{L}$.

Remark 3.1.1. If K'/K is finite then $\widehat{h}_{K',\mathcal{L}_{K'}} = [K' : K]\widehat{h}_{K,\mathcal{L}}$. This is an immediate consequence of the definition of the naive height on projective spaces.

3.2 The canonical height pairing

Definition 3.2.1. The *canonical height pairing* is the bi-additive function

$$A(\overline{K}) \times \widehat{A}(\overline{K}) \to \mathbf{R}$$

defined by

$$(a, \mathcal{L}) \mapsto \frac{1}{[K' : K]} \widehat{h}_{K', \mathcal{L}}(a)$$

where K'/K is a finite extension for which $\mathcal{L} \in \widehat{A}(K')$ and $a \in A(K')$ (the choice of such K'/K clearly does not matter, due to the division by field degree). Note that $\mathcal{L}$ is antisymmetric because $\mathcal{L}$ comes from $\widehat{A}$, so this construction really is additive in a for fixed $\mathcal{L}$.

Remark 3.2.2. If $\mathcal{P}$ denotes the Poincaré bundle on $A \times \widehat{A}$ then it can be shown that $\widehat{h}_{K,\mathcal{P}}(a, \mathcal{L}) = \langle a, \mathcal{L} \rangle$.

Theorem 3.2.3. *For $\mathcal{N}$ invertible, the "Mumford construction" $\phi_{\mathcal{N}} \colon A \to \widehat{A}$ defined by $x \mapsto t_x^*(\mathcal{N}) \otimes \mathcal{N}^{-1}$ induces a bi-additive map*

$$\langle \cdot, \cdot \rangle_{A,} \circ (\mathrm{Id} \times \phi_{\mathcal{N}}) \colon A(\overline{K}) \times A(\overline{K}) \to \mathbf{R}$$

and this is equal to

$$(a_1, a_2) \mapsto \widehat{h}_{K,\mathcal{N}}(a_1 + a_2) - \widehat{h}_{K,\mathcal{N}}(a_1) - \widehat{h}_{K,\mathcal{N}}(a_2).$$

Observe that if $\mathcal{N}$ is symmetric then up to a factor of 2 this is the bilinear form associated to the quadratic form $\widehat{h}_{K,\mathcal{N}}$. Recall also that if $\mathcal{N}$ is ample then $\phi_{\mathcal{N}}$ is an isogeny, so it is a finite-to-one surjective on $\overline{K}$-points.

We conclude that for $\mathcal{N}$ ample and symmetric the commutative diagram

$$\begin{array}{ccc} A(\overline{K})_{\mathbf{R}} \times A(\overline{K})_{\mathbf{R}} & & \\ {\scriptstyle 1 \times \phi_{\mathcal{N}}} \downarrow \simeq & \searrow & \\ A(\overline{K})_{\mathbf{R}} \times \widehat{A}(\overline{K})_{\mathbf{R}} & \longrightarrow & \mathbf{R} \end{array}$$

(which is clearly an isomorphism vertically) has diagonal that on K'-points for any finite subextension K'/K is the bilinear form associated to the quadratic form $\widehat{h}_{K,\mathcal{N}}$ that is shown to be *positive-definite* in the basic arithmetic theory of abelian varieties (via Minkowski convex-body arguments).

Corollary 3.2.4. *The pairing $A(\overline{K})_{\mathbf{R}} \times \widehat{A}(\overline{K})_{\mathbf{R}} \to \mathbf{R}$ is non-degenerate on K'-points for all finite K'/K. In particular, the kernel on each side of the pairing is trivial.*

We have *Mordell-Weil lattices* $A(K)/\text{torsion} \subset \mathrm{A(K)}_{\mathbf{R}}$ and $\widehat{A}(K)/\text{torsion} \subset \widehat{\mathrm{A}}(\mathrm{K})_{\mathbf{R}}$ in these finite-dimensional $\mathbf{R}$-vector spaces of the same dimension, and the preceding shows that these $\mathbf{R}$-vector spaces are in perfect duality via the canonical height pairing. This establishes the *non-vanishing* in:

Definition 3.2.5. We define the *regulator* of A to be

$$R_A = |\det(\langle a_i, \mathcal{L}_j \rangle)| \neq 0$$

for ordered bases $\{a_i\}, \{\mathcal{L}_j\}$ of the Mordell-Weil lattices. Note that this is independent of the choice of bases.

Remark 3.2.6. What happens under base change of the ground field? The L-function can change dramatically, and so does the regulator.

On the other hand, nothing much changes under Weil restriction. Recall that BSD predicts

$$L(A,s) \sim \frac{Ш_A R_A \Omega_A}{\#A(K)_{\text{tors}} \# \widehat{A}(K)_{\text{tors}}} (s-1)^{\operatorname{rank} A} \ldots$$

If K/K_0 is a finite separable extension with K_0 a global field then the Weil restriction $A_0 := R_{K/K_0}(A)$ is an abelian variety over K_0 of dimension $[K : K_0] \dim A$ with a natural identification $A_0(K_0) = A(K)$ (hence the ranks coincide) and $V_\ell(A_0) = \operatorname{Ind}_K^{K_0}(V_\ell(A))$ as Galois modules. The invariance of the Artin formalism under induction then gives that $L(A_0/K_0, s) = L(A/K, s)$, so a first test of BSD is whether the leading coefficients for A_0 and A agree. In fact they agree term by term (so this is a much weaker "test" than isogeny-invariance, for which we will see that individual factors are not isogeny-invariant)! The invariance under Weil restriction involves several checks:

(i) The cohomological definition of $Ш_A$ to be discussed next time will yield (via Shapiro's Lemma) that $Ш_A \simeq Ш_{A_0}$. Alternatively, we will have a way to express $Ш_A$ in terms of the Néron model, from which this isomorphism can also be explained via exactness of finite-pushforward for the étale topology. (This alternative proof is truly killing a fly with a sledgehammer.)

(ii) In general for a field k and finite étale k-algebra k' (such as a finite separable extension field, or product of copies of such) and a projective k'-scheme X' with geometrically integral fibers, there is a canonical isomorphism

$$\mathrm{R}_{k'/k}(\mathrm{Pic}^0_{X'/k'}) \simeq \mathrm{Pic}^0_{\mathrm{R}_{k'/k}(X')/k}.$$

[The precise definition of this map requires some genuine thought and is left to the reader as an exercise. Once the map is defined, the proof that it is an isomorphism reduce to the fact that the formation of Picard schemes of geometrically integral schemes commutes with direct products. But this compatibility is not at all trivial since it is *false* for the entire Picard scheme and in positive characteristic requires the full force of the scheme-theoretic version of the Theorem of the Cube.]

Applying this compatibility to abelian varieties, the torsion factors match through Weil restriction even for the dual term.

(iii) The regulators match via the compatibility of duality with Weil restriction *provided* that this duality is compatible with the formation of canonical heights. This latter compatibility is not immediately obvious just from the definitions! It is however most efficiently proved by exploiting *extension of the ground field* to split K/K_0 (so it was good that we set up the theory of heights on geometric points!) in order to reduce to the analogous much easier compatibility of the formation of canonical heights with respect to direct products of abelian varieties.

The only nontrivial compatibility for Weil restriction is invariance of the volume term Ω_A (to be defined next time). Since the volume term is defined using Haar measures built via top-degree differential forms, which is to say global sections of the determinant of the cotangent bundle, the difficulty is related to the nontrivial interaction between finite flat pushforward and the formation of determinant bundles. After some unraveling, this amounts to the following (surprisingly challenging!) algebraic geometry problem: for a finite locally free map of schemes $f : X \to Y$ and a vector bundle $\mathcal{E}$ on X, the line bundle $\det_X(\mathcal{E})$ has an associated line bundle $\mathrm{N}_{X/Y}(\det_X(\mathcal{E}))$ on Y and there is also the line bundle $\det_Y(f_*(\mathcal{E}))$ on Y. How are these related? Even in the affine case, where it is a question in commutative algebra, the task is not so easy.

These two line bundles on Y are *not* generally isomorphic, and the canonical discrepancy is governed by $\det_Y(f_*(O_X))$ via a formula inspired by the transitivity of discriminants in number theory. However, the proof is necessarily totally different from the situation in number theory, making no use of features of Dedekind domains, because there is no "universal finite flat ring extension" and hence this module-theoretic problem *cannot* be reduced to a "universal case" over some polynomial ring over $\mathbf{Z}$ (to check at height-1 primes). We refer the reader to §4–§5 in Chapter II of Oesterlé's marvelous 1984 Inventiones article on Tamagawa numbers for an elegant discussion of the general task of relating Haar measures to Weil restriction over local rings, which includes a solution to the preceding problem with determinants and finite flat pushforward.