ARTICLE IN PRESS

Available online at www.sciencedirect.com

SCIENCE DIRECT®

Advances in Mathematics ■■■ (■■■■) ■■■–■■■

ADVANCES IN Mathematics

www.elsevier.com/locate/aim

Organizing the arithmetic of elliptic curves

Barry Mazur[a,1], Karl Rubin[b,*,1]

[a]*Department of Mathematics, Harvard University, Cambridge, MA 02138, USA*
[b]*Department of Mathematics, University of California Irvine, Irvine, CA 92697, USA*

Received 28 November 2004; accepted 11 May 2005

Communicated by Johan De Jong

Abstract

Suppose that E is an elliptic curve defined over a number field K, p is a rational prime, and K_∞ is the maximal $\mathbf{Z}_p$-power extension of K. In previous work [B. Mazur, K. Rubin, Elliptic curves and class field theory, in: Ta Tsien Li (Ed.), Proceedings of the International Congress of Mathematicians, ICM 2002, vol. II, Higher Education Press, Beijing, 2002, pp. 185–195; B. Mazur, K. Rubin, Pairings in the arithmetic of elliptic curves, in: J. Cremona et al. (Eds.), Modular Curves and Abelian Varieties, Progress in Mathematics, vol. 224, 2004, pp. 151–163] we discussed the possibility that much of the arithmetic of E over K_∞ (i.e., the Mordell–Weil groups and their p-adic height pairings, the Shafarevich–Tate groups and their Cassels pairings, over all finite extensions of K in K_∞) can be described efficiently in terms of a single skew-Hermitian matrix with entries drawn from the Iwasawa algebra of K_∞/K.

In this paper, using work of Nekovář [J. Nekovář, Selmer complexes. Preprint available at ⟨http://www.math.jussieu.fr/∼nekovar/pu/⟩], we show that under not-too-stringent conditions such an "organizing" matrix does in fact exist. We also work out an assortment of numerical instances in which we can describe the organizing matrix explicitly.

Keywords: Elliptic curves; Iwasawa theory; Selmer groups; Mordell–Weil groups; Cohomological pairings

[*] Corresponding author. Fax: +1 949 824 7993.
E-mail addresses: mazur@math.harvard.edu (B. Mazur), krubin@math.uci.edu (K. Rubin).

[1] The authors are supported by NSF grants DMS-0403374 and DMS-0140378, respectively.

0001-8708/$ - see front matter © 2005 Elsevier Inc. All rights reserved.
doi:10.1016/j.aim.2005.05.024

1. Introduction

Fix the data (p, K, E) where p is a prime number, K a number field, and E an elliptic curve over $\mathbf{Q}$. Let K_∞/K denote the maximal $\mathbf{Z}_p$-power extension of K. Recent work[2] provides, in some instances, detailed information about p-adic completions of Mordell–Weil groups and their associated p-adic height pairings, and the p-primary Shafarevich–Tate groups and their associated Cassels pairings, over intermediate fields in K_∞/K. Added to this information we also have a constellation of conjectures telling us even more precisely how all this arithmetic should behave.

In previous articles [MR1,MR2] we have considered the possibility that, under some not too stringent assumptions, much of this arithmetic data can be packaged efficiently in terms of a single skew-Hermitian matrix with entries drawn from the Iwasawa algebra of the $\mathbf{Z}_p$-power extension K_∞/K. We say that such a matrix H *organizes the arithmetic of* (p, K, E) if it plays this role vis-à-vis the arithmetic of (p, K, E). For a detailed discussion of this, see §7. In the special case where there is no nontrivial p-torsion in the Shafarevich–Tate group of E over K, our skew-Hermitian matrix may be thought of as a (skew-Hermitian) lifting to the Iwasawa algebra of the matrix describing the p-adic height pairing on the Mordell–Weil group $E(K)$.

The main result. Theorems 7.5 and 7.7 provide a construction of such skew-Hermitian “organizing matrices” in a fairly general context. Our construction depends heavily on work of Nekovář [N] (which in turn makes use of work of Greenberg). An example of what we can prove is the following.

Let (p, K, E) be such that

- $K/\mathbf{Q}$ is abelian,
- the integers p, $\mathrm{disc}(K)$, $\mathrm{cond}(E)$ are pairwise relative prime,
- E has ordinary reduction at p,
- p does not divide $\#E(k_v)$ for any of the residue fields k_v at places v of K lying above p,
- the Tamagawa numbers of E/K are all prime to p.

Then an organizing matrix H for the arithmetic of (p, K, E) exists, and is unique up to (noncanonical) equivalence.

We work out an assortment of numerical instances in which we can describe the organizing matrix explicitly. In §9 we consider the case where the base field K is $\mathbf{Q}$. For example, if E is either of the curves denoted 1058C1 or 1058D1 in [Cr] (and assuming the Birch and Swinnerton–Dyer conjecture for $E/\mathbf{Q}$) then using calculations by William Stein we can give the organizing matrix H exactly for all 337 primes less than 2400 that satisfy the conditions listed above. We also show that a congruence modulo 5 between the modular forms corresponding to these two curves is matched by a congruence modulo 5 between their organizing matrices.

[2] Advances here have been made be many people, including Bertolini and Darmon [BD1,BD2], Cornut [Co], Greenberg [G1,G2], Howard [Ho2,Ho1], Kato [Ka], Nekovář [N], Perrin-Riou [PR1,PR2,PR3,PR4], and Vatsal [V].

In §10 we consider the case where E is defined over $\mathbf{Q}$ and K is an imaginary quadratic field satisfying the "Heegner condition". We find, among other things, examples of Iwasawa modules X^{anti} attached to elliptic curves over anti-cyclotomic $\mathbf{Z}_p$-extensions such that X^{anti} contains nontrivial finite submodules, and we also give a *counterexample* to a prior conjecture of ours.

To describe the structure we deal with in more detail, put $\Lambda := \mathbf{Z}_p[[\mathrm{Gal}(K_\infty/K)]]$, and denote by $\iota : \Lambda \to \Lambda$ the standard involution (that sends every group element γ in Λ to its inverse and is the identity on $\mathbf{Z}_p$). If M is a Λ-module, its *conjugate* M^ι is the Λ-module with the same underlying group as M but with Λ-module structure obtained from that of M by composition with ι. By a *basic skew-Hermitian Λ-module* Φ we mean a free Λ-module of finite rank equipped with a skew-Hermitian pairing,

$$\Phi \otimes_\Lambda \Phi^\iota \to \mathfrak{m} \subset \Lambda,$$

where $\mathfrak{m}$ is the maximal ideal in Λ, and such that this pairing is nondegenerate after extending scalars to the field of fractions of Λ. If the arithmetic of (p, K, E) is organized by Φ, we can derive Mordell–Weil and Shafarevich–Tate information at all layers of K_∞/K together with their self-pairings from the structure of the basic skew-Hermitian Λ-module Φ, as described in §7.

Given an organizing module Φ for (p, K, E) as above, consider the free Λ-module of rank one $\Delta := \det_\Lambda \Phi^{-1}$, i.e., the inverse of the determinant module of Φ over Λ. Define $L_p^{\mathrm{arith}}(K, E)$, the *arithmetic p-adic L function attached to* (p, K, E) (relative to the organizing module Φ) to be the discriminant of the skew-Hermitian module Φ. (The definition of a p-adic L-function as a determinant of a complex in a derived category has already appeared in the work of Nekovář; see the footnote at the end of the introduction to [N].) Given our hypotheses above, the arithmetic p-adic L-function is a nonzero element

$$L_p^{\mathrm{arith}}(K, E) \in \Delta \otimes_\Lambda \Delta^\iota.$$

How canonical is this construction? First, the Λ-module $\Delta \otimes_\Lambda \Delta^\iota$ is canonically isomorphic to the determinant Λ-module of Nekovář's "Selmer complex," which is represented in the derived category by a finite complex of projective modules of finite rank (under the hypotheses listed above). Therefore, the free Λ-module of rank one $\Delta \otimes_\Lambda \Delta^\iota$ is canonically determined by our initial data (p, K, E), as is the element $L_p^{\mathrm{arith}}(K, E)$ in it.

There is also a canonical *orientation* on $\Delta \otimes_\Lambda \Delta^\iota$. By an orientation of a free Λ-module of rank one let us a mean a choice of generator up to multiplication by an element of the form $u \cdot u^\iota$ where $u \in \Lambda^\times$ is a unit. Since the organizing module Φ is determined up to (noncanonical) equivalence, we have that $\Delta \otimes_\Lambda \Delta^\iota$ inherits a canonical orientation.

There is, of course, the p-adic analytic side of this story. For simplicity fix $K = \mathbf{Q}$. We have the standard (modular symbols) construction of the p-adic analytic L-function of the elliptic curve, $L_p^{\mathrm{anal}}(K, E)$, which can be viewed, again canonically, as an element of $H_1(E(\mathbf{C}), \mathbf{Z})^+ \otimes_{\mathbf{Z}} \Lambda$, where the superscript $+$ refers to the $+$-eigenspace of the homology group in question under the action of complex conjugation. Given the modular parametrization $X_0(\mathrm{cond}(E)) \to E$ we may even make a canonical choice of a “positive” generator of the infinite cyclic group $H_1(E(\mathbf{C}), \mathbf{Z})^+$. Identifying $H_1(E(\mathbf{C}), \mathbf{Z})^+$ with $\mathbf{Z}$ via the canonical generator, we may view $L_p^{\mathrm{anal}}(K, E)$ as an element of Λ, this being one of the accidental bonuses (as we shall see below) of working with elliptic curves rather than abelian varieties of higher dimension, or modular eigenforms of higher weight. The *expectation* here (the *main conjecture*, in this context) for which there is now much evidence, is that (giving $L_p^{\mathrm{anal}}(K, E)$ a natural normalization) there is a unique generator g of the free Λ-module of rank one $\Delta \otimes_\Lambda \Delta^\iota$ such that

$$L_p^{\mathrm{anal}}(K, E) \cdot g \;=\; L_p^{\mathrm{arith}}(K, E).$$

It is natural to wonder whether this unique generator g might bear some clear relationship to the orientation structure of $\Delta \otimes_\Lambda \Delta^\iota$; it might make sense to make use of the theory of Shimura’s lift to half-integral weight modular forms to study this question.

Questions about variation. We feel that our result might be but the first hint of some kind of *generic purity* phenomenon regarding Nekovář’s Selmer complexes. The remainder of this introduction section is completely speculative, and is offered to give a sense of what we might mean by this.

Let $p \geqslant 5$ be a prime number. Put $\mathbf{W} = \mathbf{Z}_p[[\mathbf{Z}_p^\times]]$, which we take as p-adic *weight space,* where for $k \in \mathbf{Z}$, we have $s_k : \mathbf{W} \to \mathbf{Z}_p$, the natural *projection to weight k and nebentypus character* ω^k. Here ω is the standard Teichmüller character, and s_k is the $\mathbf{Z}_p$-algebra homomorphism that sends a group element $x \in \mathbf{Z}_p^\times$ to $x^k \in \mathbf{Z}_p^\times \subset \mathbf{Z}_p$.

Let $\mathbf{T}$ denote Hida’s Hecke algebra for ordinary p-adic modular eigenforms on $\Gamma_0(p)$. Hida’s Hecke algebra $\mathbf{T}$ is a finite flat $\mathbf{W}$-algebra with the following property. For $k = 2, 3, 4, \ldots$ if we make the base change from $\mathbf{W}$ to $\mathbf{Z}_p$ via s_k we have that $\mathbf{T} \otimes_{\mathbf{W}} \mathbf{Z}_p$ is naturally isomorphic to the (classical) Hecke algebra that acts faithfully on p-adic cuspidal ordinary modular eigenforms on $\Gamma_1(p)$ of weight k and nebentypus character ω^k. Let $m \subset \mathbf{T}$ denote a maximal ideal associated to an absolutely irreducible residual representation of the Galois group $\bar{\rho} : \mathrm{Gal}(\bar{\mathbf{Q}}/\mathbf{Q}) \to \mathrm{GL}_2(\mathbf{T}/m)$ and let $\mathbf{T}_m$ denote the completion of $\mathbf{T}$ at m. Put

$$R := \mathbf{T}_m \hat{\otimes}_{\mathbf{Z}_p} \Lambda,$$

and let $\iota : R \to R$ denote the involution $1 \hat{\otimes} \iota$. There is a canonical representation

$$\rho : \mathrm{Gal}(\bar{\mathbf{Q}}/\mathbf{Q}) \longrightarrow \mathrm{GL}_2(R),$$

unramified outside p, uniquely characterized by the requirement that if

$$f = q + \sum_{n \geqslant 2} a_n(f) q^n$$

is an ordinary eigenform on $\Gamma_1(p)$ whose associated residual representation is equivalent to $\bar{\rho}$ and if $\chi : \mathrm{Gal}(\mathbf{Q}_\infty/\mathbf{Q}) \to \mathbf{C}_p^\times$ is a wild p-adic character, then the Galois representation

$$\mathrm{Gal}(\bar{\mathbf{Q}}/\mathbf{Q}) \longrightarrow \mathrm{GL}_2(\mathbf{C}_p)$$

attached to $f \otimes \chi$ is the one induced from ρ by the homomorphism $R \to \mathbf{C}_p$ which, for positive integers n prime to p, takes $T_n \hat{\otimes} \gamma$ to $a_n(f)\chi(\gamma)$ and takes $U_p \hat{\otimes} \gamma$ to $a_p(f)\chi(\gamma)$.

Attached to ρ there is a (finitely generated) Selmer R-module S, which we wish to view as coherent sheaf $\mathcal{S}$ over $X := \mathrm{Spec}(R)$. Moreover, there is a "two-variable" p-adic L function L_p^{anal} that is naturally a section of a certain line bundle[3] over X that we will denote P.

In view of the main result of this article, we might wonder whether there are fairly general conditions under which one may find a Zariski open subscheme $Y \subset X = \mathrm{Spec}(R)$ stable under ι, and a skew-Hermitian vector bundle Φ of finite rank over Y with these two properties:

- The skew-Hermitian vector bundle Φ over Y bears an "organizing" relationship to the coherent sheaf $\mathcal{S} \otimes_{\mathcal{O}_X} \mathcal{O}_Y$ (analogous to the relationship that the organizing skew-Hermitian module Φ in the context of elliptic curves above bears to the classical Selmer module).
- Forming $\Delta := \det \Phi^{-1}$, which is a line bundle over Y, and

  $$L_p^{\mathrm{arith}} := \mathrm{discriminant}(\Phi),$$

  viewed as a section of the line bundle $\Delta \otimes \Delta^{\iota}$ over Y, there is a (unique) isomorphism of line bundles

  $$g : P \otimes_{\mathcal{O}_X} \mathcal{O}_Y \cong \Delta \otimes \Delta^{\iota}$$

  that brings the section L_p^{anal} (restricted to Y) to L_p^{arith} (this being analogous to the "main conjecture" relationship between arithmetic and analytic p-adic L-functions of elliptic curves described above).

2. The setup

Fix a number field K, an elliptic curve E defined over K, and a rational prime p such that E has good ordinary reduction at all primes of K above p.

[3] Usually one defines L_p^{anal} to be a bona fide function (cf. [GS,Ki]) but the natural construction of this two-variable L-function—independent of any choice—is as a section of a specific line bundle that we refer to above as P, which one must trivialize to express L_p^{anal} as a function.

For every finite extension L of K we have the p-power Selmer group

$$\mathrm{Sel}_p(E, L) := \ker(H^1(L, E[p^\infty]) \longrightarrow \prod_v H^1(L_v, E)),$$

where $E[p^\infty]$ is the Galois module of p-power torsion on E, and the product is over all places v of L. This Selmer group sits in an exact sequence

$$0 \longrightarrow E(L) \otimes \mathbf{Q}_p/\mathbf{Z}_p \longrightarrow \mathrm{Sel}_p(E, L) \longrightarrow Ш(E, L)[p^\infty] \longrightarrow 0, \tag{2.1}$$

where $Ш(E, L)[p^\infty]$ is the p-primary part of the Shafarevich–Tate group of E over L.

Let K_∞ denote the maximal $\mathbf{Z}_p$-power extension of K, i.e., $\mathrm{Gal}(K_\infty/K) \cong \mathbf{Z}_p^d$ for some $d \in \mathbf{Z}^+$ and K_∞ contains all $\mathbf{Z}_p$-extensions of K. By class field theory we have $r_2 + 1 \leqslant d \leqslant [K : \mathbf{Q}]$, where r_2 is the number of complex places of K, and $d = r_2 + 1$ if Leopoldt's Conjecture holds for K. In particular $d = 1$ if $K = \mathbf{Q}$ and $d = 2$ if K is quadratic imaginary. Let $\Gamma := \mathrm{Gal}(K_\infty/K)$, and define the Iwasawa algebra

$$\Lambda := \mathbf{Z}_p[[\Gamma]].$$

If $K \subset L \subset K_\infty$ we let $\Gamma_L := \mathrm{Gal}(L/K)$ and $\Lambda_L := \mathbf{Z}_p[[\Gamma_L]]$ for the corresponding quotients of Γ and Λ.

As in the introduction, we let $\iota : \Lambda_L \to \Lambda_L$ denote the involution that sends $\gamma \mapsto \gamma^{-1}$ for $\gamma \in \Gamma_L$, and if M is a Λ_L-module we let M^ι be the *conjugate module*, the Λ_L-module with the same underlying abelian group as M, but with Λ_L-module structure obtained from that of M by composition with ι.

If $K \subset L \subset K_\infty$ we define

$$\mathrm{Sel}_p(E, L) := \varinjlim \mathrm{Sel}_p(E, F),$$

direct limit (with respect to restriction maps on Galois cohomology) over finite extensions F of K in L, and the Pontrjagin dual

$$\mathcal{S}_p(E, L) := \mathrm{Hom}(\mathrm{Sel}_p(E, L), \mathbf{Q}_p/\mathbf{Z}_p).$$

We will frequently make the following assumption.

Perfect control assumption. *If $K \subset L \subset K_\infty$ then the canonical restriction map*

$$\mathrm{Sel}_p(E, L) \longrightarrow \mathrm{Sel}_p(E, K_\infty)^{\mathrm{Gal}(K_\infty/L)}$$

is an isomorphism.

Remark 2.1. The Perfect Control assumption does not always hold. However, the kernel and cokernel of the map $\mathrm{Sel}_p(E, L) \to \mathrm{Sel}_p(E, K_\infty)^{\mathrm{Gal}(K_\infty/L)}$ are usually small and bounded independently of L. (This is the "Control Theorem", see for example [M1,G1].) In a case where the Perfect Control assumption does not hold, we can either localize Λ to avoid the support of these kernels and cokernels, or else work with the collection of $\mathrm{Sel}_p(E, K_\infty)^{\mathrm{Gal}(K_\infty/L)}$ instead of the classical Selmer groups $\mathrm{Sel}_p(E, L)$.

See Appendix A for a discussion of sufficient conditions that will guarantee that the Perfect Control assumption holds.

Lemma 2.2. *If the Perfect Control assumption holds and $K \subset L \subset K_\infty$, then*

$$\mathcal{S}_p(E, K_\infty) \otimes_\Lambda \Lambda_L \cong \mathcal{S}_p(E, L),$$

$$\mathcal{S}_p(E, L) \otimes_{\Lambda_L} (\Lambda_L/\mathfrak{m}_L) \cong \mathcal{S}_p(E, K) \otimes \mathbf{Z}/p\mathbf{Z},$$

where $\mathfrak{m}_L$ is the maximal ideal of Λ_L. In particular $\mathcal{S}_p(E, L)$ is a finitely generated Λ_L-module.

Proof. The two isomorphisms are clear, and then since $\mathcal{S}_p(E, K) \otimes \mathbf{Z}/p\mathbf{Z}$ is finite, Nakayama's Lemma shows that $\mathcal{S}_p(E, L)$ is finitely generated over Λ_L. □

Lemma 2.3. *Suppose L is a finite extension of K in K_∞.*

(i) *There is a canonical isomorphism*

$$\mathcal{S}_p(E, L)_{\mathrm{tors}} \cong Ш(E, L)[p^\infty]/Ш(E, L)[p^\infty]_{\mathrm{div}},$$

where $Ш(E, L)[p^\infty]_{\mathrm{div}}$ is the maximal divisible subgroup of $Ш(E, L)[p^\infty]$. If $Ш(E, L)[p^\infty]$ is finite then this isomorphism becomes

$$\mathcal{S}_p(E, L)_{\mathrm{tors}} \cong Ш(E, L)[p^\infty].$$

(ii) *There is a canonical inclusion*

$$(E(L)/E(L)_{\mathrm{tors}}) \otimes \mathbf{Z}_p \hookrightarrow \mathrm{Hom}(\mathcal{S}_p(E, L), \mathbf{Z}_p)$$

which is an isomorphism if $Ш(E, L)[p^\infty]$ is finite.

Proof. Clear. (In the isomorphism of (i) we have used the Cassels pairing to identify $Ш(E, L)[p^\infty]/Ш(E, L)[p^\infty]_{\mathrm{div}}$ with its Pontrjagin dual.) □

Definition 2.4. If $K \subset L \subset K_\infty$ we define the Λ_L-module of universal norms

$$\mathcal{M}_p(E, L) := \varprojlim \mathrm{Hom}(\mathcal{S}_p(E, F), \mathbf{Z}_p),$$

the inverse limit (with respect to the maps induced by corestriction) being taken over finite extensions F of K in L. We have

$$\mathcal{M}_p(E,L) \supset \varprojlim (E(F)/E(F)_{\mathrm{tors}}) \otimes \mathbf{Z}_p$$

(inverse limit with respect to the trace maps) by Lemma 2.3(ii), with equality if $Ш(E,F)[p^\infty]$ is finite for the intermediate fields F.

If L/K is finite then $\mathcal{M}_p(E,L) = \mathrm{Hom}(\mathcal{S}_p(E,L), \mathbf{Z}_p) \supset (E(L)/E(L)_{\mathrm{tors}}) \otimes \mathbf{Z}_p$, and if further $Ш(E,L)[p^\infty]$ is finite then $\mathcal{M}_p(E,L) = (E(L)/E(L)_{\mathrm{tors}}) \otimes \mathbf{Z}_p$.

Remark 2.5. When L/K is infinite, one often expects that $\mathcal{M}_p(E,L) = 0$ (for example, when L contains the cyclotomic $\mathbf{Z}_p$-extension of K). However, $\mathcal{M}_p(E,L)$ can be nonzero for certain infinite extensions L/K, for example [Co,V] when K is imaginary quadratic and L is the anti-cyclotomic $\mathbf{Z}_p$-extension of K. See [MR3] for a further discussion of this.

Proposition 2.6. *If the Perfect Control assumption holds and $K \subset L \subset K_\infty$, then*

$$\mathrm{Hom}_\Lambda(\mathcal{S}_p(E,K_\infty), \Lambda_L)^\iota = \mathrm{Hom}_{\Lambda_L}(\mathcal{S}_p(E,L), \Lambda_L)^\iota \cong \mathcal{M}_p(E,L).$$

Proof. The first equality is Lemma 2.2.

If L/K is finite, then Lemma B.1 of Appendix B shows that

$$\mathrm{Hom}_{\Lambda_L}(\mathcal{S}_p(E,L), \Lambda_L)^\iota \cong \mathrm{Hom}_{\mathbf{Z}_p}(\mathcal{S}_p(E,L), \mathbf{Z}_p),$$

which proves the proposition in this case. The general case follows by passing to the inverse limit. □

3. Hermitian and skew-Hermitian modules

Definition 3.1. A *semi-linear* Λ-module is a Λ-module M endowed with an involution $i : M \to M$ such that $i(\lambda m) = \iota(\lambda) \cdot i(m)$ for all $\lambda \in \Lambda$ and $m \in M$. Equivalently, we may think of the involution i as a Λ-module isomorphism $i : M \to M^\iota$ such that $i^\iota \circ i : M \to (M^\iota)^\iota = M$ is the identity. We refer to such a pair (M,i) as a *semi-linear module*, for short. The involution ι of the free Λ-module Λ endows that module with a natural semi-linear structure. If M is a Λ-module and N is a semi-linear Λ-module, the Λ-module $\mathrm{Hom}_\Lambda(M,N)$ inherits a semi-linear structure as follows. For $f \in \mathrm{Hom}(M,N)$ let $i(f) \in \mathrm{Hom}_\Lambda(M,N)$ be given by $i(f) := i \circ f$. For a free Λ-module Φ of finite rank, by the *semi-linear conjugate Λ-dual* Φ^* of Φ we mean the Λ-module $\Phi^* := \mathrm{Hom}_\Lambda(\Phi^\iota, \Lambda)$ with the semi-linear structure as given above.

If $I \subset \Lambda$ is an ideal that is stable under the action ι then the quotient Λ/I inherits an involution compatible with ι; we denote it again ι.

Example 3.2. If $K \subset L \subset K_\infty$, let $I_L \subset \Lambda$ be the closed ideal generated by all elements of the form $h-1 \in \Lambda$ for $h \in \mathrm{Gal}(K_\infty/L)$. That is, I_L is the kernel of the natural projection $\Lambda \to \Lambda_L$. We have a canonical isomorphism of Λ_L-modules

$$\mathrm{Gal}(K_\infty/L) \otimes_{\mathbf{Z}_p} \Lambda_L \cong I_L/I_L^2$$

characterized by the property that the element $h \otimes 1$ is sent to $h-1$ modulo I_L^2 for all $h \in \mathrm{Gal}(K_\infty/L)$.

Definition 3.3. If Φ is a Λ-module, and M a semi-linear Λ-module, a pairing

$$h : \Phi \otimes_\Lambda \Phi^\iota \to M$$

is called *Hermitian* if

$$h(a \otimes b) = +\iota(h(b \otimes a)),$$

and *skew-Hermitian* if

$$h(a \otimes b) = -\iota(h(b \otimes a)).$$

A *skew-Hermitian* Λ*-module* is a free Λ-module of finite rank with a skew-Hermitian Λ-valued pairing, where we view Λ as semi-linear Λ-module via its involution ι.

4. Derived pairings

Suppose from now on that Φ is a skew-Hermitian Λ-module as in Definition 3.3, with a nondegenerate Λ-valued skew-Hermitian pairing $h : \Phi \otimes \Phi^\iota \to \Lambda$. Such a pairing corresponds to an injective Λ-homomorphism (which we will also denote by h)

$$h : \Phi \longrightarrow \Phi^*$$

and the skew-Hermitian property of the pairing is then equivalent to the fact that the induced map

$$\Phi^\iota = \mathrm{Hom}(\Phi^*, \Lambda) \xrightarrow{h^*} \mathrm{Hom}(\Phi, \Lambda) = (\Phi^*)^\iota$$

is identified with $-h$ under the canonical isomorphism

$$\mathrm{Hom}_\Lambda(\Phi, \Phi^*) = \mathrm{Hom}_\Lambda(\Phi^\iota, (\Phi^*)^\iota).$$

Let S denote the cokernel of h, so that

$$0 \longrightarrow \Phi \xrightarrow{h} \Phi^{*} \longrightarrow S \longrightarrow 0 \tag{4.1}$$

is a free resolution of the Λ-module S, giving, in particular that the A-modules $\mathrm{Tor}^{i}_{\Lambda}(S, A)$ and $\mathrm{Ext}^{i}_{\Lambda}(S, A)$ vanish for every Λ-algebra A and every $i > 1$. If $K \subset L \subset K_{\infty}$, put

$$M(L) := \mathrm{Tor}^{1}_{\Lambda}(S, \Lambda_L) = \ker(h \otimes \Lambda_L),$$

$$S(L) := S \otimes_{\Lambda} \Lambda_L = \mathrm{coker}(h \otimes \Lambda_L)$$

(the letter M is chosen to remind us of *Mordell–Weil*, while the letter S is chosen to remind us of *Selmer;* see Section 7). These definitions give us an exact sequence of Λ_L-modules

$$0 \longrightarrow M(L) \longrightarrow \Phi \otimes_{\Lambda} \Lambda_L \xrightarrow{h \otimes \Lambda_L} \Phi^{*} \otimes_{\Lambda} \Lambda_L \longrightarrow S(L) \longrightarrow 0. \tag{4.2}$$

We have that $h^{*} = -h$ on Φ^{ι}, and using this along with (4.2) (for the upper exact sequence) and (4.1) (for the lower exact sequence) gives a commutative diagram of Λ_L-modules,

$$\begin{tikzcd}
0 \arrow[r] & M(L)^{\iota} \arrow[r] & \Phi^{\iota} \otimes \Lambda_L \arrow[r, "-h"] \arrow[d, "\cong"] & (\Phi^{*})^{\iota} \otimes \Lambda_L \arrow[r] \arrow[d, "\cong"] & S(L)^{\iota} \arrow[r] & 0 \\
0 \arrow[r] & \mathrm{Hom}_{\Lambda}(S, \Lambda_L) \arrow[r] & \mathrm{Hom}(\Phi^{*}, \Lambda_L) \arrow[r, "h^{*}"] & \mathrm{Hom}(\Phi, \Lambda_L) \arrow[r] & \mathrm{Ext}^{1}_{\Lambda}(S, \Lambda_L) \arrow[r] & 0.
\end{tikzcd}$$

Thus we obtain canonical isomorphisms

$$M(L)^{\iota} \cong \mathrm{Hom}_{\Lambda}(S, \Lambda_L), \tag{4.3}$$

$$S(L)^{\iota} \cong \mathrm{Ext}^{1}_{\Lambda}(S, \Lambda_L). \tag{4.4}$$

Recall (Example 3.2) that I_L is the kernel of the map $\Lambda \twoheadrightarrow \Lambda_L$. Tensoring the exact sequence

$$0 \longrightarrow I_L \longrightarrow \Lambda \longrightarrow \Lambda_L \longrightarrow 0$$

with S gives a canonical injection

$$0 \longrightarrow \mathrm{Tor}^1_\Lambda(S, \Lambda_L) \longrightarrow I_L \otimes_\Lambda S$$

and composing this with the natural pairing

$$(I_L \otimes_\Lambda S) \otimes_\Lambda \mathrm{Hom}_\Lambda(S, \Lambda_L) \longrightarrow I_L \otimes_\Lambda \Lambda_L = I_L/I_L^2$$

we get the pairing

$$\mathrm{Tor}^1_\Lambda(S, \Lambda_L) \otimes_{\Lambda_L} \mathrm{Hom}_\Lambda(S, \Lambda_L) \longrightarrow I_L/I_L^2.$$

Now, using the definition of $M(L)$ and (4.3), we obtain the pairing:

$$M(L) \otimes_{\Lambda_L} M(L)^\iota \longrightarrow I_L/I_L^2 \cong \mathrm{Gal}(K_\infty/L) \otimes \Lambda_L. \tag{4.5}$$

The pairing (4.5) is skew-Hermitian with respect to the involution on I_L/I_L^2 induced by ι. The identification $I_L/I_L^2 \cong \mathrm{Gal}(K_\infty/L) \otimes \Lambda_L$ sends this involution to $-1 \otimes \iota$ on $\mathrm{Gal}(K_\infty/L) \otimes \Lambda_L$. By Proposition B.2 of Appendix B, if L/K is finite then the pairing (4.5) induces a *symmetric* pairing

$$M(L) \otimes_{\mathbf{Z}_p} M(L) \to \mathrm{Gal}(K_\infty/L). \tag{4.5$'$}$$

Remark 4.1. Here is a more direct description of the pairing (4.5). Let $\langle\ ,\ \rangle$ denote the skew-Hermitian pairing corresponding to h, and if $m \in M(L) \subset \Phi/I_L\Phi$ let $\tilde{m} \in \Phi$ denote any choice of lifting of m. Then, from the definition of $M(L)$, we have $\langle \tilde{m}, x\rangle \in I_L \subset \Lambda$ for every $x \in \Phi$. If $m_1, m_2 \in M(L)$ we see that the value $\langle \tilde{m}_1, \tilde{m}_2\rangle \in I_L$, when taken modulo I_L^2, is dependent only upon the elements $m_1, m_2 \in M(L)$ and independent of the choices of liftings $\tilde{m}_1, \tilde{m}_2 \in \Phi$. Then the Λ_L-bilinear pairing (4.5) is defined by the rule

$$m_1 \otimes m_2 \mapsto \langle \tilde{m}_1, \tilde{m}_2\rangle (\mathrm{mod}\ I_L^2) \in I_L/I_L^2.$$

Let $\mathcal{K}_L$ denote the total ring of fractions of Λ_L. If M is a Λ_L-module, M_{tors} will denote the kernel of the natural map $M \to M \otimes \mathcal{K}_L$ (the set of elements of M annihilated by a non-zero-divisor of Λ_L).

Applying the functor $\mathrm{Hom}_\Lambda(S, \cdot)$ to the exact sequence of Λ-modules

$$0 \to \Lambda_L \to \mathcal{K}_L \to \mathcal{K}_L/\Lambda_L \to 0,$$

we obtain an exact sequence

$$\mathrm{Hom}_{\Lambda_L}(S(L), \mathcal{K}_L) \to \mathrm{Hom}_{\Lambda_L}(S(L), \mathcal{K}_L/\Lambda_L) \to \mathrm{Ext}^1_\Lambda(S, \Lambda_L) \to \mathrm{Ext}^1_\Lambda(S, \mathcal{K}_L).$$

The kernel of the right-hand map contains $\mathrm{Ext}^1_\Lambda(S, \Lambda_L)_{\mathrm{tors}}$, and there is a natural map from the cokernel of the left-hand to $\mathrm{Hom}_{\Lambda_L}(S(L)_{\mathrm{tors}}, \mathcal{K}_L/\Lambda_L)$. Thus using (4.4) we get a map

$$S(L)^\iota_{\mathrm{tors}} \cong \mathrm{Ext}^1_\Lambda(S, \Lambda_L)_{\mathrm{tors}} \longrightarrow \mathrm{Hom}_{\Lambda_L}(S(L)_{\mathrm{tors}}, \mathcal{K}_L/\Lambda_L)$$

and hence a Λ-bilinear pairing

$$S(L)_{\mathrm{tors}} \otimes_{\Lambda_L} S(L)^\iota_{\mathrm{tors}} \longrightarrow \mathcal{K}_L/\Lambda_L. \tag{4.6}$$

The pairing (4.6) is skew-Hermitian with respect to the involution on $\mathcal{K}_L/\Lambda_L$ induced by ι. If L/K is finite, the identification $\mathcal{K}_L/\Lambda_L \cong \mathbf{Q}_p/\mathbf{Z}_p \otimes \Lambda_L$ sends this involution to $1 \otimes \iota$ on $\mathbf{Q}_p/\mathbf{Z}_p \otimes \Lambda_L$. By Proposition B.2 of Appendix B, the pairing (4.6) induces a skew-symmetric pairing

$$S(L)_{\mathrm{tors}} \otimes_{\mathbf{Z}_p} S(L)_{\mathrm{tors}} \to \mathbf{Q}_p/\mathbf{Z}_p. \tag{4.6$'$}$$

Remark 4.2. Here is a more direct description of the pairing (4.6). Suppose $s \in S(L)_{\mathrm{tors}}$, say $as = 0$ with a nonzero-divisor $a \in \Lambda_L$. From the definition (4.2) of $S(L)$, we can choose $\tilde{s} \in \Phi \otimes \Lambda_L$ and $\tilde{s}^* \in \Phi^* \otimes \Lambda_L$ such that $\tilde{s}^*$ lifts s (under (4.2)) and $\tilde{s}$ lifts $a\tilde{s}^*$. Similarly, if $t \in S(L)^\iota_{\mathrm{tors}}$ and $bt = 0$ we can lift to $\tilde{t} \in \Phi^\iota \otimes \Lambda_L$ whose image in $(\Phi^*)^\iota \otimes \Lambda_L$ is b times a lift of t.

Let $\langle\, , \rangle_L$ denote the skew-Hermitian pairing $(\Phi \otimes \Lambda_L) \otimes (\Phi^\iota \otimes \Lambda_L) \to \Lambda_L$ induced by h. Then the pairing (4.6) is given by

$$s \otimes t \mapsto (ab)^{-1}\langle \tilde{s}, \tilde{t}\rangle_L (\mathrm{mod}\, \Lambda_L) \in \mathcal{K}_L/\Lambda_L.$$

This is independent of all the choices that were made.

In summary, given a skew-Hermitian module Φ over Λ, with the hypotheses above, for every extension L of K in K_∞ we get a Λ_L-bilinear pairing (4.5) on $M(L)$ with values in I_L/I_L^2 and a Λ_L-bilinear pairing (4.6) on $S(L)_{\mathrm{tors}}$ with values in $\mathcal{K}/\Lambda_L$.

5. Complexes

Fix a noetherian local ring R with maximal ideal $\mathfrak{m}$ and residue field $\Bbbk = R/\mathfrak{m}$. We will be interested in the case where $R = \Lambda$, but the results of this section are more general.

Definition 5.1. By a *complex* of R-modules we mean an infinite *co-complex*, i.e., a sequence of R-modules and R-homomorphisms

$$C^\bullet : \quad \ldots C^{-n} \to C^{1-n} \to \cdots \to C^n \to C^{n+1} \to \ldots$$

with (co-)boundary operators raising degrees by 1 and such that the composition of any two successive coboundaries vanishes. For an integer k, the complex $C^\bullet[k]$ will denote the complex $C^\bullet$ shifted by k

$$\dots (C')^{-n} \to (C')^{1-n} \to \cdots \to (C')^{n} \to (C')^{n+1} \to \dots$$

where $(C')^m := C^{m+k}$.

If $C^\bullet$ is a complex, its R-dual $\mathrm{Hom}(C^\bullet, R)$ is again a complex, where, as usual the gradation on $\mathrm{Hom}(C^\bullet, R)$ is given by $\mathrm{Hom}(C^\bullet, R)^n := \mathrm{Hom}(C^{-n}, R)$.

If all of the modules C^n are free of finite rank over R, then the natural identification of a free R-module of finite rank with its double R-dual,

$$M \xrightarrow{\sim} \mathrm{Hom}(\mathrm{Hom}(M, R), R) \quad \text{by } m \mapsto \{\phi \mapsto \phi(m)\}$$

extends to a natural identification of $C^\bullet$ with its double R-dual.

Let $\mathcal{C} = \mathcal{C}(R)$ denote the category of complexes of R-modules, where morphisms are morphisms (of degree zero) of complexes of R-modules. A *quasi-isomorphism* $f : C^\bullet \to D^\bullet$ of complexes is a morphism that induces an isomorphism on cohomology $H^*(f) : H^*(C^\bullet) \xrightarrow{\sim} H^*(D^\bullet)$.

Definition 5.2. A two-term complex of free R-modules of finite rank, $F^\bullet$, concentrated in degrees 1 and 2

$$\cdots \to 0 \to F^1 \xrightarrow{\partial} F^2 \to 0 \to \cdots$$

will be called a *basic complex* if the coboundary homomorphism ∂ is injective and if, when we form the short exact sequence of R-modules,

$$0 \to F^1 \to F^2 \to H \to 0,$$

the induced homomorphism $F^2 \otimes_R \Bbbk \to H \otimes_R \Bbbk$ is an isomorphism. (The latter condition is equivalent to requiring that the image of F^1 is contained in $\mathfrak{m}F^2$.)

Such a basic complex has cohomology concentrated in degree 2 with $H^2(F^\bullet) = H$.

Lemma 5.3. *Suppose that $C^\bullet$ is a complex of free R-modules concentrated in degrees 1 and 2, with injective coboundary map $C^1 \xrightarrow{\partial} C^2$. Then $C^\bullet$ is quasi-isomorphic to a basic complex.*

Proof. Let $H = H^2(C^\bullet)$ and consider the exact sequence

$$C^1 \otimes \Bbbk \xrightarrow{\partial \otimes \Bbbk} C^2 \otimes \Bbbk \longrightarrow H \otimes \Bbbk \longrightarrow 0.$$

Let $\bar{\Sigma}_2$ be a $\Bbbk$-basis for $\mathrm{image}(\partial \otimes \Bbbk) = \ker(C^2 \otimes \Bbbk \to H \otimes \Bbbk)$. Pull each element of $\bar{\Sigma}_2$ back to $C^1 \otimes \Bbbk$ via $\partial \otimes \Bbbk$ and then lift each of these elements to C^1. Denote the resulting sets by $\bar{\Sigma}_1 \subset C^1 \otimes \Bbbk$ and $\Sigma_1 \subset C^1$, and let $\Sigma_2 := \partial(\Sigma_1) \subset C^2$, a set lifting $\bar{\Sigma}_2$.

For $i = 1, 2$ let $D^i \subset C^i$ be the Λ-module generated by Σ_i, and let $B^i := C^i / D^i$. Complete $\bar{\Sigma}_i$ to a $\Bbbk$-basis $\bar{\Sigma}_i \cup \bar{\Sigma}'_i$ of $C^i \otimes \Bbbk$, and lift $\bar{\Sigma}'_i$ to $\Sigma'_i \subset C^i$. By Nakayama's Lemma $\Sigma_i \cup \Sigma'_i$ generates C^i, and since C^i is free (of rank $\dim_{\Bbbk}(C^i \otimes \Bbbk)$) $\Sigma_i \cup \Sigma'_i$ must be a Λ-basis of C^i. Hence Σ'_i projects to a Λ-basis of B^i, and in particular B^i is free over Λ.

The map $\partial : C^1 \to C^2$ induces an injection $B^1 \to B^2$ with cokernel equal to H. Since by definition $D^2 \otimes \Bbbk$ and $C^1 \otimes \Bbbk$ have the same image in $C^2 \otimes \Bbbk$, the induced map $B^1 \otimes \Bbbk \to B^2 \otimes \Bbbk$ is the zero map. Thus, if we set $B^i := 0$ for $i \neq 1, 2$ then $B^\bullet$ is a basic complex, and the projection map $C^\bullet \to B^\bullet$ is a quasi-isomorphism. □

Lemma 5.4. *Suppose that $F^\bullet$ and $G^\bullet$ are basic complexes, and $f : H^2(F^\bullet) \to H^2(G^\bullet)$ is an R-homomorphism.*

(i) *There is a morphism of complexes $\phi : F^\bullet \to G^\bullet$ such that $H^2(\phi) = f$, and any two such morphisms of complexes are homotopic.*

(ii) *If f is an isomorphism then the morphism ϕ of* (i) *is an isomorphism of complexes.*

Proof. We are given a diagram

$$\begin{tikzcd}
0 \arrow[r] & F^1 \arrow[r] & F^2 \arrow[r] & H^2(F^\bullet) \arrow[r] \arrow[d, "f"] & 0 \\
0 \arrow[r] & G^1 \arrow[r] & G^2 \arrow[r] & H^2(G^\bullet) \arrow[r] & 0.
\end{tikzcd}$$

Since F^2 is free we can pull f back to a map $\phi_2 : F^2 \to G^2$, which in turn restricts to a map $\phi_1 : F^1 \to G^1$. This gives a morphism of complexes $\phi : F^\bullet \to G^\bullet$ with $H^2(\phi) = f$, and it is clear that any two such morphisms are homotopic.

Using the definition of basic complex we see that $\ker(\phi_2 \otimes \Bbbk) = \ker(f \otimes \Bbbk)$ and $\mathrm{coker}(\phi_2 \otimes \Bbbk) = \mathrm{coker}(f \otimes \Bbbk)$. Thus, if f is an isomorphism then so is $\phi_2 \otimes \Bbbk$, and by Nakayama's Lemma so is ϕ_2 (and therefore ϕ_1 as well). This proves (ii). □

Definition 5.5. Let $\mathcal{D} = \mathcal{D}(R)$ denote the derived category of complexes of R-modules. That is, $\mathcal{D}(R)$ is the category usually denoted $D(A)$ where A is the abelian category of R-modules (see for example [Hart]).

Recall that $\mathcal{D}$ is constructed as follows ([Hart] Chapter I). Let $\mathcal{K} = \mathcal{K}(A)$ be the category whose objects are complexes of R-modules, and whose morphisms are homotopy classes of morphisms of complexes. The category $\mathcal{D}$ is obtained from $\mathcal{K}$ by "localizing quasi-isomorphisms." That is, every morphism in $\mathcal{K}$ that induces an isomorphism on cohomology groups becomes an isomorphism in the category $\mathcal{D}$. The categories $\mathcal{K}$ and $\mathcal{D}$ are triangulated categories.

Corollary 5.6. *Suppose that $F^\bullet$ and $G^\bullet$ are basic complexes, and $\psi : F^\bullet \to G^\bullet$ is an isomorphism in the derived category $\mathcal{D}$. Then there is an isomorphism of complexes (i.e., in the category $\mathcal{C}$) $\phi : F^\bullet \xrightarrow{\sim} G^\bullet$ that gives rise to ψ. The isomorphism ϕ is unique up to homotopy.*

Proof. The $\mathcal{D}$-isomorphism ψ induces an isomorphism $f : H^2(F^\bullet) \to H^2(G^\bullet)$. The desired isomorphism of complexes is then provided by Lemma 5.4. □

6. Skew-Hermitian structures on complexes

Keep the noetherian local ring R of §5, and suppose further that R possesses an involution $\iota : R \to R$. Denote by $M \mapsto M^\iota$ the induced involution on the categories of R-modules, complexes of R-modules, etc.

Definition 6.1. Suppose $C^\bullet$ is an R-complex of free R-modules of finite rank.

A *skew-Hermitian, degree n, perfect pairing in the category $\mathcal{C}$* on $C^\bullet$ is an isomorphism

$$\phi : C^\bullet \to \mathrm{Hom}_R(C^\bullet, R)^\iota[-n]$$

of R-complexes such that after the natural identification of the complex $C^\bullet$ with its R-double dual, the morphism $\mathrm{Hom}_R(\phi^\iota)$, which may be viewed as a morphism

$$\mathrm{Hom}_R(\phi^\iota) : C^\bullet \to \mathrm{Hom}_R(C^\bullet, R)^\iota[-n],$$

is equal to $-\phi$.

A *skew-Hermitian, degree n, perfect pairing in the category $\mathcal{D}$* on $C^\bullet$ is an isomorphism

$$\phi : C^\bullet \to \mathrm{Hom}_R(C^\bullet, R)^\iota[-n]$$

in $\mathcal{D}$ such that after the natural identification of the complex $C^\bullet$ with its R-double dual, the morphism $\mathrm{Hom}_R(\phi^\iota)$ is equal in $\mathcal{D}$ to $-\phi$.

We have the evident notion of *equivalence* of skew-Hermitian, degree n, perfect pairings, for each of the two categories $\mathcal{C}$ and $\mathcal{D}$.

An isomorphism $C^\bullet \to E^\bullet$ in either of the two categories transports—in the evident manner—skew-Hermitian, degree n, perfect pairings on $C^\bullet$ to skew-Hermitian, degree n, perfect pairings on $E^\bullet$.

Corollary 6.2. *If a basic complex $F^\bullet$ possesses a skew-Hermitian degree* 3 *perfect pairing*

$$\psi : F^\bullet \to \mathrm{Hom}_R(F^\bullet, R)^\iota[-3]$$

in the category $\mathcal{D}$ then there is a degree 3 *perfect pairing*

$$\phi : F^\bullet \to \mathrm{Hom}_R(F^\bullet, R)^\iota[-3]$$

in the category $\mathcal{C}$ of R-complexes, inducing ψ, such that the morphisms $\mathrm{Hom}_R(\phi^\iota)$ *and* $-\phi$ *in $\mathcal{C}$ are homotopic. The degree* 3 *perfect pairing ϕ with these properties is unique up to homotopy.*

Proof. If $F^\bullet$ is a basic complex, then so is $\mathrm{Hom}_R(F^\bullet, R)^\iota[-3]$. Thus, the corollary is immediate from Corollary 5.6. □

Let Φ be a skew-Hermitian R-module as defined in Definition 3.3 (for the case $R = \Lambda$). Thus, Φ is a free R-module of finite rank, endowed with a skew-Hermitian pairing, i.e., an R-homomorphism

$$h : \Phi \to \mathrm{Hom}_R(\Phi^\iota, R)$$

such that the induced homomorphism $\mathrm{Hom}(h^\iota)$ is identified with

$$-h : \Phi \to \mathrm{Hom}_R(\Phi^\iota, R)$$

when we identify $\mathrm{Hom}_R(\mathrm{Hom}_R(\Phi, R), R) \cong \Phi$. Recall that $\Phi^* := \mathrm{Hom}_R(\Phi^\iota, R) = \mathrm{Hom}_R(\Phi, R)^\iota$, and let $h^* := \mathrm{Hom}(h^\iota)$. We have natural identifications of "double-duals" $\Phi^{**} = \Phi$ and $h^{**} = h$.

Definition 6.3. Given a skew-Hermitian R-module Φ, we form a complex $\Phi^\bullet$, concentrated in degrees 1 and 2, by putting $\Phi^1 := \Phi$, $\Phi^2 := \Phi^*$, and setting the coboundary $\partial : \Phi^1 \to \Phi^2$ to be $h : \Phi \to \Phi^*$.

We will say that Φ is a *basic skew-Hermitian module* if h is injective, and $h \otimes \Bbbk = 0$ (or equivalently, if h is injective and $h(\Phi) \subset \mathfrak{m}\Phi^*$). Thus, Φ is basic if and only if $\Phi^\bullet$ is a basic complex.

For example, if R is an integral domain, then Φ is basic if and only if

- the skew-Hermitian pairing over the field of fractions of R obtained from Φ is nondegenerate,
- there are no unimodular pieces that can be split off from Φ (i.e., Φ is minimal for our purposes).

Suppose Φ is a basic skew-Hermitian module, and let $N^\bullet := \mathrm{Hom}_R(\Phi^\bullet, R)^\iota[-3]$. We have canonical identifications

$$N^1 = \mathrm{Hom}_R(\mathrm{Hom}_R(\Phi, R), R) \cong \Phi, \quad N^2 = \mathrm{Hom}_R(\Phi, R)^\iota \cong \Phi^*,$$

where the coboundary is given by $h^* = -h$. The isomorphism of basic complexes $j : \Phi^\bullet \to N^\bullet$ given by putting $j^1 = -1$ and $j^2 = +1$ (after the identifications we have just made) is a skew-Hermitian degree 3 perfect pairing of the basic R-complex $\Phi^\bullet$.

Definition 6.4. A skew-Hermitian, degree 3, perfect pairing on a complex $C^\bullet$ in the category $\mathcal{D}$ *comes from the basic skew-Hermitian R-module* Φ if Φ is a basic skew-Hermitian R-module and there is an isomorphism in the derived category $\mathcal{D}$

$$\Phi^\bullet \overset{\sim}{\to} C^\bullet$$

such that the skew-Hermitian, degree 3, perfect pairing on $C^\bullet$ is the one obtained by transport of structure from the pairing on $\Phi^\bullet$.

Proposition 6.5. *Suppose that the residual characteristic of R is not* 2, *that* $C^\bullet$ *is a complex of free R-modules concentrated in degrees* 1 *and* 2, *and the coboundary map* $C^1 \to C^2$ *is injective. Then every skew-Hermitian, degree* 3, *perfect pairing on* $C^\bullet$ *in the category* $\mathcal{D}$ *comes from a basic skew-Hermitian R-module* Φ.

Proof. By Lemma 5.3, $C^\bullet$ is isomorphic in $\mathcal{D}$ to a basic complex $F^\bullet$, so we may as well assume that $C^\bullet$ is a basic complex in the statement of the proposition. By Corollary 6.2 we can lift the skew-Hermitian degree 3 pairing on $C^\bullet$ in $\mathcal{D}$ to a skew-Hermitian degree 3 pairing on $C^\bullet$ in $\mathcal{C}$, so in particular we get isomorphisms α and β in a commutative diagram

$$\begin{tikzcd}
C^1 \arrow[r, "\partial"] \arrow[d, "\alpha"'] & C^2 \arrow[d, "\beta"] \\
(C^2)^* \arrow[r, "\partial^*"] & (C^1)^*
\end{tikzcd}$$

Passing to the dual, we get the diagram

$$\begin{tikzcd}
C^1 \arrow[r, "\partial"] \arrow[d, "\beta^*"'] & C^2 \arrow[d, "\alpha^*"] \\
(C^2)^* \arrow[r, "\partial^*"] & (C^1)^*
\end{tikzcd}$$

By Corollary 6.2, these two maps of complexes are homotopic (after replacing (α, β) by $(-\alpha, -\beta)$ in the first diagram), so there exists an R-homomorphism $w : C^2 \to (C^2)^*$ such that

$$\alpha^* = -\beta + \partial^* w \quad \text{and} \quad \beta^* = -\alpha + w\partial.$$

This implies (among other things) that $\partial^* w = (w\partial)^*$.

If the residual characteristic of R is different from 2, we can modify the morphism of complexes (α, β) by a homotopy, replacing (α, β) by (α', β') where

$$\alpha' := \alpha - w\partial/2 \quad \text{and} \quad \beta' := \beta - \partial^* w/2.$$

Since $\alpha^* + \beta = \partial^* w$, we get that

$$(\alpha')^* + \beta' = \partial^* w - (w\partial)^*/2 - \partial^* w/2 = 0.$$

It follows that the perfect degree 3 skew-Hermitian pairing in the derived category $\mathcal{D}$ comes from the pairing on $C^\bullet$ in the category $\mathcal{C}$ described by the diagram

$$\begin{tikzcd}
C^1 \arrow[r, "\partial"] \arrow[d, "\alpha'"'] & C^2 \arrow[d, "-(\alpha')^*"] \\
(C^2)^* \arrow[r, "\partial^*"] & (C^1)^*
\end{tikzcd}$$

Now put $\Phi := C^1$, and consider the homomorphism

$$h := (\alpha')^* \circ \partial : \Phi \to \Phi^*.$$

We have that $h^* = -h$, giving Φ the structure of a basic skew-Hermitian R-module. The basic complex $\Phi^\bullet$ is isomorphic to the basic complex $C^\bullet$ by the mapping

$$(1, (\alpha')^*) : C^\bullet \to \Phi^\bullet$$

and this isomorphism respects skew-Hermitian structures. □

Proposition 6.6. *Suppose that the residual characteristic of R is not 2. Suppose further that Φ and Ψ are basic skew-Hermitian modules, and there is an isomorphism $\Phi^\bullet \xrightarrow{\sim} \Psi^\bullet$ in the derived category $\mathcal{D}$ that induces an equivalence of degree 3 perfect skew-Hermitian pairings. Then Φ and Ψ are isomorphic as skew-Hermitian modules.*

In other words, if a skew-Hermitian, degree 3, *perfect pairing on a complex* $C^\bullet$ *in* $\mathcal{D}$ *comes from a basic skew-Hermitian module* Φ, *then* Φ (*with its skew-Hermitian structure*) *is unique up to* (*noncanonical*) *isomorphism.*

Proof. By Corollary 5.6 there is an actual isomorphism of complexes $\Phi^\bullet \xrightarrow{\sim} \Psi^\bullet$ giving rise to the isomorphism in $\mathcal{D}$. In other words there is a commutative diagram

$$\begin{tikzcd}
\Phi \arrow[r, "h"] \arrow[d, "\alpha"'] & \Phi^* \arrow[d, "\beta"] \\
\Psi \arrow[r, "g"] & \Psi^*
\end{tikzcd}$$

with isomorphisms α, β. Further, since the isomorphism in $\mathcal{D}$ induces an equivalence of skew-Hermitian pairings, there is a homotopy between this diagram and the "dual diagram" (after replacing $h^* = -h$ and $g^* = -g$ by h and g)

$$\begin{tikzcd}
\Phi \arrow[r, "h"] & \Phi^* \\
\Psi \arrow[r, "g"] \arrow[u, "\beta^*"] & \Psi^* \arrow[u, "\alpha^*"']
\end{tikzcd}$$

Thus, there is a map $w : \Phi^* \to \Psi$ such that

$$(\beta^*)^{-1} = \alpha + wh \quad \text{and} \quad (\alpha^*)^{-1} = \beta + gw. \tag{6.1}$$

In particular, since Φ and Ψ are basic skew-Hermitian modules, we have $h(\Phi) \subset \mathfrak{m}\Phi^*$ and $g(\Psi) \subset \mathfrak{m}\Psi^*$ and so

$$\alpha\beta^* \equiv 1_\Phi (\operatorname{mod} \mathfrak{m} \operatorname{Hom}(\Phi, \Phi)).$$

Suppose now that

$$\alpha\beta^* \equiv 1_\Phi (\operatorname{mod} \mathfrak{m}^k \operatorname{Hom}(\Phi, \Phi)) \tag{6.2}$$

for some $k \geqslant 1$. We will show that we can replace the isomorphism of complexes (α, β) by a homotopic one, congruent to (α, β) modulo $\mathfrak{m}^k$, and satisfying (6.2) with k replaced by $2k$.

Let $\alpha' = \alpha + wh/2$, $\beta' = \beta + gw/2$. Then

$$\begin{aligned}\alpha'(\beta')^* &= (\alpha + wh/2)(\beta^* + (gw)^*/2)\\ &= \alpha\beta^* + wh\beta^*/2 + \alpha(gw)^*/2 + wh(gw)^*/4.\end{aligned}$$

By (6.1) we have

$$\alpha\beta^* + wh\beta^* = 1_\Phi = (1_{\Phi^*})^* = (\beta\alpha^* + gw\alpha^*)^* = \alpha\beta^* + \alpha(gw)^*$$

so $wh\beta^* = \alpha(gw)^*$ and

$$\alpha'(\beta')^* = 1_\Phi + wh(gw)^*/4.$$

By (6.1) and (6.2) we see that $wh \in \mathfrak{m}^k\ \mathrm{Hom}(\Phi, \Psi)$ and $gw \in \mathfrak{m}^k\ \mathrm{Hom}(\Phi^*, \Psi^*)$, so $\alpha'(\beta')^* \equiv 1_\Phi (\mathrm{mod}\, \mathfrak{m}^{2k}\ \mathrm{Hom}(\Phi, \Phi))$.

Proceeding by induction and passing to the limit, we may assume that $\beta^* = \alpha^{-1}$. In other words, the isomorphism of complexes (with skew-Hermitian pairings) $(\alpha, \beta) : \Phi^\bullet \xrightarrow{\sim} \Psi^\bullet$ is induced by the isomorphism $\alpha : \Phi \xrightarrow{\sim} \Psi$. □

Although we will not need it, we have the following corollary.

Corollary 6.7. *Suppose that Φ and Ψ are basic skew-Hermitian modules, with pairings h_Φ and h_Ψ, and let $\mathcal{L} \subset \Lambda$ be the ideal generated by the determinant of h_Φ with respect to any Λ-bases of Φ and Φ^*. If Φ and Ψ are equivalent modulo $\mathcal{L}^2$, then they are equivalent.*

In other words, if there is an isomorphism $\tilde{\rho} : \Psi \otimes (\Lambda/\mathcal{L}^2) \xrightarrow{\sim} \Phi \otimes (\Lambda/\mathcal{L}^2)$ such that $\tilde{h}_\Psi = \tilde{\rho}^*\tilde{h}_\Phi\tilde{\rho}$ (where $\tilde{h}_\Phi = h_\Phi \otimes (\Lambda/\mathcal{L}^2)$ and $\tilde{h}_\Psi = h_\Psi \otimes (\Lambda/\mathcal{L}^2)$), then there is an isomorphism $\rho : \Psi \xrightarrow{\sim} \Phi$ such that *$h_\Psi = \rho^* h_\Phi \rho$.*

Proof. Since Φ and Ψ are free over Λ, we can lift $\tilde{\rho}$ to a map $\alpha : \Psi \to \Phi$. Nakayama's Lemma shows that α is an isomorphism, and we have

$$h_\Psi \equiv \alpha^*\, h_\Phi\, \alpha (\mathrm{mod}\, \mathcal{L}^2\ \mathrm{Hom}(\Psi, \Psi^*)).$$

Let $\lambda \in \Lambda$ be a generator of $\mathcal{L}$. Since $\mathcal{L}$ is the determinant of h_Φ and α, α^* are isomorphisms, there is a homomorphism $g : \Psi^* \to \Psi$ such that $(\alpha^* h_\Phi \alpha)g = \lambda \cdot \mathrm{id}_{\Psi^*}$ and $g(\alpha^* h_\Phi \alpha) = \lambda \cdot \mathrm{id}_\Psi$. Thus

$$h_\Psi\, g \equiv (\alpha^* h_\Phi \alpha) g = \lambda \cdot \mathrm{id}_{\Psi^*} (\mathrm{mod}\, \mathcal{L}^2\ \mathrm{Hom}(\Psi^*, \Psi^*)),$$

so we see that $\lambda^{-1}h_\Psi g \in \mathrm{Hom}(\Psi^*, \Psi^*)$ and $\lambda^{-1}h_\Psi g \equiv \mathrm{id}_{\Psi^*} (\mathrm{mod}\, \mathcal{L}\, \mathrm{Hom}(\Psi^*, \Psi^*))$. Let $\beta = (\lambda^{-1}h_\Psi g)\alpha^* \in \mathrm{Hom}(\Phi^*, \Psi^*)$. Then

$$\beta \equiv \alpha^* (\mathrm{mod}\, \mathcal{L}\, \mathrm{Hom}(\Phi^*, \Psi^*)) \tag{6.3}$$

and

$$\beta\, h_\Phi\, \alpha = (\lambda^{-1}h_\Psi g)(\alpha^* h_\Phi \alpha) = h_\Psi. \tag{6.4}$$

Using the fact that h_Φ and h_Ψ are skew-Hermitian, we obtain from (6.4) two isomorphisms of complexes $\Phi^\bullet \xrightarrow{\sim} \Psi^\bullet$

$$\begin{tikzcd}
\Phi \arrow[r, "h_\Phi"] \arrow[d, "\alpha^{-1}"'] & \Phi^* \arrow[d, "\beta"] \\
\Psi \arrow[r, "h_\Psi"] & \Psi^*
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
\Phi \arrow[r, "h_\Phi"] \arrow[d, "\beta^{*-1}"'] & \Phi^* \arrow[d, "\alpha^*"] \\
\Psi \arrow[r, "h_\Psi"] & \Psi^*
\end{tikzcd} \tag{6.5}$$

It follows from (6.3) that these two morphisms induce the same isomorphism

$$\mathrm{coker}(h_\Phi) \xrightarrow{\sim} \mathrm{coker}(h_\Psi),$$

so by Lemma 5.4(i) they are homotopic. It follows that $\Phi^\bullet$ and $\Psi^\bullet$ are isomorphic in $\mathcal{D}$ as complexes with skew-Hermitian, degree 3, perfect pairings, and so the corollary follows from Proposition 6.6. □

7. Organization

We now return to the elliptic curve E/K and $\mathbf{Z}_p^d$-extension K_∞/K, and we take R to be the Iwasawa algebra Λ. We will make the following hypotheses:

$$p > 2 \text{ and } E \text{ has good ordinary reduction at all primes above } p, \tag{7.1}$$

$$\mathcal{S}_p(E, K_\infty) \text{ is a torsion } \Lambda\text{-module}, \tag{7.2}$$

$$E(K)[p] = 0, \tag{7.3}$$

$$\text{for every prime } v \text{ of bad reduction, } p \nmid [E(K_v) : E_0(K_v)], \tag{7.4}$$

$$\text{the Perfect Control assumption holds} \tag{7.5}$$

(recall that $[E(K_v) : E_0(K_v)]$ is the Tamagawa number in the Birch and Swinnerton–Dyer conjecture for E/K).

Definition 7.1. Let $C^\bullet_{\mathrm{Nek}}$ be Nekovář's *Selmer complex* in the derived category $\mathcal{D}$, the complex denoted $\widetilde{\mathbf{R}\Gamma}_{f,\mathrm{Iw}}(K_\infty/K, T_p(E))$ in [N] §9.7.1, where $T_p(E) := \varprojlim E[p^n]$ is the p-adic Tate module of E.

Remark 7.2. Let S be a finite set of places of K and let $G_{K,S}$ denote the Galois group of K unramified outside S. For the general definition of "Nekovář–Selmer complexes" (of complexes of $G_{K,S}$-modules $X^\bullet$ with local conditions $\Delta(X^\bullet) = \{\Delta(X^\bullet)_v\}_{v\in S}$ imposed) see §6 of [N]. These Nekovář–Selmer complexes are canonical complexes in the appropriate derived category that compute the cohomology of $X^\bullet$ subject to specified local conditions $\Delta(X^\bullet)$. The classical Selmer module of an abelian variety over a number field, with ordinary reduction above p, falls into this rubric (see the preparation for this, in particular "control theorems," discussed in §7 of [N], and the study of such modules in the context of Iwasawa theory in [N] §8. Section 9 of [N] defines the complexes we call $C^\bullet_{\mathrm{Nek}}$ (Definition 7.1 above) with a close study of the self-dualities such complexes enjoy; the relationship between this self-duality and the various derived self-pairings obtained from the self-duality on the level of complexes is studied in [N] §10 (where the classical Cassels–Tate pairing is treated) and §11 (for the classical p-adic height pairing).

Nekovář's complex $C^\bullet_{\mathrm{Nek}}$ is a canonical complex in $\mathcal{D}$, with a skew-Hermitian pairing in $\mathcal{D}$, and with second cohomology

$$H^2(C^\bullet_{\mathrm{Nek}}) = \mathcal{S}_p(E, K_\infty)$$

(see [N] §9.6.7 and §9.7). Under our hypotheses above $C^\bullet_{\mathrm{Nek}}$ has the following additional useful properties.

Theorem 7.3 (*Nekovář*). *Suppose that hypotheses* (7.1-4) *hold. Then* $C^\bullet_{\mathrm{Nek}}$ *can be represented by a complex concentrated in degrees* 1 *and* 2, *with free* Λ*-modules* C^1, C^2 *of finite rank and an injective coboundary map* $C^1 \to C^2$. *Further,* $C^\bullet_{\mathrm{Nek}}$ *has a canonical skew-Hermitian, degree* 3, *perfect pairing in the derived category.*

Proof. By Proposition 9.7.7(iii) of [N], our hypotheses (7.1), (7.3), and (7.4) imply that $C^\bullet_{\mathrm{Nek}}$ can be represented by a complex concentrated in degrees 1 and 2, with free Λ-modules C^1, C^2 of finite rank. The additional hypothesis (7.2) ensures ([N] Proposition 9.7.7(iv)) that the coboundary map $C^1 \to C^2$ is injective.

By [N] Proposition 9.7.3(ii), $C^\bullet_{\mathrm{Nek}}$ has a degree three pairing in the derived category, and by [N] Propositions 9.7.3(iv) and 9.7.7(ii), respectively, the pairing is perfect and skew-Hermitian. □

Definition 7.4. Suppose that Φ is a basic skew-Hermitian Λ-module as in Definition 6.3. Thus Φ is free over Λ of finite rank, with an injective Λ-valued skew-Hermitian pairing

$$h : \Phi \longrightarrow \Phi^*$$

that is the zero map after tensoring with the residue field $\Lambda/\mathfrak{m}$. We will say that Φ *organizes the arithmetic of* E *over* K_∞ if the complex $C^\bullet_{\mathrm{Nek}}$, with its skew-Hermitian pairing, comes from Φ in the sense of Definition 6.4: i.e., if there is an isomorphism $C^\bullet_{\mathrm{Nek}} \xrightarrow{\sim} \Phi^\bullet$ in $\mathcal{D}$ preserving the skew-Hermitian structures. In this case we will call Φ an *organizing module*.

Theorem 7.5. *Suppose that hypotheses* (7.1-4) *hold. Then there is a basic skew-Hermitian module* Φ *that organizes the arithmetic of* E *over* K_∞.

If Ψ *is another organizing module for* E *over* K_∞*, then there is a* (*noncanonical*) *isomorphism* $\Phi \xrightarrow{\sim} \Psi$ *which takes the skew-Hermitian pairing on* Φ *to the one on* Ψ.

Proof. The existence of an organizing module is immediate from Theorem 7.3 and Proposition 6.5. The uniqueness is Proposition 6.6. □

Remark 7.6. Although the organizing module is not unique up to canonical equivalence, there is a canonical rank-one Λ-module, containing a canonical discriminant, defined as follows. If Φ is an organizing module let Δ_Φ be the free, rank-one Λ-module

$$\Delta_\Phi := \det{}_\Lambda \Phi^{-1} = \bigwedge\nolimits^{\mathrm{rank}_\Lambda \Phi} \mathrm{Hom}(\Phi, \Lambda)$$

and $\mathrm{disc}(\Phi)$ the discriminant

$$\mathrm{disc}(\Phi) := \det{}_\Lambda h_\Phi \in \mathrm{Hom}(\det{}_\Lambda \Phi, \det{}_\Lambda \Phi^*) = \det{}_\Lambda \Phi^{-1} \otimes_\Lambda \det{}_\Lambda \Phi^* = \Delta_\Phi \otimes_\Lambda \Delta^\iota_\Phi.$$

Note that $\mathrm{disc}(\Phi)$ is the determinant of the complex $\Phi^\bullet$ as defined in §4 of [D]. In particular $\mathrm{disc}(\Phi) \cong \det(C^\bullet_{\mathrm{Nek}})$ is independent of the organizing module Φ. (Concretely, if Ψ is another organizing module, then the noncanonical isomorphism of Theorem 7.5 induces a *canonical* isomorphism $\Delta_\Phi \otimes_\Lambda \Delta^\iota_\Phi \xrightarrow{\sim} \Delta_\Psi \otimes_\Lambda \Delta^\iota_\Psi$ which sends $\mathrm{disc}(\Phi)$ to $\mathrm{disc}(\Psi)$.)

Theorem 7.7. *Suppose that hypotheses* (7.1-5) *hold and that the basic skew-Hermitian module* Φ *organizes the arithmetic of* E *over* K_∞. *Let*

$$S = \mathrm{coker}(\Phi \xrightarrow{h} \Phi^*) = H^2(\Phi^\bullet).$$

(i) *There are natural isomorphisms*

$$S \cong \mathcal{S}_p(E, K_\infty),$$

and for every intermediate field $K \subset L \subset K_\infty$

$$S \otimes \Lambda_L \cong \mathcal{S}_p(E, L), \quad \mathrm{Tor}^1_\Lambda(S, \Lambda_L) \cong \mathcal{M}_p(E, L),$$

where $\mathcal{M}_p(E, L)$ *is the universal norm module of Definition* 2.4.

(ii) *If L is a finite extension of K in K_∞ then the isomorphisms of* (i) *induce a surjection and injection, respectively*

$$Ш(E,L)[p^\infty] \twoheadrightarrow (S \otimes \Lambda_L)_{\mathrm{tors}},$$
$$(E(L) \otimes \mathbf{Z}_p) \hookrightarrow \mathrm{Tor}^1_\Lambda(S, \Lambda_L),$$

which are isomorphisms if $Ш(E,L)[p^\infty]$ is finite.

(iii) *If L is a finite extension of K in K_∞ then the pairings*

$$Ш(E,L)[p^\infty] \otimes Ш(E,L)[p^\infty] \longrightarrow \mathbf{Q}_p/\mathbf{Z}_p,$$
$$(E(L) \otimes \mathbf{Z}_p) \otimes (E(L) \otimes \mathbf{Z}_p) \longrightarrow \mathrm{Gal}(K_\infty/L),$$

obtained by combining the derived pairings (4.6′) *and* (4.5′) *with the maps of* (ii), *coincide* (*up to sign*) *with the classical Cassels and p-adic height pairing, respectively*.

Proof. We have $S = H^2(\Phi^\bullet) \cong H^2(C^\bullet_{\mathrm{Nek}}) \cong \mathcal{S}_p(E, K_\infty)$. This gives the first isomorphism of (i), the second follows by Lemma 2.2, and the third by Proposition 2.6 and (4.3).

The first map of (ii) comes from (i) and Lemma 2.3(i), and the second comes from (i) and the inclusion $(E(L)/E(L)_{\mathrm{tors}}) \otimes \mathbf{Z}_p \subset \mathcal{M}_p(E,L)$.

For assertion (iii), we need to check two things. The first is that our derived pairings (4.5) and (4.6), defined directly from the basic skew-Hermitian module Φ, coincide (up to sign) with the corresponding pairings made by Nekovář via the skew-Hermitian degree three perfect duality enjoyed by the basic complex $\Phi^\bullet$ obtained from Φ. The second is to relate these derived pairings to the corresponding (various) classical pairings.

For every intermediate field extension L/K in K_∞/K the Iwasawa algebra Λ_L is a quotient of a (complete) regular noetherian local ring by an ideal generated by a regular sequence, and so is a Gorenstein ring. For each of the intermediate fields L we identify the dualizing complex $\omega^\bullet_{\Lambda_L}$ of the ring Λ_L with the complex concentrated in degree zero, and given in degree zero by the free Λ_L-module of rank one, Λ_L itself.

Suppose $X^\bullet$ and $Y^\bullet$ are complexes of Λ-modules with cohomology of finite type equipped with a morphism of complexes

$$\eta : X^\bullet \otimes_\Lambda Y^\bullet \longrightarrow \omega^\bullet_\Lambda[-3].$$

Consider the following two pairings of cohomology of $X^\bullet$ and $Y^\bullet$. First, for all intermediate fields L we have ([N] 2.10.14) the morphism defined via cup-product

$$H^2(X^\bullet \otimes_\Lambda \Lambda_L)_{\mathrm{tors}} \otimes_{\Lambda_L} H^2(Y^\bullet \otimes_\Lambda \Lambda_L)_{\mathrm{tors}} \to H^0(\omega^\bullet_{\Lambda_L}) \otimes_{\Lambda_L} \mathcal{K}_L/\Lambda_L = \mathcal{K}_L/\Lambda_L, \tag{7.6}$$

where $\mathcal{K}_L$ is the field of fractions of Λ_L.

Second, we have the "derived (1, 1) cup-product"

$$H^1(X^\bullet \otimes_\Lambda \Lambda_L) \otimes_{\Lambda_L} H^1(Y^\bullet \otimes_\Lambda \Lambda_L) \longrightarrow H^0(\omega^\bullet_{\Lambda_L}) \otimes_{\Lambda_L} I_L/I_L^2. \tag{7.7}$$

This pairing can be defined in the following elementary way. For cohomology classes $(a, b) \in H^1(X^\bullet \otimes_\Lambda \Lambda_L) \times H^1(Y^\bullet \otimes_\Lambda \Lambda_L)$, choose 1-cochains $(x, y) \in X^1 \times Y^1$ such that the projection $(\tilde{x}, \tilde{y}) \in (X^1 \otimes_\Lambda \Lambda_L) \times (Y^1 \otimes_\Lambda \Lambda_L)$ is a pair of 1-cocycles representing the pair of cohomology classes (a, b). Note that $\partial x \in I_L X^2$ and $\partial y \in I_L Y^2$. So $\eta(x, \partial y) = -\eta(\partial x, y) \in \Lambda$ projects to zero in Λ_L, and hence lies in I_L. Let $\pi_L : I_L \to I_L^2$ be the natural projection, and put

$$\langle a, b \rangle := \pi_L(\eta(x, \partial y)) = -\pi_L(\eta(\partial x, y)) \in I_L/I_L^2. \tag{7.8}$$

To show that this is well-defined, first note that if $e \in I_L X^1$ then $\eta(e, \partial y) \in I_L^2$ (and, if $e \in I_L Y^1$ then $\eta(\partial x, e) \in I_L^2$) which tells us that $\pi_L(\eta(x, \partial y)) = -\pi_L(\eta(\partial x, y))$ depends only on $(\tilde{x}, \tilde{y})$. Next, if $\tilde{x} = \partial \tilde{v}$ for $\tilde{v} \in X^0 \otimes_\Lambda \Lambda_L$ lifting $\tilde{v}$ to $v \in X^0$ and taking $x = \partial v$ to be our lifting of $\tilde{x}$ gives us that $\pi_L(\eta(x, \partial y)) = \pi_L(\eta(\partial v, \partial y))$ vanishes; this, and the symmetrical argument when $y = \partial w$, gives us that the pairing (7.8) is well-defined.

The basic complex $\Phi^\bullet$ associated to Φ has a skew-Hermitian pairing

$$\Phi^\bullet \otimes_\Lambda (\Phi^\bullet)^\iota \to \omega^\bullet_\Lambda[-3], \tag{7.9}$$

so for each intermediate field L we have the induced pairing

$$(\Phi^\bullet \otimes_\Lambda \Lambda_L) \otimes_\Lambda (\Phi^\bullet \otimes_\Lambda \Lambda_L)^\iota \longrightarrow \omega^\bullet_{\Lambda_L}[-3].$$

In the notation of §4 we have $S(L) = H^2(\Phi^\bullet \otimes_\Lambda \Lambda_L)$ and $M(L) = H^1(\Phi^\bullet \otimes_\Lambda \Lambda_L)$, so the cup-product pairing (7.6) obtained from (7.9) may be written

$$S(L)_{\mathrm{tors}} \otimes_{\Lambda_L} S(L)^\iota_{\mathrm{tors}} \longrightarrow \mathcal{K}_L/\Lambda_L, \tag{7.10}$$

and the derived (1, 1) pairing may be written

$$M(L) \otimes_{\Lambda_L} M(L)^\iota \longrightarrow I_L/I_L^2. \tag{7.11}$$

It is straightforward to compute that the pairing (4.6) is, up to sign, equal to the pairing (7.10) and the pairing (4.5) is, up to sign, equal to the pairing (7.11).

Now, using the equivalence in the derived category $\mathcal{D}$ between the perfect degree three skew-Hermitian self-dualities on $C^\bullet_{\mathrm{Nek}}$ and $\Phi^\bullet$, one can check that the pairing (7.10) is, up to sign, equal to the ("Cassels–Tate") pairing

$$\cup_{\bar{\pi},0,2,2} : H^2(C^\bullet_{\mathrm{Nek}} \otimes_\Lambda \Lambda_L)_{\mathrm{tors}} \otimes_{\Lambda_L} H^2(C^\bullet_{\mathrm{Nek}} \otimes_\Lambda \Lambda_L)^\iota_{\mathrm{tors}} \to H^0(\omega^\bullet_{\Lambda_L}) \otimes_{\Lambda_L} \mathcal{K}_L/\Lambda_L$$

of ([N] §10.3.3.3), and that (7.11) is, up to sign, equal to the ("height") pairing

$$\tilde{h}_{\pi,L/K,1,1} : H^1(C^{\bullet}_{\mathrm{Nek}} \otimes_{\Lambda} \Lambda_L) \otimes_{\Lambda_L} H^1(C^{\bullet}_{\mathrm{Nek}} \otimes_{\Lambda} \Lambda_L)^{\iota} \to H^0(\omega^{\bullet}_{\Lambda_L}) \otimes_{\Lambda_L} I_L/I_L^2$$

of [N] (11.1.7.5) (see also [N] §§11.1.4,11.1.7,11.1.8).

Finally, assertion (iii) follows from the discussion in §10 and §11 of [N] that makes the connection between the Cassels–Tate and height pairings defined there and the classical pairings of the same name. □

Remark 7.8. There are indeed many different approaches to defining what may be called the *classical p-adic height pairing* and the somewhat ample discussion in [N] is a welcome addition to the literature comparing some of these approaches. The next step that remains to be done is a systematic expository account of all this.

Remark 7.9. Note that because Φ is a basic skew-Hermitian module, we have

$$\mathrm{rank}_{\Lambda}(\Phi) = \dim_{\mathbf{F}_p}(\mathrm{Sel}_p(E,K)[p]) = \mathrm{rank}_{\mathbf{Z}}(E(K)) + \dim_{\mathbf{F}_p} Ш(E,K)[p].$$

If we choose a basis of the organizing module Φ then the pairing h is equivalent to a skew-Hermitian matrix H with entries in Λ. We then have that the characteristic ideal $\mathrm{char}(\mathcal{S}_p(E,K_\infty)) = \det(H)\Lambda$, and the matrix H contains complete information about the Selmer modules $\mathcal{S}_p(E,L)$ and the Cassels and p-adic height pairings on $Ш(E,L)[p^\infty]$ and $E(L)\otimes\mathbf{Z}_p$, for every finite extension L of K in K_∞.

Remark 7.10. Thanks to the Perfect Control assumption (see Lemma 2.2), if $\mathcal{S}_p(E,L)$ is a torsion Λ_L-module for some $\mathbf{Z}_p^d$-extension L of K with $d \geqslant 0$, then $\mathcal{S}_p(E,K_\infty)$ is a torsion Λ-module. In particular

- if $\mathrm{Sel}_p(E,K)$ is finite (i.e., if $E(K)$ is finite, since we are assuming that $Ш(E,K)[p^\infty]$ is finite) then $\mathcal{S}_p(E,K_\infty)$ is a torsion Λ-module,
- if E is defined over $\mathbf{Q}$ and $K/\mathbf{Q}$ is abelian, then by work of Kato [Ka] $\mathcal{S}_p(E,K\mathbf{Q}_\infty)$ is a torsion $\Lambda_{K\mathbf{Q}_\infty}$-module, where $K\mathbf{Q}_\infty$ denotes the cyclotomic $\mathbf{Z}_p$-extension of K, so $\mathcal{S}_p(E,K_\infty)$ is a torsion Λ-module.

Remark 7.11. Corollary A.3 shows that the Perfect Control assumption follows from hypotheses (7.3), (7.4) along with the additional assumption that $E(k_v)[p] = 0$ for every prime v of K above p, where k_v is the residue field at v.

The following proposition, which combines some of the observations above, allows us to verify hypotheses (7.1-5) in many interesting cases.

Proposition 7.12. *Suppose that E is defined over $\mathbf{Q}$ and K is a finite abelian extension of $\mathbf{Q}$. Suppose p is a rational prime such that*

(i) *for every prime v of K above p, E has good reduction at v and $\#E(k_v) \not\equiv 0$ or $1 \pmod p$ where k_v is the residue field at v,*

(ii) *for every prime v of K where E has bad reduction, p does not divide the Tamagawa number $[E(K_v) : E_0(K_v)]$, and*
(iii) *p is unramified in $K/\mathbf{Q}$.*

Then hypotheses (7.1-5) *hold.*

Proof. If (i) holds then p cannot be 2, and further E has good ordinary reduction at each v dividing p. This is (7.1), and (ii) is (7.4).

Fix a prime v of K above p. It follows from (iii) that K_v^{unr} has no pth roots of unity, so (7.3) follows from Lemma A.6. Now the Perfect Control assumption (7.5) follows from (i) and Corollary A.3, as in Remark 7.11, and then (7.2) follows as in Remark 7.10. □

For example, we have the following corollary mentioned in the introduction.

Corollary 7.13. *Suppose that E is defined over $\mathbf{Q}$, with conductor N_E and minimal discriminant Δ_E. Suppose further that K is a finite abelian extension of $\mathbf{Q}$ with discriminant D_K prime to N_E, and p is a rational prime such that*
(i) $p \nmid 3N_E D_K \prod_{\ell | N_E} \mathrm{ord}_\ell(\Delta_E)$,
(ii) $a_p \not\equiv 0$ *and* $a_p^{[K:\mathbf{Q}]} \not\equiv 1 (\mathrm{mod}\, p)$, *where as usual* $a_p = 1 + p - \#E(\mathbf{Z}/p\mathbf{Z})$.
Then there is a basic skew-Hermitian module Φ, unique up to (noncanonical) isomorphism, that organizes the arithmetic of E over K_∞. We can recover from Φ as in Theorem 7.7 *the Selmer modules, p-adic height pairings, and Cassels pairings over every finite extension of K in K_∞.*

Proof. We will verify that the hypotheses of Proposition 7.12 hold. Proposition 7.12(iii) holds since $p \nmid D_K$.

Suppose first that v is a prime of K above p. Since $p \nmid N_E$, E has good reduction at v. Further, if α_p and β_p are the roots of the Frobenius polynomial $x^2 - a_p x + p$, and $f = [k_v : \mathbf{F}_p]$, then

$$\#E(k_v) = 1 + p^f - \alpha^f - \beta^f \equiv 1 - (\alpha + \beta)^f = 1 - a_p^f (\mathrm{mod}\, p).$$

Since $f \mid [K : \mathbf{Q}]$ and $a_p^{[K:\mathbf{Q}]} \not\equiv 0, 1 (\mathrm{mod}\, p)$, Proposition 7.12(i) holds.

Next suppose v is a prime of K where E has bad reduction, and let ℓ be the rational prime below v. If E has either additive or nonsplit multiplicative reduction at v then $[E(K_v) : E_0(K_v)]$ divides 12 (see [T]), but condition (i) rules out $p = 3$ and condition (ii) rules out $p = 2$, so $p \nmid [E(K_v) : E_0(K_v)]$. On the other hand, if E has multiplicative reduction at v then $[E(K_v) : E_0(K_v)]$ is the order at v of the discriminant of E/K ([T] step 2). Since by assumption ℓ is unramified in $K/\mathbf{Q}$, we have $[E(K_v) : E_0(K_v)] = \mathrm{ord}_\ell(\Delta_E)$ which is prime to p. Thus, Proposition 7.12(ii) holds.

Now by Proposition 7.12, hypotheses (7.1-5) hold. Thus, the existence and uniqueness of Φ follow from Theorem 7.5, and that fact that we can recover the arithmetic of E over finite extensions of K in K_∞ follows from Theorem 7.7(iii). □

8. A generic example

In the next three sections we consider several families of examples where we can give some information about the organizing module. We first consider the “generic” situation where $\text{Ш}(E/K)[p] = 0$, so that $\mathrm{Sel}_p(E/K) = E(K) \otimes \mathbf{Q}_p/\mathbf{Z}_p$.

Suppose that E is an elliptic curve defined over K, and let $r = \mathrm{rank}(E(K))$. Let p be a rational prime for which hypotheses (7.1-5) are satisfied (see for example Proposition 7.12), and suppose in addition that $\text{Ш}(E/K)[p] = 0$. (Conjecturally this last condition is satisfied for all but finitely many p.) Then we have $\mathrm{Sel}_p(E, K) \cong (\mathbf{Q}_p/\mathbf{Z}_p)^r$, and by Theorem 7.5 there is a basic skew-Hermitian Λ-module Φ, free of rank r, that organizes the arithmetic of E/K_∞.

If $r = 0$ then Φ is trivial, the Selmer modules over all intermediate fields are trivial, and there is nothing more to study. Suppose, then, that $r > 0$. We want to describe the $r \times r$ skew-Hermitian matrix H for the pairing h corresponding to a suitable basis of Φ.

Let I denote the augmentation ideal $I_K \subset \Lambda$, and identify $\Lambda_K = \Lambda/I = \mathbf{Z}_p$. The skew-Hermitian pairing h induces an exact sequence

$$\Phi \otimes_\Lambda \mathbf{Z}_p \xrightarrow{h \otimes \mathbf{Z}_p} \Phi^* \otimes_\Lambda \mathbf{Z}_p \to \mathrm{Hom}(E(K), \mathbf{Z}_p) \to 0 \tag{8.1}$$

in which the first three $\mathbf{Z}_p$-modules are all free of rank r. It follows that the map $\Phi^* \otimes_\Lambda \mathbf{Z}_p \to \mathrm{Hom}(E(K), \mathbf{Z}_p)$ is an isomorphism, and using the identification

$$\Phi^* \otimes_\Lambda \mathbf{Z}_p \cong \mathrm{Hom}(\Phi/I\Phi, \mathbf{Z}_p)$$

we obtain an isomorphism

$$\Phi/I\Phi \cong (E(K)/E(K)_{\mathrm{tors}}) \otimes \mathbf{Z}_p.$$

Thus, we can take the organizing module Φ to be $(E(K)/E(K)_{\mathrm{tors}}) \otimes_{\mathbf{Z}} \Lambda$.

It also follows from (8.1) that the matrix H has entries in I. In addition, the image of H in $\mathrm{M}_r(I/I^2)$ is the p-adic height pairing matrix for a basis of $(E(K)/E(K)_{\mathrm{tors}}) \otimes \mathbf{Z}_p$ corresponding to the chosen basis of Φ. Hence we can view H as a lift of the $(I/I^2$-valued) p-adic height pairing on $(E(K)/E(K)_{\mathrm{tors}}) \otimes \mathbf{Z}_p$ to an I-valued skew-Hermitian pairing on $\Phi \otimes_\Lambda \Phi^\iota$, with $\Phi = (E(K)/E(K)_{\mathrm{tors}}) \otimes \Lambda$.

9. Examples over Q

For this section we take $K = \mathbf{Q}$. Fix a generator γ of $\Gamma = \mathrm{Gal}(\mathbf{Q}_\infty/\mathbf{Q})$ and let $\theta := \gamma - \gamma^{-1}$. Then we have $\Lambda = \mathbf{Z}_p[[\gamma - 1]] = \mathbf{Z}_p[[\theta]]$, and the augmentation ideal $I = \theta\Lambda$. If we write $\Lambda^\pm$ for the ± 1 eigenspaces of ι on Λ, then $\Lambda_+ = \mathbf{Z}_p[[\theta^2]]$ and $\Lambda_- = \theta\Lambda_+$.

Fix an elliptic curve E defined over $\mathbf{Q}$.

Definition 9.1. We say that a prime p is *admissible* if it satisfies the following two conditions:

- E has good reduction at p, p does not divide the order of the torsion subgroup of $E(\mathbf{Q})$, and p does not divide any of the Tamagawa numbers of E over $\mathbf{Q}$,
- E has ordinary and nonanomalous reduction at p (i.e., $\#E(\mathbf{F}_p) \not\equiv 1 (\mathrm{mod}\, p)$ and $\#E(\mathbf{F}_p) \not\equiv 0 (\mathrm{mod}\, p)$).

Note that the first condition rules out only a finite set of primes, and the second only rules out a set of Dirichlet density $1/2$ or 0 depending upon whether E has CM (over $\bar{\mathbf{Q}}$) or not.

9.1. *The case* Ш$(E, \mathbf{Q})[p] = 0$

Suppose now that p is admissible, and suppose further that Ш$(E, \mathbf{Q})[p] = 0$. (If the Shafarevich–Tate group of E is finite, then this is true for all but finitely many admissible primes.) Then we are in the situation of §8, and there is a skew-Hermitian pairing on $\Phi := (E(\mathbf{Q})/E(\mathbf{Q})_{\mathrm{tors}}) \otimes \Lambda$ so that Φ organizes the arithmetic of $E/\mathbf{Q}_\infty$.

Let $r = \mathrm{rank}(E(\mathbf{Q}))$. We want to describe the $r \times r$ skew-Hermitian matrix H for the pairing h corresponding to a suitable basis of Φ. As discussed in §8, H has entries in $I = \theta\Lambda$ and H is a lift to $\mathrm{M}_r(I)$ of the height pairing matrix in $\mathrm{M}_r(I/I^2)$ for $E(\mathbf{Q}) \otimes \mathbf{Z}_p$. Let

$$H' := \theta^{-1} H \in \mathrm{M}_r(\Lambda),$$

so H' is a Hermitian matrix in $\mathrm{M}_r(\Lambda)$ and its reduction in $\mathrm{M}_r(\Lambda/I) = \mathrm{M}_r(\mathbf{Z}_p)$ is a symmetric matrix describing the height pairing (divided by θ)

$$\eta : (E(\mathbf{Q}) \otimes \mathbf{Z}_p) \otimes (E(\mathbf{Q}) \otimes \mathbf{Z}_p) \longrightarrow I/I^2 \xrightarrow{\theta^{-1}} \Lambda/I \xrightarrow{\sim} \mathbf{Z}_p.$$

Definition 9.1.1. Choose a $\mathbf{Z}_p$-basis $\mathbf{b} := \{e_1, e_2, \ldots, e_r\}$ of $(E(\mathbf{Q})/E(\mathbf{Q})_{\mathrm{tors}}) \otimes \mathbf{Z}_p$ and compute the discriminant of η, i.e.,

$$\mathrm{disc}(\eta, \mathbf{b}) = \det(\eta(e_i, e_j)) \in \mathbf{Z}_p.$$

This discriminant is well-defined, independent of the chosen basis $\mathbf{b}$ up to multiplication by the square of an element in $\mathbf{Z}_p^\times$. In particular, if $\mathrm{disc}(\eta, \mathbf{b})$ does not vanish (i.e., if the p-adic height pairing is nondegenerate), then we can define two numerical invariants

- a nonnegative integer $\rho := \mathrm{ord}_p(\mathrm{disc}(\eta, \mathbf{b}))$, the *irregularity* of η,
- the Legendre symbol $(\frac{p^{-\rho}\mathrm{disc}(\eta, \mathbf{b})}{p}) \in \{\pm 1\}$, the *sign* of η.

If the irregularity of h is zero, we will say that p is *regular* for E. If it ever happens that $\mathrm{disc}(\eta, \mathbf{b}) = 0$, we will just say then that the irregularity is ∞ (and not try to ascribe a “sign” to η).

Note that the irregularity of η depends only on Φ and its skew-Hermitian pairing. The same is true of $\mathrm{sign}(\eta)$ if r is even, but if r is odd then $\mathrm{sign}(\eta)$ also depends on the choice of γ.

Proposition 9.1.2. *If p is regular for E, then Φ has a basis for which the matrix H' is diagonal with all but the last entry equal to* 1, *and the last entry can be taken to be any $u \in \mathbf{Z}_p^\times$ with $(\frac{u}{p}) = \mathrm{sign}(\eta)$. In particular if* $\mathrm{sign}(\eta) = +1$ *then H' can be take to be the identity matrix.*

Proof. Let h' denote the Hermitian pairing $\theta^{-1}h$ on Φ. Since p is regular, h' is a perfect pairing.

If $\mathrm{rank}_\Lambda \Phi > 1$, then h' represents a square in $\Lambda^\times$, i.e., we can choose $x \in \Phi$ such that $h'(x,x) = \beta^2$ with $\beta \in \Lambda^\times$. Since h' is Hermitian, we have $(\beta^2)^\iota = \beta^2$, so $\beta^\iota = \pm\beta$. But $\beta \notin \Lambda^-$ since β is a unit, so $\beta \in \Lambda^+$. Replacing x by $x_1 = \beta^{-1}x$ we have $h'(x_1, x_1) = 1$.

Let $M_1 = \Lambda x_1$ and let $N_1 \subset \Phi$ be the orthogonal complement of M_1. Then $M_1 \oplus N_1 = \Phi$. Continuing by induction we get a basis $\{x_1, \ldots, x_{r-1}, x_r\}$ of Φ such that $h'(x_i, x_j) = 0$ if $i \neq j$, and $h'(x_i, x_i) = 1$ if $i < r$. We have $h'(x_r, x_r) \in \Lambda^+$, and we may change it by any square in Λ^+. In this way we obtain the desired basis of Φ. □

It would be interesting to gather numerical data for particular elliptic curves E to learn something about the distribution, among admissible primes, of sign and irregularity. Some examples and conjectures concerning irregularity are given by Wuthrich in [W].

Example 9.1.3. Let E be the elliptic curve $y^2 + xy + y = x^3 + 2$, 1058C1 in Cremona's tables [Cr]. For this curve we have $E(\mathbf{Q}) \cong \mathbf{Z}^2$, the Tamagawa numbers at the bad primes 2 and 23 are 2 and 1, respectively, and the Birch and Swinnerton–Dyer conjecture predicts that $Ш(E, \mathbf{Q}) = 0$.

Using the basis $\mathbf{b} = \{(-1, 1), (0, 1)\}$ for $E(\mathbf{Q})$, William Stein (using methods described in a forthcoming paper by Stein, Tate, and the first author [MST]) computed $\mathrm{disc}(\eta, \mathbf{b})$ for the 337 admissible primes $p < 2400$. The computation shows that all of these primes are regular, and 175 have sign $= +1$ and 162 have sign $= -1$.

For example, if $p = 5$ and we take γ to be the generator of Γ satisfying $\varepsilon(\gamma) = 6$, where $\varepsilon : \Gamma \xrightarrow{\sim} 1 + 5\mathbf{Z}_5$ is the cyclotomic character, then the height pairing matrix for the basis $\mathbf{b}$ above is

$$H' \equiv \begin{pmatrix} 33 & 105 \\ 105 & 83 \end{pmatrix} \mathrm{mod}(5^3 + I).$$

Thus the sign is $+1$, so by Proposition 9.1.2 we can choose a new basis with

$$H = \begin{pmatrix} \theta & 0 \\ 0 & \theta \end{pmatrix} \tag{9.1}$$

9.2. The case $\operatorname{rank}(E(\mathbf{Q}))=0$

At the opposite extreme from §9.1, we consider here a case where $E(\mathbf{Q})$ has rank zero so that the Selmer group is the Shafarevich–Tate group. We will make some additional assumptions so that we can analyze this example in detail.

Suppose that $\operatorname{rank}(E(\mathbf{Q}))=0$, p is admissible, and $Ш(E,\mathbf{Q})[p^\infty]\cong(\mathbf{Z}/p\mathbf{Z})^2$. Suppose further that $\mathcal{S}_p(E,\mathbf{Q}_\infty)$ has $\mathbf{Z}_p$-rank 2. In this case we have an organizing module Φ with $\operatorname{rank}_\Lambda(\Phi)=2$.

Proposition 9.2.1. *There is a basis of* Φ *such that the corresponding skew-Hermitian matrix has the form* $\begin{pmatrix}\theta & -p\\ p & \alpha\theta\end{pmatrix}$ *with* $\alpha\in\mathbf{Z}_p[[\theta^2]]^\times$.

Sketch of proof. Fix a basis of Φ and let $f\in\Lambda_+=\mathbf{Z}_p[[\theta^2]]$ be the determinant of the corresponding skew-Hermitian matrix. Write $f=a_0+a_2\theta^2+\cdots$ with $a_i\in\mathbf{Z}_p$.

We have $f\Lambda=\operatorname{char}(\mathcal{S}_p(E,\mathbf{Q}_\infty))$. Thus

$$a_0\in p^2\mathbf{Z}_p^\times,\quad a_2\in\mathbf{Z}_p^\times \tag{9.2}$$

because $\mathcal{S}_p(E,\mathbf{Q})$ has order p^2 and $\operatorname{rank}_{\mathbf{Z}_p}(\mathcal{S}_p(E,\mathbf{Q}_\infty))=2$, respectively.

If $x,y\in\Phi$ let $\langle x,y\rangle$ denote $h(x\otimes y)$.

We first claim that there is an $x\in\Phi$ such that $\langle x,x\rangle\notin\theta\mathfrak{m}$, where $\mathfrak{m}$ is the maximal ideal of Λ. Suppose on the contrary that $\langle x,x\rangle\in\theta\mathfrak{m}$ for every x. Then if $\{u,v\}$ is the chosen basis of Φ, we have modulo $\theta\mathfrak{m}$

$$\begin{aligned} f &= \langle u,u\rangle\langle v,v\rangle-\langle u,v\rangle\langle v,u\rangle\equiv-\langle u,v\rangle\langle v,u\rangle\\ &= \tfrac14(\langle u,v\rangle-\langle v,u\rangle)^2-\tfrac14(\langle u,v\rangle+\langle v,u\rangle)^2\\ &= \tfrac14(\langle u,v\rangle-\langle v,u\rangle)^2-\tfrac14(\langle u+v,u+v\rangle-\langle u,u\rangle-\langle v,v\rangle)^2\\ &\equiv \tfrac14(\langle u,v\rangle+\langle u,v\rangle^\iota)^2. \end{aligned}$$

Since $(\langle u,v\rangle+\langle u,v\rangle^\iota)/2\in\Lambda_+=\mathbf{Z}_p[[\theta^2]]$, this is incompatible with (9.2). This proves the claim.

Fix a basis $\{x,y\}$ of Φ with $\langle x,x\rangle\notin\theta\mathfrak{m}$. Since $\langle x,x\rangle\in\Lambda_-=\theta\Lambda_+$, we have $\langle x,x\rangle\in\theta\Lambda_+^\times$. By adding a multiple of x to y we may assume that $\langle x,y\rangle\in\mathbf{Z}_p$, and by (9.2) we must have $\langle x,y\rangle\in p\mathbf{Z}_p^\times$ and $\langle y,y\rangle\in\theta\Lambda_+^\times$. Now scaling y by a unit we may assume further that $\langle x,y\rangle=p$.

Finally, by considering $ax+by$ with $a,b\in\mathbf{Z}_p$, we can see now that there is a $z\in\Phi$ such that $\langle z,z\rangle=\theta\beta$ with β a square in $\Lambda_+^\times$. Scaling z by $\sqrt{\beta}$ we find that $\langle z,z\rangle=1$. Repeating the argument of the previous paragraph starting with $x=z$ proves the proposition. □

Example 9.2.2. Let E be the elliptic curve $y^2+xy=x^3-x^2-332311x-73733731$, 1058D1 in Cremona's tables [Cr]. For this curve we have $E(\mathbf{Q})=0$, $Ш(E,\mathbf{Q})\cong(\mathbf{Z}/5\mathbf{Z})^2$, and all Tamagawa numbers are 1. If p is an admissible prime different from 5, then $\Phi = 0$ is an organizing module.

Now take $p = 5$. Since $\#E(\mathbf{Z}/5\mathbf{Z}) = 4$, Proposition 7.12 shows that hypotheses (7.1-5) are satisfied. In particular the Perfect Control assumption holds, so $\mathcal{S}_5(E, \mathbf{Q}_\infty)$ is not a cyclic Λ-module. By Greenberg's Theorem C.2, $\mathcal{S}_5(E, \mathbf{Q}_\infty)$ has no finite Λ-submodules (this can also be seen directly from the existence of an organizing module for $E/\mathbf{Q}_\infty$), so the sum of the λ- and μ-invariants $\lambda_{\mathrm{alg}} + \mu_{\mathrm{alg}}$ of $\mathcal{S}_5(E, \mathbf{Q}_\infty)$ is at least 2.

Let $\mathcal{L}_5(E) \in \Lambda$ denote the 5-adic L-function attached to E. Let $\mathbf{1}$ denote the trivial character of Γ, $\zeta \in \boldsymbol{\mu}_5$ a primitive 5th root of unity, and χ the character of Γ that sends γ to ζ. The definition of $\mathcal{L}_5(E)$ and a computation of $L(E, 1)$ and $L(E, \chi, 1)$ show that $\mathbf{1}(\mathcal{L}_5(E)) \in 5^2\mathbf{Z}_5^\times$ and

$$\chi(\mathcal{L}_5(E)) = (-3\zeta^3 - 25\zeta^2 - 3\zeta)(\zeta - \zeta^{-1})^2 \frac{\mathbf{1}(\mathcal{L}_5(E))}{5^2}. \tag{9.3}$$

Since $-3\zeta^3 - 25\zeta^2 - 3\zeta \equiv -1 (\mathrm{mod}\,(\zeta - 1))$ is a unit in $\mathbf{Z}_5[\zeta]$, we see that the λ- and μ-invariants of $\mathcal{L}_5(E)$ are $\lambda_{\mathrm{an}} = 2$ and $\mu_{\mathrm{an}} = 0$.

One can check that the representation $G_{\mathbf{Q}} \to \mathrm{Aut}(E[5]) \cong \mathrm{GL}_2(\mathbf{F}_5)$ is surjective, so a theorem of Kato [Ka] shows that $\mathrm{char}(\mathcal{S}_5(E, \mathbf{Q}_\infty))$ divides $\mathcal{L}_5(E)$. In particular $\lambda_{\mathrm{alg}} \leqslant 2$ and $\mu_{\mathrm{alg}} = 0$, so $\lambda_{\mathrm{alg}} = 2$, $\mathcal{S}_5(E, \mathbf{Q}_\infty)$ is free of rank 2 over $\mathbf{Z}_p$, and the assumptions at the beginning of §9.2 are satisfied. Further, we conclude that the Main Conjecture is true for E, i.e.,

$$\mathcal{L}_5(E)\Lambda = \mathrm{char}(\mathcal{S}_5(E, \mathbf{Q}_\infty)). \tag{9.4}$$

Let H be the skew-symmetric matrix of Proposition 9.2.1. We will show that α is a square in Λ_+.

By (9.4) there is a $\beta \in \Lambda^\times$ such that

$$\mathcal{L}_5(E) = \det(H)\beta = \beta(\alpha\theta^2 + 5^2).$$

It follows that $\mathbf{1}(\mathcal{L}_5(E)) = \mathbf{1}(\beta)5^2$ and that

$$\chi(\mathcal{L}_5(E)) = \chi(\beta)(\chi(\alpha)\chi(\theta)^2 + 5^2) \equiv \mathbf{1}(\beta)\chi(\alpha)(\zeta - \zeta^{-1})^2 (\mathrm{mod}\,(\zeta - 1)^3)$$

in the ring $\mathbf{Z}_5[\zeta]$ (with maximal ideal generated by $\zeta - 1$). Comparing this with (9.3) we conclude that

$$\chi(\alpha) \equiv -3\zeta^3 - 25\zeta^2 - 3\zeta \equiv -1 (\mathrm{mod}\,(\zeta - 1))$$

so α is the square of a unit in $\Lambda_+ = \mathbf{Z}_5[[\theta^2]]$.

Fix $a \in \Lambda_+$ with $a^2 = \alpha$. Replacing the basis $\{x, y\}$ of Proposition 9.2.1 by $\{x, a^{-1}y\}$ gives a new matrix

$$H' = \begin{pmatrix} \theta & -5a^{-1} \\ 5a^{-1} & \theta \end{pmatrix}. \tag{9.5}$$

With more work one can modify the basis to obtain $H'' = \left(\begin{smallmatrix} \theta & -5b \\ 5b & \theta \end{smallmatrix}\right)$ with $b \in \mathbf{Z}_5^\times$.

9.3. *A congruence*

The curves of Examples 9.1.3 and 9.2.2 have a congruence modulo 5. More precisely, their corresponding modular forms are congruent modulo 5 (and have the same conductor). In particular, the Shafarevich–Tate group $(\mathbf{Z}/5\mathbf{Z})^2$ in Example 9.2.2 is "visible" in the sense of [CM] thanks to this congruence and the Mordell–Weil group $\mathbf{Z}^2$ of Example 9.2.2.

Examples 9.1.3 and 9.2.2, and in particular (9.1) and (9.5), show that this congruence is matched by a congruence modulo 5 between the two organizing modules.

10. Examples over an imaginary quadratic field

Suppose now that E is defined over $\mathbf{Q}$, and that K is an imaginary quadratic field in which all primes dividing the conductor of E split. Suppose p is a prime where E has good ordinary reduction, not dividing any of the Tamagawa numbers c_ℓ for primes ℓ of bad reduction. Suppose further that p is unramified in $K/\mathbf{Q}$, $a_p \not\equiv 1 (\mathrm{mod}\, p)$ where a_p is the pth Fourier coefficient of the modular form corresponding to E, and if p is inert in K then $a_p \not\equiv -1 (\mathrm{mod}\, p)$ as well. Then by Proposition 7.12, hypotheses (7.1-5) all hold, so we have an organizing module Φ by Theorem 7.5.

Let K^{anti} denote the anti-cyclotomic $\mathbf{Z}_p$-extension of K, and $\Lambda_{\mathrm{anti}} := \Lambda_{K^{\mathrm{anti}}}$. Fix a topological generator γ of $\mathrm{Gal}(K_\infty/K^{\mathrm{anti}}) \cong \mathrm{Gal}(\mathbf{Q}_\infty/\mathbf{Q})$ and let $\theta := \gamma - \gamma^{-1}$, a generator of the augmentation ideal $I_{K^{\mathrm{anti}}} \subset \Lambda$.

Let $X_\infty = \mathcal{S}_p(E, K_\infty)$ and $X^{\mathrm{anti}} := X_\infty \otimes_\Lambda \Lambda_{\mathrm{anti}} = \mathcal{S}_p(E, K^{\mathrm{anti}})$. Writing $\mathcal{U} := M(K^{\mathrm{anti}})$ as defined in §4, the exact sequence (4.2) becomes

$$0 \longrightarrow \mathcal{U} \longrightarrow \Phi \otimes_\Lambda \Lambda_{\mathrm{anti}} \xrightarrow{h \otimes \Lambda_{\mathrm{anti}}} \Phi^* \otimes_\Lambda \Lambda_{\mathrm{anti}} \longrightarrow X^{\mathrm{anti}} \longrightarrow 0. \tag{10.1}$$

By Proposition 2.6, (4.3), and the Perfect Control assumption, $\mathcal{U}$ is canonically isomorphic to the module of anti cyclotomic universal norms

$$\mathcal{M}_p(E, K^{\mathrm{anti}}) = \varinjlim_{L} (E(L) \otimes \mathbf{Z}_p),$$

inverse limit over finite extensions L of K in K^{anti}. Let $r := \mathrm{rank}_\Lambda \Phi$.

It follows from the work of Cornut [Co] and Vatsal [V] that, under the hypotheses above on K, we have $\mathrm{rank}_{\Lambda_{\mathrm{anti}}} X^{\mathrm{anti}} = 1$. Hence we conclude from (10.1) that

$$\begin{aligned}
&\mathcal{U} \text{ is free of rank one over } \Lambda_{\mathrm{anti}},\\
&(\Phi \otimes \Lambda_{\mathrm{anti}})/\mathcal{U} \text{ is torsion-free of rank } r-1 \text{ over } \Lambda_{\mathrm{anti}},\\
&(\Phi \otimes \Lambda_{\mathrm{anti}})/\mathcal{U} \text{ is free } \iff X^{\mathrm{anti}} \text{ has no nonzero finite submodules.}
\end{aligned} \tag{10.2}$$

Suppose first that $(\Phi \otimes \Lambda_{\mathrm{anti}})/\mathcal{U}$ is free. Choose a Λ-basis $\{u_1, \ldots, u_r\}$ of Φ such that u_1 projects to a Λ_{anti}-generator of $\mathcal{U}$.

With this basis, the skew-Hermitian matrix H has the form

$$H = \left(\begin{array}{c|c} \theta a & \theta(\mathbf{w}^{\iota})^{\mathrm{tr}} \\ \hline \theta\mathbf{w} & B \end{array}\right), \tag{10.3}$$

where $a \in \Lambda$, $B \in \mathrm{M}_{r-1}(I_K)$, and $\mathbf{w} \in \Lambda^{r-1}$ is a column vector. (To see this, note that the left-hand column is divisible by θ because the image of u_1 in $\Phi \otimes \Lambda_{\mathrm{anti}}$ lies in $\mathcal{U} = \ker(h \otimes \Lambda_{\mathrm{anti}})$, and everything else follows from the fact that H is skew-Hermitian.)

Let $\mathrm{ht}_{\mathrm{anti}} : \mathcal{U} \otimes \mathcal{U}^{\iota} \to I_{K^{\mathrm{anti}}}/I^2_{K^{\mathrm{anti}}}$ denote the derived pairing (4.5). By definition of "organizing module", this is the same as the inverse limit of the p-adic height pairings over finite extensions of K in K^{anti}. We easily deduce the following:

$$\mathrm{char}(X_\infty) = \det(H)\Lambda \quad \text{and} \quad \det(H) \equiv \theta\, a \, \det(B) (\mathrm{mod}\, \theta^2), \tag{10.4}$$

$$\mathrm{char}(X^{\mathrm{anti}}_{\mathrm{tors}}) = \det(B)\Lambda_{\mathrm{anti}}, \tag{10.5}$$

$$\mathrm{ht}_{\mathrm{anti}}(\mathcal{U} \otimes \mathcal{U}^{\iota}) = a(I_{K^{\mathrm{anti}}}/I^2_{K^{\mathrm{anti}}}), \tag{10.6}$$

where the third assertion is immediate from the definition of the derived pairing (see Remark 4.1).

Note that the matrix H makes it easy to compute the Fitting ideals of X_∞. We see that

$$\begin{aligned}
\mathrm{Fitt}_0(X_\infty) &= \det(H)\Lambda = \mathrm{char}(X_\infty),\\
\mathrm{Fitt}_1(X_\infty)\Lambda_{\mathrm{anti}} &= \det(B)\Lambda_{\mathrm{anti}} = \mathrm{char}(X^{\mathrm{anti}}_{\mathrm{tors}}).
\end{aligned}$$

Remark 10.1. We will call the image in Λ_{anti} of the element a of (10.6) the *anticyclotomic regulator* of E/K^{anti}, and we will say that p is *regular* for E/K^{anti} if the anticyclotomic regulator is a unit (or equivalently if $\mathrm{ht}_{\mathrm{anti}}(\mathcal{U} \otimes \mathcal{U}^{\iota}) = I_{K^{\mathrm{anti}}}/I^2_{K^{\mathrm{anti}}}$). In Conjecture 6.1 of [MR2] (see also Conjecture 6 of [MR1]), we conjectured that every prime p (satisfying our hypotheses above) is regular for E/K^{anti}. This turns out to be false in general; see Example 10.10 for a counterexample.

One can still hope to predict some properties of the anti-cyclotomic regulator. For example, the nondegeneracy of the p-adic height pairing in the cyclotomic direction over all finite extensions of K in K^{anti} would imply that $\chi(a) \neq 0$ for all characters χ of finite order of $\mathrm{Gal}(K_{\mathrm{anti}}/K)$.

Theorem 10.2. *The characteristic ideal* $\mathrm{char}(X_\infty)$ *is contained in* $I_{K^{\mathrm{anti}}}$ *and*

$$\mathrm{char}(X_\infty) \equiv \mathrm{ht}_{\mathrm{anti}}(\mathcal{U} \otimes \mathcal{U}^\iota)\mathrm{char}(X_{\mathrm{tors}}^{\mathrm{anti}})(\mathrm{mod}\ I_{K^{\mathrm{anti}}}^2).$$

Proof. If $(\Phi \otimes \Lambda_{\mathrm{anti}})/\mathcal{U}$ is free, then $\mathrm{char}(X_\infty) \subset I_{K^{\mathrm{anti}}}$ by (10.4) and the congruence of the theorem is a consequence of (10.4), (10.5), and (10.6).

If $(\Phi \otimes \Lambda_{\mathrm{anti}})/\mathcal{U}$ is not free, then it injects into a free module with finite cokernel. With more care, that is sufficient to follow the argument above and deduce the theorem. □

The literature contains the following conjectures, and theorems concerning them.

Conjecture 10.3 (*Main conjecture*). $\mathrm{char}(X_\infty) = \mathcal{L}_p(E)$, *where* $\mathcal{L}_p(E) \in \Lambda$ *is the* 2-*variable p-adic L-function of Haran* [Hara], *Hida* [Hi], *and Perrin-Riou* [PR2].

Conjecture 10.4. $\mathrm{char} X_{\mathrm{tors}}^{\mathrm{anti}} = \mathrm{char}(\mathrm{ht}_{\mathrm{anti}}(\mathcal{U} \otimes \mathcal{U}^\iota)/\mathrm{ht}_{\mathrm{anti}}(\mathcal{H} \otimes \mathcal{H}^\iota))$ where $\mathcal{H} \subset \mathcal{U}$ is the submodule of universal norms of Heegner points.

Theorem 10.5 (*Howard [Ho2]*). $\mathcal{L}_p(E)\Lambda_{\mathrm{anti}} = \mathrm{ht}_{\mathrm{anti}}(\mathcal{H} \otimes \mathcal{H}^\iota)$ *in* $(I_{K^{\mathrm{anti}}}/I_{K^{\mathrm{anti}}}^2)$.

Theorem 10.6 (*Howard [Ho1]*). *If the p-adic representation on* $E[p^\infty]$

$$\mathrm{Gal}(\bar{K}/K) \longrightarrow \mathrm{Aut}_{\mathbf{Z}_p}(E[p^\infty]) \longrightarrow \mathrm{GL}_2(\mathbf{Z}_p)$$

is surjective, then

$$\mathrm{char}(X_{\mathrm{tors}}^{\mathrm{anti}})\ \textit{divides}\ \mathrm{char}(\mathrm{ht}_{\mathrm{anti}}(\mathcal{U} \otimes \mathcal{U}^\iota)/\mathrm{ht}_{\mathrm{anti}}(\mathcal{H} \otimes \mathcal{H}^\iota)).$$

Corollary 10.7. *If the p-adic representation* $\mathrm{Gal}(\bar{K}/K) \to \mathrm{GL}_2(\mathbf{Z}_p)$ *is surjective, then*

$$(\theta^{-1}\mathrm{char}(X_\infty))\Lambda_{\mathrm{anti}}\ \textit{divides}\ (\theta^{-1}\mathcal{L}_p(E))\Lambda_{\mathrm{anti}},$$

with equality if and only if Conjecture 10.4 *holds*.

Proof. Combine Howard's Theorems 10.5 and 10.6 with Theorem 10.2. □

Proposition 10.8. *If* $X_{\mathrm{tors}}^{\mathrm{anti}} = 0$, *then* X_∞ *is a cyclic* Λ-module, and $\mathcal{S}_p(E, K) \cong \mathbf{Z}_p$.

Proof. If $X^{\mathrm{anti}}_{\mathrm{tors}} = 0$ then by (10.2) the Λ-module $(\Phi \otimes \Lambda_{\mathrm{anti}})/\mathcal{U}$ is free. Hence the organizing matrix H has the form given by (10.3), and the submatrix B of (10.3) is invertible by (10.5). But all the entries of H are in the maximal ideal $\mathfrak{m}$ of Λ, so this is possible only if $r = 1$, i.e., H is a 1×1 matrix. Thus

$$\dim_{\mathbf{F}_p} \mathcal{S}_p(E,K)/p\mathcal{S}_p(E,K) = \dim_{\mathbf{F}_p} X_\infty/\mathfrak{m}X_\infty = 1.$$

Since X^{anti} has positive Λ_{anti}-rank, $\mathcal{S}_p(E,K)$ must be infinite and the proposition follows. □

Example 10.9. *An example of a nonzero submodule in* X^{anti}. Let E be the elliptic curve

$$y^2 + xy = x^3 + x^2 - 34x - 135,$$

1913B1 in Cremona's tables [Cr]. We take $p = 3$, and $K = \mathbf{Q}(\sqrt{-2})$. Note that E has good ordinary reduction at 3, both 3 and 1913 split in K, the Tamagawa number $c_{1913} = 2$, and the Fourier coefficient $a_3 = 2$. Thus, all of our hypotheses (7.1-5) hold. We have $E(K) \cong \mathbf{Z} \oplus \mathbf{Z}/2\mathbf{Z}$ and $Ш(E,K) \cong (\mathbf{Z}/3\mathbf{Z})^2$. Thus, the organizing matrix H is 3×3 in this case.

For every $n \geqslant 0$ let K_n denote the extension of K of degree 3^n inside K^{anti}. Let $\mathcal{H}_n \subset E(K_n) \otimes \mathbf{Z}_3$ be the $\mathbf{Z}_3[\mathrm{Gal}(K_n/K)]$-submodule generated by Heegner points in $E(K_n)$. A computation shows that the Heegner point in $E(K)$ is

$$\left(-\tfrac{71}{18} - \tfrac{29}{18}\sqrt{-2}, \tfrac{299}{54} + \tfrac{145}{108}\sqrt{-2}\right)$$

and from this it follows easily that $\mathcal{H}_0 = 3E(K) \otimes \mathbf{Z}_3$. By computing the Heegner points in K_1, and dividing by 3 where possible, one can compute generators of $E(K_1)/3E(K_1)$ and verify that

$$\mathrm{Tr}_{K_1/K} E(K_1) = 3E(K).$$

Thus, the image of the projection $\mathcal{U} \to E(K) \otimes \mathbf{Z}_3$ is $\mathcal{H}_0$. Since the Fourier coefficient $a_3 = 2$, every Heegner point is a universal norm of Heegner points (see for example [M2]), so the projection $\mathcal{H} \to \mathcal{H}_0$ is surjective. Since $\mathcal{U}$ is free of rank one over Λ_{anti}, it follows that $\mathcal{U} = \mathcal{H}$.

We also compute, using the techniques of [Se] (especially §IV.3.2), that the 3-adic representation $\mathrm{Gal}(\bar{K}/K) \to \mathrm{GL}_2(\mathbf{Z}_3)$ is surjective, so we deduce from Howard's Theorem 10.6 that $X^{\mathrm{anti}}_{\mathrm{tors}}$ is finite. But $\mathcal{S}_p(E,K) \cong \mathbf{Z} \times (\mathbf{Z}/3\mathbf{Z})^2$, so by Proposition 10.8 we cannot have $X^{\mathrm{anti}}_{\mathrm{tors}} = 0$. Thus X^{anti} has a nonzero finite submodule, namely $X^{\mathrm{anti}}_{\mathrm{tors}}$.

For related work on the possibility of nonzero finite submodules of X^{anti}, see [B].

Example 10.10. Counterexamples to one of our conjectures from [MR2]. Let E be the elliptic curve

$$y^2 + y = x^3 - x,$$

37A1 in Cremona's tables [Cr]. We have $E(\mathbf{Q}) \cong \mathbf{Z}$, generated by $P := (0, 0)$, and $Ш(E, \mathbf{Q}) = 0$.

Let $K := \mathbf{Q}(\sqrt{-3})$. Then 37 splits in K, and $E(K) = E(\mathbf{Q}) \cong \mathbf{Z}$, $Ш(E, K) = 0$. If $p > 3$, $p \neq 37$, and the Fourier coefficient $a_p \neq 0, 1$ then all of our hypotheses (7.1-5) are satisfied.

Since $\mathrm{Sel}_p(E/K) \cong \mathbf{Q}_p/\mathbf{Z}_p$, the skew-Hermitian organizing matrix is 1×1, i.e., $H = (\theta a)$ in (10.3) for some $a \in \Lambda$. Arguing as in §9.1, if $I = I_K$ denotes the augmentation ideal of Λ then the image of θa in I/I^2 is (up to a unit) the p-adic height of P.

Let $h_p(P)$ denote the p-adic height of P. William Stein (using methods of [MST]) has computed $h_p(P)$ for all primes less than 100 of good ordinary reduction and with $a_p \neq 1$, and in all cases except $p = 13$ and 67, $h_p(P)$ generates (the free, rank-one $\mathbf{Z}_p$-module) $\theta\Lambda/I^2 \subset I/I^2$. In the two exceptional cases $h_p(P)$ generates $p(\theta\Lambda/I^2)$. Thus 13 and 67 are irregular in the sense of Definition 9.1.1.

Recall that by (10.6), $\mathrm{ht}_{\mathrm{anti}}(\mathcal{U} \otimes \mathcal{U}^\iota) = a(I_{K^{\mathrm{anti}}}/I^2_{K^{\mathrm{anti}}})$. Thus, if p is one of the 17 primes less than 100 and different from 13 and 67 where E has good ordinary reduction and $a_p \neq 1$, then $a \in \Lambda^\times$ and $\mathrm{ht}_{\mathrm{anti}}(\mathcal{U} \otimes \mathcal{U}^\iota) = I_{K^{\mathrm{anti}}}/I^2_{K^{\mathrm{anti}}}$. But if $p = 13$ or 67 then $a \notin \Lambda^\times$ and $\mathrm{ht}_{\mathrm{anti}}(\mathcal{U} \otimes \mathcal{U}^\iota) \neq I_{K^{\mathrm{anti}}}/I^2_{K^{\mathrm{anti}}}$ (so p is irregular for E/K^{anti} in the sense of Remark 10.1). These last two cases give counterexamples to Conjecture 6.1 of [MR2] (see also Conjecture 6 of [MR1]).

Appendix A. Perfect control assumption

We keep the notation of the body of the paper. In particular E is an elliptic curve over a number field K, with good ordinary reduction at all primes above p, and K_∞ is the maximal $\mathbf{Z}_p$-power extension of K.

We will use the following theorem of Greenberg ([G1] §5.I).

Theorem A.1 (*Greenberg*). *Suppose that F is a finite extension of K and L/F is a $\mathbf{Z}_p$-extension. Suppose further that*

(i) *$E(F)$ has no point of order p,*
(ii) *for every prime w of F above p, $E(f_w)$ has no point of order p, where f_w is the residue field of F at w,*
(iii) *for every prime w of F where E has bad reduction, either $E(F_w)$ has no point of order p or $E(F_w^{\mathrm{unr}})[p^\infty]$ is divisible.*

Then the natural map $\mathrm{Sel}_p(E, F) \to \mathrm{Sel}_p(E, L)^{\mathrm{Gal}(L/F)}$ *is an isomorphism.*

Lemma A.2. *Suppose A is an elliptic curve defined over a field k, p is a prime, and ℓ is an abelian (pro-)p-extension of k. If $A(k)$ has no point of order p then $A(\ell)$ has no point of order p.*

Proof. By Nakayama's Lemma, if $A(\ell) \cap A[p] \neq 0$, then

$$A(k) \cap A[p] = (A(\ell) \cap A[p])^{\mathrm{Gal}(\ell/k)} \neq 0. \qquad \square$$

Corollary A.3. *Suppose*

(i) *$E(K)$ has no point of order p,*
(ii) *for every prime v of K above p, $E(k_v)$ has no point of order p, where k_v is the residue field of K at v,*
(iii) *for every prime v of K where E has bad reduction, either $E(K_v)$ has no point of order p or $E(K_v^{\mathrm{unr}})[p^\infty]$ is divisible.*

If $K \subset F \subset F' \subset K_\infty$ then the natural map $\mathrm{Sel}_p(E, F) \to \mathrm{Sel}_p(E, F')^{\mathrm{Gal}(F'/F)}$ is an isomorphism.

In particular, the Perfect Control assumption holds.

Proof. Suppose v is a prime of K, $K \subset F \subset K_\infty$, and w is a prime of F above v.

If $v \nmid p$ then $F_w^{\mathrm{unr}} = K_v^{\mathrm{unr}}$, so assumption (iii) and Lemma A.2 imply assumption (iii) of Theorem A.1 for F. If $v \mid p$ then the residue field f_w is a p-extension of k_v, so assumption (ii) and Lemma A.2 imply assumption (ii) of Theorem A.1 for F. Finally, assumption (i) and Lemma A.2 imply assumption (i) of Theorem A.1 for F.

It is enough to prove the corollary when F is a finite extension of K, and then pass to the limit for general F. Further, it is enough to consider the case where F'/F is cyclic, because every extension of F in K_∞ can be given as a finite chain of cyclic extensions.

So suppose that F'/F is cyclic. Then there is a $\mathbf{Z}_p$-extension L of F in K_∞ containing F'. The hypotheses of Theorem A.1 are satisfied for F, so if $F' = L$ then the statement of the corollary is just the conclusion of Theorem A.1. If $F' \neq L$ then the hypotheses of Theorem A.1 are satisfied for F' as well, and we conclude from Theorem A.1 that

$$\begin{aligned}\mathrm{Sel}_p(E, F) &= \mathrm{Sel}_p(E, L)^{\mathrm{Gal}(L/F)} = (\mathrm{Sel}_p(E, L)^{\mathrm{Gal}(L/F')})^{\mathrm{Gal}(F'/F)} \\ &= \mathrm{Sel}_p(E, F')^{\mathrm{Gal}(F'/F)}. \qquad \square\end{aligned}$$

Remark A.4. There are a few comments to make about the hypotheses in Corollary A.3.

For a fixed elliptic curve E, hypotheses (i) and (iii) hold for all but finitely many primes p. Condition (ii) can fail to hold; this is the *anomalous* case of [M1]. Condition (ii) should hold for "most" p, but it could fail for infinitely many p. However, we have the following lemma.

Lemma A.5. *Suppose that $E(K)$ has a point of finite order $\ell > 1$. Then for every rational prime $p > 5$, $p \neq \ell$ and every prime v of K of degree one dividing p where E has good reduction, $E(k_v)$ has no point of order p.*

Proof. Fix such a v and suppose that $E(k_v)$ has a point of order p. Our assumptions guarantee that $E(k_v)$ has a point of order ℓ as well, so $\#E(k_v) \geqslant p\ell$. Since v has degree one we have $\#E(k_v) - (p+1) < 2\sqrt{p}$, and this is impossible if $p > 5$. □

We also have the following lemma relating hypotheses (i) and (ii) of Corollary A.3.

Lemma A.6. *Suppose that for some prime v of K above p with residue field k_v* (*where as usual we suppose that E has good, ordinary reduction*), *$E(k_v)$ has no point of order p. If K_v^{unr} does not contain a primitive pth root of unity then $E(K_v)$ has no point of order p.*

In particular if the ramification of $K_v/\mathbf{Q}_p$ is not divisible by $p-1$ then $E(K_v)$ has no point of order p, and so $E(K)$ has no point of order p.

Proof. If $E(K_v)$ has a point of order p, it must be in the kernel of reduction. But since E has good ordinary reduction at v, the inertia group at v acts on the kernel of reduction via the cyclotomic character. This proves the lemma. □

Appendix B. Some commutative algebra with group rings

For this appendix suppose that G is a finite abelian group, R is a commutative ring, and let $\iota : R[G] \to R[G]$ be the R-linear involution that sends $g \mapsto g^{-1}$ for $g \in G$. As in §2, if M is an $R[G]$-module we let M^ι denote the $R[G]$-module whose underlying abelian group is M, but with the action of G obtained from that if M by composition with ι.

Suppose that A is an $R[G]$-module and B is an R-module with trivial G-action.

Lemma B.1. *There is a natural isomorphism*

$$\mathrm{Hom}_{R[G]}(A, B \otimes R[G]) \longrightarrow \mathrm{Hom}_R(A, B)^\iota.$$

Proof. Let $\pi : R[G] \to R$ denote the projection map $\pi(\sum_g a_g g) := a_1$. Composition with π defines an R-module homomorphism

$$\mathrm{Hom}_{R[G]}(A, B \otimes_R R[G]) \longrightarrow \mathrm{Hom}_R(A, B)^\iota \tag{B.1}$$

and it is straightforward to check that this is a morphism of $R[G]$-modules. The inverse of (B.1) is given by sending $f \in \mathrm{Hom}_R(A, B)^\iota$ to the map

$$a \mapsto \sum_g f(a^g) \otimes g^{-1},$$

and it follows that (B.1) is an isomorphism. □

Now consider the composition

$$\begin{aligned}
&\mathrm{Hom}_G(A \otimes_{R[G]} A^\iota, B \otimes_R R[G]) \overset{\sim}{\to} \mathrm{Hom}_G(A^\iota, \mathrm{Hom}_G(A, B \otimes_R R[G])) \\
&\overset{\sim}{\to} \mathrm{Hom}_G(A, \mathrm{Hom}_G(A, B \otimes_R R[G])^\iota) \overset{\sim}{\to} \mathrm{Hom}_G(A, \mathrm{Hom}_R(A, B)) \\
&\to \mathrm{Hom}_R(A \otimes_R A, B), \qquad \text{(B.2)}
\end{aligned}$$

where the third isomorphism comes from Lemma B.1. This composition sends a $B \otimes R[G]$-valued, $R[G]$-bilinear pairing on $A \times A^\iota$ to a B-valued, R-bilinear pairing on $A \times A$.

Proposition B.2. *Suppose that $i : B \to B$ is an R-linear involution, and that $\pi : A \otimes_{R[G]} A^\iota \to B \otimes_R R[G]$ is a skew-Hermitian pairing, i.e.,*

$$\pi(a' \otimes a) = -(i \otimes \iota)(\pi(a \otimes a')).$$

Then the pairing $\pi_0 : A \otimes_R A \to B$ induced from π via (B.2) *is i-skew symmetric, i.e.,*

$$\pi_0(a' \otimes a) = -i(\pi_0(a \otimes a')).$$

In particular if i is the identity then π_0 is skew-symmetric, and if i is multiplication by -1 then π_0 is symmetric.

Proof. Straightforward. □

Appendix C. The structure of Selmer modules

One weak consequence of the existence of a skew-Hermitian module Φ that organizes the arithmetic of E over K_∞ is that the Λ-module $\mathcal{S}_p(E, K_\infty)$ has a free resolution of length two. In this appendix we give a direct proof of this fact, under some mild hypotheses, without appealing to the work of Nekovář [N] which was the basis for our proof of Theorem 7.5.

We continue to suppose that E has good ordinary reduction at all primes above p, the Perfect Control assumption holds, and we will make the following two additional assumptions for this section.

Torsion assumption. $\mathcal{S}_p(E, K_\infty)$ *is a torsion Λ-module.*

Local Nontriviality assumption. *For some prime $\mathfrak{p}$ of K above p, the decomposition group of $\mathfrak{p}$ in G_K acts nontrivially on the kernel of reduction modulo $\mathfrak{p}$ in $E[p]$.*

Remark C.1. If $K(E[p])/K$ is ramified at some prime above p then the Local Nontriviality assumption holds, so in particular (since $\boldsymbol{\mu}_p \subset K(E[p])$) it holds if p is odd and unramified in $K/\mathbf{Q}$.

Theorem C.2 (*Greenberg*). *If L is a $\mathbf{Z}_p^d$-extension of K, then $\mathcal{S}_p(E, L)$ has no nonzero pseudo-null Λ_L-submodules.*

Proof. This is proved by Greenberg [G2], using the Torsion and Local Nontriviality assumptions. □

Proposition C.3. *Suppose L is a $\mathbf{Z}_p^d$-extension of K, $\mathcal{S}_p(E, L)$ is a torsion Λ_L-module, M is a free Λ_L module of finite rank, and $f : M \twoheadrightarrow \mathcal{S}_p(E, L)$ is a surjective map of Λ_L-modules. Then* $\ker(f)$ *is free over Λ_L.*

Proof. The proof will be by induction on d, where $\mathrm{Gal}(L/K) \cong \mathbf{Z}_p^d$. If $d = 0$ then $L = K$, $\Lambda_L = \mathbf{Z}_p$, and there is nothing to prove.

Let $N := \ker(f)$. Then N is a finitely generated torsion-free Λ_L-module, so the structure theorem for such modules says that there is an exact sequence

$$0 \longrightarrow N \longrightarrow S \longrightarrow Z \longrightarrow 0,$$

where S is a reflexive Λ-module and Z is pseudo-null.

Let $\mathcal{K}$ denote the field of fractions of Λ_L. The inclusion $N \hookrightarrow M$ extends uniquely to an inclusion $S \hookrightarrow M \otimes \mathcal{K}$. Since S/N is pseudo-null and $\mathcal{K}/\Lambda_L$ has no nonzero pseudo-null Λ_L-submodules, we must have $S \hookrightarrow M \subset M \otimes \mathcal{K}$. But then

$$Z = S/N \hookrightarrow M/N \cong \mathcal{S}_p(E, L)$$

so by Greenberg's Theorem C.2 we must have $Z = 0$, and so $N = S$ is reflexive.

It remains to show that N is free. If $d = 1$ then every reflexive module is free, so we may assume that $d \geqslant 2$. Since $\mathcal{S}_p(E, L)$ is a torsion Λ_L-module, for all but finitely many $\mathbf{Z}_p^{d-1}$ extensions F of K contained in L we have (using the Perfect Control assumption) that $\mathcal{S}_p(E, F) = \mathcal{S}_p(E, L) \otimes_{\Lambda_L} \Lambda_F$ is a torsion Λ_F-module. For such an F, writing $H := \mathrm{Gal}(L/F) \cong \mathbf{Z}_p$, we have an exact sequence

$$0 \longrightarrow \mathcal{S}_p(E, L)^H \longrightarrow N \otimes \Lambda_F \longrightarrow M \otimes \Lambda_F \longrightarrow \mathcal{S}_p(E, F) \longrightarrow 0,$$

Since $\mathcal{S}_p(E, L) \otimes_{\Lambda_L} \Lambda_F$ is a torsion Λ_F-module, $\mathcal{S}_p(E, L)^H$ is a pseudo-null Λ_L-module (see for example Lemma 4 of §I.1.3 of [PR1]). Again using Greenberg's Theorem C.2 we conclude that $\mathcal{S}_p(E, L)^H = 0$, and so

$$\mathcal{S}_p(E, F) \cong (M \otimes \Lambda_F)/(N \otimes \Lambda_F).$$

We conclude from our induction hypothesis that $N \otimes \Lambda_F$ is a free Λ_F-module of rank $t := \mathrm{rank}_{\Lambda_F}(M \otimes \Lambda_F) = \mathrm{rank}_{\Lambda_L} M$. By Nakayama's Lemma N can be generated over Λ_L by t generators, and since (by the Torsion assumption) $\mathrm{rank}_{\Lambda_L} N = \mathrm{rank}_{\Lambda_L} M = t$, N must be free as claimed. □

Theorem C.4. *There are free Λ-modules $N \subset M$ such that $\mathcal{S}_p(E, K_\infty) \cong M/N$. If $t := \dim_{\mathbf{F}_p} \mathrm{Sel}_p(E, K)[p]$ then we can take M and N to have Λ-rank t.*

Proof. By Lemma 2.2, we have

$$\mathcal{S}_p(E, K_\infty)/\mathfrak{m}\mathcal{S}_p(E, K_\infty) \cong \mathrm{Hom}(\mathrm{Sel}_p(E, K)[p], \mathbf{F}_p) \cong \mathbf{F}_p^t,$$

where $\mathfrak{m}$ is the maximal ideal of Λ. By Nakayama's Lemma there is a surjection $\Lambda^t \twoheadrightarrow \mathcal{S}_p(E, K_\infty)$, and by Proposition C.3 the kernel of this surjection is also free. □

References

[B] M. Bertolini, Iwasawa theory for elliptic curves over imaginary quadratic fields, J. Théor. Nombres Bordeaux 13 (2001) 1–25.

[BD1] M. Bertolini, H. Darmon, Derived p-adic heights, Amer. J. Math. 117 (1995) 1517–1554.

[BD2] M. Bertolini, H. Darmon, Derived heights and generalized Mazur–Tate regulators, Duke Math. J. 76 (1994) 75–111.

[Co] C. Cornut, Mazur's conjecture on higher Heegner points, Invent. Math. 148 (2002) 495–523.

[Cr] J.E. Cremona, Algorithms for Modular Elliptic Curves, Cambridge University Press, Cambridge, 1992.

[CM] J.E. Cremona, B. Mazur, Visualizing elements in the Shafarevich–Tate group, Experiment. Math. 9 (2000) 13–28.

[D] P. Deligne, Le déterminant de la cohomologie, in: Current Trends in Arithmetical Algebraic Geometry, Arcata, CA, 1985, Contemp. Math. 67 (1987) 93–177.

[G1] R. Greenberg, Galois theory for the Selmer group of an abelian variety, Compositio Math. 136 (2003) 255–297.

[G2] R. Greenberg, in preparation.

[GS] R. Greenberg, G. Stevens, On the conjecture of Mazur, Tate, and Teitelbaum, in: B. Mazur, G. Stevens (Eds.), p-adic Monodromy and the Birch and Swinnerton–Dyer Conjecture, Contemp. Math. 165 (1994) 183–211.

[Hara] S. Haran, p-adic L-functions for elliptic curves over CM fields, Thesis, MIT, 1983.

[Hart] R. Hartshorne, Residues and Duality, Lecture Notes in Mathematics, vol. 20, Springer, Berlin, 1960.

[Hi] H. Hida, A p-adic measure attached to the zeta functions associated with two elliptic modular forms I, Invent. Math. 79 (1985) 159–195.

[Ho1] H. Howard, The Heegner point Kolyvagin system, Compositio Math. 140 (2004) 1439–1472.

[Ho2] H. Howard, The Iwasawa theoretic Gross–Zagier theorem, Compositio Math. 141 (2005) 811–846.

[Ka] K. Kato, p-adic Hodge theory and values of zeta functions of modular forms, in: Cohomologies p-adiques et applications arithmétiques III, Astérisque 295 (2004) 117–290.

[Ki] K. Kitagawa, On standard p-adic L-functions of families of elliptic cusp forms. in: B. Mazur, G. Stevens (Eds.), p-adic Monodromy and the Birch and Swinnerton–Dyer Conjecture, Contemp. Math. 165 (1994) 81–110.

[M1] B. Mazur, Rational points of abelian varieties with values in towers of number fields, Invent. Math. 18 (1972) 183–266.

[M2] B. Mazur, Modular curves and arithmetic, in: Proceedings of the International Congress of Mathematicians (Warsaw 1983), PWN, Warsaw, 1984, pp. 185–211.

[MR1] B. Mazur, K. Rubin, Elliptic curves and class field theory, in: Ta Tsien Li (Ed.), Proceedings of the International Congress of Mathematicians ICM 2002, vol. II, Higher Education Press, Beijing, 2002, pp. 185–195.

[MR2] B. Mazur, K. Rubin, Pairings in the arithmetic of elliptic curves, in: J. Cremona et al. (Ed.), Modular Curves and Abelian Varieties, Prog. Math., 224, 2004, pp. 151–163.

[MR3] B. Mazur, K. Rubin, Studying the growth of Mordell–Weil, Documenta Math. extra volume (2003) 585–607.

[MST] B. Mazur, W. Stein, J. Tate, Computation of p-adic heights and log convergence, to appear.

[N] J. Nekovář, Selmer complexes, Preprint available at <http://www.math.jussieu.fr/~nekovar/pu/>.

[PR1] B. Perrin-Riou, Arithmétique des courbes elliptiques et théorie d'Iwasawa, Bull. Soc. Math. Suppl. Mémoire 17 (1984).

[PR2] B. Perrin-Riou, Fonctions L p-adiques associées à une forme modulaire et à un corps quadratique imaginaire, J. London Math. Soc. 38 (1988) 1–32.

[PR3] B. Perrin-Riou, Théorie d'Iwasawa et hauteurs p-adiques (cas des variétés abéliennes), Séminaire de Théorie des Nombres, Paris, 1990–1991. Prog. Math., 108, 1993, pp. 203–220.

[PR4] B. Perrin-Riou, Groupes de Selmer et accouplements: cas particulier des courbes elliptiques, Documenta Math. extra volume (2003) 725–760.

[Se] J.-P. Serre, Abelian l-Adic Representations and Elliptic Curves, W. A. Benjamin Inc., New York, 1968.

[T] J. Tate, Algorithm for determining the type of a singular fiber in an elliptic pencil, Modular Functions of One Variable (IV), Lecture Notes in Mathematics, vol. 476, Springer, New York, 1975, pp. 33–52.

[V] V. Vatsal, Special values of anticyclotomic L-functions, Duke Math. J. 116 (2003) 219–261.

[W] C. Wuthrich, On p-adic heights in families of elliptic curves, J. London Math. Soc. 70 (2004) 23–40.