An averaged Iwasawa approach to Mordell over $\mathbb{Q}$

Proof candidate for independent review

20 September 2026

Abstract

We develop a proposed direct argument for finiteness of rational points on a smooth projective curve over $\mathbb{Q}$ of genus at least two. The framework is the unipotent Albanese method of Kim. Its proposed additional ingredient is an aggregate bound for algebraic multiplicities of a torsion Iwasawa module, summed over a rectangular window of symplectic highest weights and scalar characters. Character concentration for the graded surface Lie algebra then gives an averaged, rather than pointwise, bound for the global H^2 obstruction. The average suffices for a cumulative Selmer dimension inequality along a subsequence of depths.

The construction uses the actual Galois image only as a closed subgroup of the full symplectic similitude group. Transferring nonzero annihilators to the ambient Iwasawa algebra removes any dependence on Tate-module semisimplicity or reductivity of actual monodromy. Bogomolov's homothety theorem is retained. The argument uses neither Parshin's construction nor the Faltings–Vojta auxiliary-section method.

Status and scope. This manuscript records a proof candidate and supplies the calculations for mathematical review. It is not an announcement of an established new proof, and makes no priority claim. The aggregate multiplicity argument and its combination with the character estimates require independent verification. The proposed conclusion is over $\mathbb{Q}$; the extension to arbitrary number fields is not established here. The draft was prepared with AI assistance and has not undergone external mathematical refereeing.

hexagon:2611.00576v1

Contents

1 The proposed conclusion and its dependencies 3

2 Aggregate algebraic multiplicities and ambient transfer 3

2.1 The ambient group and the assertion 4
2.2 An injective completed-enveloping map 4
2.3 Integral image orders for a rectangular family 6
2.4 A mixed PBW quotient bound 7
2.5 Proof of the aggregate estimate 7
2.6 Transfer from a closed subgroup 8

3 Characters and concentration for the symplectic surface Lie algebra 9

3.1 The equivariant Hilbert series and the Witt formula 10
3.2 A compact-group estimate in the full group dimension 12
3.3 A Gaussian tail for highest weights 14

4 Averaging over depth 16

5 The arithmetic reduction over the rational numbers 17

5.1 The actual Galois image and its ambient group 17
5.2 The unramified module and induction 18
5.3 Strict localization kernels and global second cohomology 19
5.4 The local terms, including bad reduction 20
5.5 From a degree average to the cumulative dimensions 22
5.6 Complex conjugation and the final dimension comparison 22

6 Scope, dependencies and review obligations 24

6.1 What semisimplicity has been replaced by 24
6.2 Why the argument currently stops at $\mathbb{Q}$ 25
6.3 The points on which the candidate should be tested 25

References 26

1 The proposed conclusion and its dependencies

Let $C/\mathbb{Q}$ be smooth, projective and geometrically connected, of genus $g \geq 2$, and let J be its Jacobian. If $C(\mathbb{Q})$ is empty there is nothing to prove; otherwise choose a rational base point. Fix a prime $p \geq 5$ of good reduction. Let U be the geometric pro-unipotent $\mathbb{Q}_p$-fundamental group of C, equipped with its Galois action, and write

$$V_n = \mathrm{gr}_n \operatorname{Lie} U, \qquad d_n = \dim V_n, \qquad W_n = V_n^*(1).$$

The lower central quotient of depth N is denoted by U_N. Fix one finite set T of places containing p, infinity and the bad primes; we may enlarge it by a fixed finite set when making a fixed auxiliary field extension. Put

$$h_n = \dim H^2(G_{\mathbb{Q},T}, V_n), \qquad D_N = \sum_{n \leq N} d_n, \qquad \lambda = g + \sqrt{g^2 - 1}.$$

All cohomology is continuous. All implied constants may depend on the fixed curve, prime, auxiliary fields, groups and modules, but not on n, N or the highest-weight cutoffs.

Proposition 1.1 (Conclusion sought by the candidate)**.** *Assuming the aggregate multiplicity estimate and the surface-Lie character estimates proved provisionally below are valid in their stated categories, the argument gives*

$$\frac{1}{N} \sum_{N \leq n < 2N} \frac{h_n}{d_n} \longrightarrow 0. \tag{1.1}$$

Consequently, for some finite depth N, it gives

$$\dim H^1_f(G_{\mathbb{Q},T}, U_N) < \dim(U_N^{\mathrm{dR}}/F^0). \tag{1.2}$$

Kim's dimension criterion then implies that the Chabauty–Kim set at that depth is finite, and in particular that $C(\mathbb{Q})$ is finite.

The intended advance is the proof of the estimates underlying this conditional deduction, not the dimension criterion itself. Eventual finiteness of Chabauty–Kim sets should not be confused with their eventual equality with $C(\mathbb{Q})$; no equality assertion is made.

The arithmetic inputs are Bogomolov's scalar theorem [2], the unramified Iwasawa torsion theorem in Venjakob [11], global and local duality, and Kim's unipotent Albanese theory [8]. The analytic algebra input is the relation between bounded measures and completed enveloping algebras [1]. The full symplectic group provides representation-theoretic complete reducibility; semisimplicity of the Galois action is never assumed.

Kim already proposed noncommutative Iwasawa theory as a route to these cohomological bounds [9]. Fu's multiplicity results for products of SL_2 [7] motivated an initial special-case calculation, but the rectangular aggregate estimate below is not attributed to Fu. An earlier fixed-rank auxiliary-section construction in this investigation was identified with the curve case of the established Faltings–Vojta method [6]; it is not used here. Xie–Yang's new proof of semisimplicity [12] motivated the investigation, but the ambient-group argument removes it as an input.

2 Aggregate algebraic multiplicities and ambient transfer

This section supplies the algebraic estimate proposed for the argument. Its proof is recorded in detail for independent audit, particularly at the passage from an integral image order

to its reduction modulo p. The estimate concerns a family of representations, rather than an individual highest weight. The use of completed enveloping algebras is related to the methods of [1, 7]; the aggregate estimate below is not being attributed to either reference.

2.1 The ambient group and the assertion

Fix an odd prime p, and set

$$S = \mathrm{Sp}_{2g,\mathbb{Q}_p}, \qquad D = \dim S = g(2g+1).$$

Use the standard split integral model of S, and put

$$\mathcal{S} = \ker\bigl(\mathrm{Sp}_{2g}(\mathbb{Z}_p) \longrightarrow \mathrm{Sp}_{2g}(\mathbb{F}_p)\bigr), \qquad Z = 1 + p\mathbb{Z}_p, \qquad \mathcal{G} = Z \times \mathcal{S}.$$

These are uniform pro-p groups. The map $(u, s) \mapsto us$ identifies $\mathcal{G}$ with a compact open subgroup of $\mathrm{GSp}_{2g}(\mathbb{Q}_p)$. Its inverse on this subgroup uses the unique square root near 1 of the multiplier. Let χ denote the character of the first factor, and write

$$\Lambda_0 = \mathbb{Z}_p[[\mathcal{G}]], \qquad \Lambda = \Lambda_0[1/p].$$

The notation Λ means the bounded rational Iwasawa algebra. In particular, its elements have a common bound on the denominators of their ordered Iwasawa coefficients.

Choose the fundamental weights $\omega_1, \dots, \omega_g$ of S. For a dominant weight $\lambda = \sum_i a_i\omega_i$, with $a_i \geq 0$, set $\|\lambda\| = \sum_i a_i$. Let E_λ be the absolutely irreducible algebraic $\mathbb{Q}_p$-representation of highest weight λ, and put $e_\lambda = \dim E_\lambda$. The representations $E_\lambda \otimes \chi^j$ are understood as representations of $S \times \mathbb{G}_m$, restricted to $\mathcal{G}$. No parity condition is imposed in this family. The representations which descend to the algebraic group GSp_{2g} form a subfamily, which is sufficient for the application.

Proposition 2.1 (Proposed aggregate multiplicity estimate)**.** *Let M be a finitely generated torsion left Λ-module. There is a constant C_M such that, for all integers $K, N \geq 0$,*

$$\begin{aligned}\sum_{\substack{\lambda \text{ dominant} \\ \|\lambda\| \leq K}} \sum_{|j| \leq N} e_\lambda \, \dim_{\mathbb{Q}_p} \mathrm{Hom}_\Lambda(M, E_\lambda \otimes \chi^j) \\ \leq C_M\bigl((N+1)(K+1)^{D-1} + (K+1)^D\bigr).\end{aligned} \tag{2.1}$$

The constant may depend on g, p, the fixed norm on weights, and M, but not on K, N.

We use left modules throughout. For a nonzero element δ of a left Ore domain, $\Lambda\delta$ is the left ideal it generates. Thus

$$\mathrm{Hom}_\Lambda(\Lambda/\Lambda\delta, W) \simeq \ker(\delta : W \longrightarrow W).$$

2.2 An injective completed-enveloping map

Let

$$\mathfrak{s}_{\mathbb{Z}_p} = \mathrm{Lie}(\mathrm{Sp}_{2g,\mathbb{Z}_p}), \qquad \mathfrak{a} = \mathfrak{s}_{\mathbb{Z}_p} \oplus \mathbb{Z}_p z, \qquad U_0 = U_{\mathbb{Z}_p}(\mathfrak{a}).$$

Here z acts on the character χ^j as the integer j. Define

$$A_0 = \varprojlim_r U_0/p^r U_0, \qquad A = A_0[1/p].$$

An element of A_0 is an ordered PBW series whose coefficients tend to zero p-adically. In particular,

$$A_0/pA_0 \simeq U_{\mathbb{F}_p}(\mathfrak{s}_{\mathbb{F}_p})[z]. \tag{2.2}$$

Lemma 2.2. *There is an injective algebra homomorphism*

$$\iota : \Lambda \longrightarrow A$$

which carries Λ_0 into A_0. Its action on each algebraic representation $E_\lambda \otimes \chi^j$ agrees with the usual action of its restriction to $\mathcal{G}$.

Proof. We recall the ordered-coordinate input for a uniform group algebra. If $g_1, \dots, g_{D+1}$ is an ordered minimal generating set of $\mathcal{G}$ and $b_i = g_i - 1$, then every element of Λ_0 has a unique expansion

$$f = \sum_{\alpha \in \mathbb{N}^{D+1}} a_\alpha b_1^{\alpha_1} \cdots b_{D+1}^{\alpha_{D+1}}, \qquad a_\alpha \in \mathbb{Z}_p. \tag{2.3}$$

The coefficients in this expansion need not tend to zero. This standard uniform-group description and the corresponding completion construction are discussed in [1, §§10.3–10.4].

Choose an ordered integral basis $x_1, \dots, x_{D+1}$ of $\mathfrak{a}$. Since $\log \mathcal{G} = p\mathfrak{a}$, we may take $g_i = \exp(px_i)$. The exponential series converges in A_0 on $p\mathfrak{a}$, and the Campbell–Hausdorff identity gives a homomorphism from $\mathcal{G}$ to $A_0^\times$. It sends

$$b_i \longmapsto \exp(px_i) - 1 \in pA_0.$$

Consequently (2.3) has a convergent image in A_0, and gives a homomorphism $\Lambda_0 \to A_0$.

We check injectivity explicitly. Because p is odd,

$$\exp(px_i) - 1 \equiv px_i \pmod{p^2 A_0}.$$

For the ordered multi-index monomials this implies

$$\iota(b^\alpha) \equiv p^{|\alpha|} x^\alpha \pmod{p^{|\alpha|+1} A_0}. \tag{2.4}$$

No reordering is used in this congruence. For $f \neq 0$, let

$$q = \min_{a_\alpha \neq 0} \big(v_p(a_\alpha) + |\alpha|\big).$$

Only finitely many multi-indices attain this minimum, since $|\alpha| \leq q$ for all of them. Dividing the image by p^q and reducing modulo p gives

$$p^{-q}\iota(f) \bmod p = \sum_{v_p(a_\alpha)+|\alpha|=q} \overline{p^{-v_p(a_\alpha)} a_\alpha}\, \overline{x}^\alpha.$$

Every displayed coefficient is nonzero, and the distinct ordered PBW monomials are linearly independent in (2.2). Thus this is a nonzero polynomial. It follows that $\iota(f) \neq 0$. Inverting p proves injectivity on Λ.

Each E_λ has a lattice stable under the integral Chevalley Lie algebra, and the scalar character χ^j has its evident integral lattice. On their tensor product the U_0-action extends continuously to A_0. For $x \in p\mathfrak{a}$, the action of $\exp(x)$ is the exponential of the infinitesimal action. It therefore agrees with the algebraic group action. The agreement extends from group elements to Λ. □

Remark 2.3. The proof uses injectivity of $\Lambda \to A$, not faithful flatness. The latter is not a property of this particular completion in general. No assertion that a nonzero module survives this completion will be needed. Only finitely many nonzero annihilating elements are transferred.

2.3 Integral image orders for a rectangular family

Fix an ordered Chevalley basis of $\mathfrak{s}_{\mathbb{Z}_p}$: root vectors together with an integral toral basis. For $L, Q \geq 0$, let $F_{L,Q}U_0$ be the span of its PBW monomials with total degree at most L in $\mathfrak{s}_{\mathbb{Z}_p}$ and degree at most Q in z. This is a bifiltration:

$$(F_{L,Q}U_0)(F_{L',Q'}U_0) \subset F_{L+L',Q+Q'}U_0.$$

For every pair (λ, j) in the window $\|\lambda\| \leq K,\ |j| \leq N$, choose a $\mathfrak{a}$-stable lattice $\mathcal{E}_{\lambda,j}$ in $E_\lambda \otimes \chi^j$. Define the actual integral image

$$B_{K,N} = \text{image}\left(U_0 \longrightarrow \bigoplus_{\substack{\|\lambda\| \leq K \\ |j| \leq N}} \text{End}_{\mathbb{Z}_p}(\mathcal{E}_{\lambda,j}) \right).$$

Lemma 2.4 (Integral PBW cutoff)**.** *There is a constant c_g, independent of K, N, such that*

$$F_{L,Q}U_0 \twoheadrightarrow B_{K,N}, \qquad L = \lceil c_g(K+1) \rceil, \quad Q = 2N. \tag{2.5}$$

Moreover $B_{K,N}$ is a finite free $\mathbb{Z}_p$-module, the image of A_0 is also $B_{K,N}$, and

$$B_{K,N}[1/p] = \bigoplus_{\substack{\|\lambda\| \leq K \\ |j| \leq N}} \text{End}_{\mathbb{Q}_p}(E_\lambda \otimes \chi^j). \tag{2.6}$$

Proof. The weights of E_λ lie in the convex hull of the Weyl orbit of λ. Thus their values on each of the finitely many chosen toral generators are integers of absolute value at most $c(K+1)$, where c depends only on g and the chosen weight norm. Every such toral generator h consequently satisfies on the entire family the monic integral relation

$$\prod_{\substack{a \in \mathbb{Z} \\ |a| \leq \lceil c(K+1) \rceil}} (h - a) = 0.$$

The same weight bound bounds the length of every root string. Each root vector therefore satisfies a nilpotence relation $x^r = 0$, with $r \leq c'(K+1)$, uniformly in the family. Finally, z satisfies the monic relation

$$\prod_{j=-N}^{N} (z - j) = 0.$$

These are simultaneous identities on all the lattices in the family. In an ordered PBW monomial, replace every excessive power of a single generator by its monic relation. All coefficients remain integral, and all remaining factors stay in their original order. The resulting exponents in each of the D symplectic generators are $O_g(K+1)$, and the exponent of z is at most $2N$. This proves (2.5) after increasing c_g.

The image $B_{K,N}$ is a submodule of a finite free $\mathbb{Z}_p$-module and is itself finite free. It is closed in the ambient finite matrix algebra and is complete for its p-adic topology. Since U_0 is dense in A_0, continuity shows that the latter has precisely the same image.

Over $\mathbb{Q}_p$, the representations $E_\lambda \otimes \chi^j$ in the window are pairwise nonisomorphic absolutely irreducible modules for $U_{\mathbb{Q}_p}(\mathfrak{a}_{\mathbb{Q}_p})$. Their endomorphism rings as modules are $\mathbb{Q}_p$. The finite-dimensional form of the density theorem gives the full matrix algebra on each one. Their annihilators are distinct maximal two-sided ideals: if two such annihilators agreed, the corresponding full matrix algebra would have two nonisomorphic simple modules, which is impossible. The Chinese remainder theorem for these finitely many comaximal ideals gives their full product. This proves (2.6). □

Remark 2.5. It is essential to use $B_{K,N}$ itself. It need not be saturated in the displayed integral matrix algebra. In particular, $B_{K,N}/pB_{K,N}$ must not be identified with the image of the enveloping algebra on $\bigoplus_{\lambda,j} \mathcal{E}_{\lambda,j}/p\mathcal{E}_{\lambda,j}$. The rational image statement (2.6) does not justify such an identification.

2.4 A mixed PBW quotient bound

Write

$$\overline{U} = U_{\mathbb{F}_p}(\mathfrak{s}_{\mathbb{F}_p})[z]$$

and equip it with the bifiltration just defined. PBW gives

$$\dim_{\mathbb{F}_p} F_{L,Q}\overline{U} = (Q+1)\binom{L+D}{D}. \tag{2.7}$$

For a quotient of $\overline{U}$, the notation $F_{L,Q}$ will mean the image of this subspace.

Lemma 2.6. *For a fixed nonzero $\eta \in \overline{U}$, one has*

$$\dim_{\mathbb{F}_p} F_{L,Q}(\overline{U}/\overline{U}\eta) \ll_\eta (Q+1)(L+1)^{D-1} + (L+1)^D. \tag{2.8}$$

Proof. Let a be the largest symplectic PBW degree of a monomial occurring in η, and b its largest z-degree. These maxima need not occur in the same monomial. The ordinary PBW associated graded algebra of $\overline{U}$ is a polynomial ring. Hence $\overline{U}$ is a domain, including in characteristic p.

For $L \geq a,\ Q \geq b$, right multiplication by η therefore injects

$$F_{L-a,Q-b}\overline{U} \longrightarrow F_{L,Q}\overline{U}.$$

Its image is killed in the quotient. Consequently

$$\begin{aligned}
\dim F_{L,Q}(\overline{U}/\overline{U}\eta) &\leq (Q+1)\binom{L+D}{D} - (Q-b+1)\binom{L-a+D}{D} \\
&= (Q+1)\left(\binom{L+D}{D} - \binom{L-a+D}{D}\right) + b\binom{L-a+D}{D} \\
&\ll_\eta (Q+1)(L+1)^{D-1} + (L+1)^D.
\end{aligned} \tag{2.9}$$

The cases $L < a$ or $Q < b$ follow from (2.7), after increasing the implied constant. For example, if $Q < b$, then Q is bounded in terms of η, and the second term on the right suffices. □

2.5 Proof of the aggregate estimate

Proof of Proposition 2.1. The Iwasawa algebra Λ of a uniform group is a noetherian Ore domain. If $m_1, \dots, m_r$ generate its torsion module M, each m_i is killed by some nonzero $\delta_i \in \Lambda$. Hence

$$\bigoplus_{i=1}^{r} \Lambda/\Lambda\delta_i \twoheadrightarrow M, \qquad \dim \mathrm{Hom}_\Lambda(M, W) \leq \sum_{i=1}^{r} \dim \ker(\delta_i \mid W).$$

It is enough to prove the assertion for one fixed nonzero $\delta \in \Lambda$.

By Lemma 2.2, $\iota(\delta) \neq 0$ in A. Multiplying by a nonzero scalar does not change any of its kernels. We may therefore normalize its image so that

$$\Delta \in A_0 \setminus pA_0.$$

Its reduction

$$\overline{\Delta} \in A_0/pA_0 = \overline{U}$$

is a nonzero finite PBW polynomial. This normalization and this polynomial are fixed independently of K, N.

Let $B = B_{K,N}$. Lemma 2.4 shows that Δ has an image $\Delta_B \in B$. Form the finite $\mathbb{Z}_p$-module

$$C = B/B\Delta_B.$$

Over $\mathbb{Q}_p$, (2.6) identifies B with a sum of full matrix algebras. For a d-dimensional vector space W, right multiplication by an operator T on $\mathrm{End}(W)$ has image of dimension $d \operatorname{rank} T$; equivalently, its cokernel has dimension $d \dim \ker T$. It follows that

$$\operatorname{rank}_{\mathbb{Z}_p} C = \sum_{\substack{\|\lambda\| \le K \\ |j| \le N}} e_\lambda \dim \ker\bigl(\delta : E_\lambda \otimes \chi^j \longrightarrow E_\lambda \otimes \chi^j\bigr). \tag{2.10}$$

The elementary structure theorem for a finite module over $\mathbb{Z}_p$ gives

$$\operatorname{rank}_{\mathbb{Z}_p} C \le \dim_{\mathbb{F}_p} C/pC.$$

Reduction of $A_0 \to B$ gives a surjection

$$\overline{U}/\overline{U}\,\overline{\Delta} \twoheadrightarrow B/(B\Delta_B + pB) = C/pC.$$

By (2.5), its restriction to the image of $F_{L,Q}\overline{U}$, with $L = \lceil c_g(K+1) \rceil$, $Q = 2N$, is still surjective. Applying Lemma 2.6 yields

$$\begin{aligned} \operatorname{rank}_{\mathbb{Z}_p} C &\le \dim_{\mathbb{F}_p} F_{L,Q}(\overline{U}/\overline{U}\,\overline{\Delta}) \\ &\ll_\delta (N+1)(K+1)^{D-1} + (K+1)^D. \end{aligned}$$

Together with (2.10), this proves the required bound for a fixed annihilator. Summing over the finite generating set proves (2.1). □

Remark 2.7. The two terms in (2.1) have separate roles. In the reduction $\overline{\Delta}$, a positive degree in the symplectic variables saves one power of K, while a relation only in the scalar variable saves a power of N. No nonvanishing assertion after each individual central specialization is used. Nor is a common two-sided annihilator of M required.

2.6 Transfer from a closed subgroup

We now explain why the actual Galois image need not itself be reductive or open in $\mathcal{G}$.

Proposition 2.8 (Closed-subgroup transfer)**.** *Let $\Gamma \subset \mathcal{G}$ be a closed uniform subgroup, and write $\Lambda_\Gamma = \mathbb{Z}_p[[\Gamma]][1/p]$. Let M be a finitely generated torsion left Λ_Γ-module. There is a fixed finitely generated torsion left Λ-module P such that every finite-dimensional Λ-module W satisfies*

$$\dim \mathrm{Hom}_{\Lambda_\Gamma}(M, W) \le \dim \mathrm{Hom}_\Lambda(P, W). \tag{2.11}$$

In particular, the left side of (2.1), with $\mathrm{Hom}_{\Lambda_\Gamma}(M, \cdot)$ in place of $\mathrm{Hom}_\Lambda(M, \cdot)$, satisfies the same bound with a constant depending on M and the fixed inclusion.

Proof. First there is an injection

$$\Lambda_\Gamma \hookrightarrow \Lambda. \tag{2.12}$$

For an open normal subgroup U of $\mathcal{G}$, the natural map

$$\mathbb{Z}_p[\Gamma/(\Gamma \cap U)] \longrightarrow \mathbb{Z}_p[\mathcal{G}/U]$$

is injective, since its basis is a subset of the target basis. As Γ is closed, the intersections $\Gamma \cap U$ induce its profinite topology. Taking inverse limits proves the integral version of (2.12); inverting p preserves injectivity.

Choose generators $m_1, \dots, m_r$ of M, with nonzero annihilators $\delta_i \in \Lambda_\Gamma$. Their images in Λ remain nonzero by (2.12). Put

$$P = \bigoplus_{i=1}^{r} \Lambda/\Lambda\delta_i.$$

This is finitely generated and torsion over Λ. Indeed, each nonzero δ_i becomes invertible in its Ore skew field, so each displayed cyclic quotient has rank zero.

Evaluation on the generators is an injection

$$\mathrm{Hom}_{\Lambda_\Gamma}(M, W) \hookrightarrow \bigoplus_{i=1}^{r} \ker(\delta_i \mid W).$$

The direct sum on the right is $\mathrm{Hom}_\Lambda(P, W)$, proving (2.11). Proposition 2.1 applied to P gives the final assertion. □

Remark 2.9 (Ordinary induction). One may equivalently use $\Lambda \otimes_{\Lambda_\Gamma} M$. This algebraic tensor product is a quotient of P, and hence is finitely generated torsion. Ordinary adjunction gives

$$\mathrm{Hom}_\Lambda(\Lambda \otimes_{\Lambda_\Gamma} M, W) \simeq \mathrm{Hom}_{\Lambda_\Gamma}(M, W).$$

The proof above avoids introducing a topology on this induced module. It requires neither a completed induction functor nor a flatness assertion for the subgroup inclusion.

Remark 2.10 (Application to the graded surface Lie algebra). For a symplectic space H, the relation line $\mathbb{Q}_p\omega \subset \bigwedge^2 H$ is stable under $\mathrm{GSp}(H)$ and transforms by its multiplier ν. Consequently

$$\mathrm{Lie}(H)/(\omega)$$

and all of its homogeneous pieces have canonical algebraic $\mathrm{GSp}(H)$-actions. The representations $V_n^* \otimes \nu$ therefore extend from the actual Galois image to the ambient group used here. Their scalar character is χ^{2-n}. Proposition 2.8 transfers the requisite Hom bounds without asserting that the actual Galois representation is semisimple. The arithmetic module remains defined over its actual Galois image; only its finitely many annihilators are used in the ambient group algebra.

3 Characters and concentration for the symplectic surface Lie algebra

This section supplies the representation-theoretic estimates used in the proposed argument. They concern the standard representation of the abstract algebraic group Sp_{2g}, with $g \geq 2$. In particular, using these estimates for an arithmetic representation requires no assertion

that its Galois image is reductive. The connection with the graded Lie algebra of a surface group is classical; the algebra calculation needed here is included in full.[1]

Let H be the standard $2g$-dimensional symplectic vector space in characteristic zero. Choose a symplectic basis $x_1, y_1, \dots, x_g, y_g$, and put

$$L = \mathrm{Lie}(H)/(\omega), \qquad \omega = \sum_{i=1}^{g} [x_i, y_i], \qquad V_n = L_n, \qquad d_n = \dim V_n. \tag{3.1}$$

Here $\mathrm{Lie}(H)$ denotes the free graded Lie algebra with H in degree one, and (ω) denotes the Lie ideal generated by ω. Throughout this section set

$$S = \mathrm{Sp}_{2g}, \qquad D = \dim S = g(2g+1), \qquad \lambda = g + \sqrt{g^2 - 1} > 1.$$

Over a splitting field, write

$$V_n = \bigoplus_{\alpha} E_\alpha^{\oplus m_{n,\alpha}},$$

where E_α is the irreducible representation of dominant highest weight α. Constants in this section may depend on g, but are independent of n, α, and the cutoff K.

3.1 The equivariant Hilbert series and the Witt formula

Lemma 3.1 (The quadratic resolution)**.** *Let $A = U(L)$. There is an exact sequence of graded, S-equivariant left A-modules*

$$0 \longrightarrow A(-2) \xrightarrow{d_2} A(-1) \otimes H \xrightarrow{d_1} A \xrightarrow{\varepsilon} \mathbb{C} \longrightarrow 0. \tag{3.2}$$

Consequently, for every $s \in S(\mathbb{C})$,

$$\sum_{n \geq 0} \mathrm{tr}(s \mid A_n) t^n = \frac{1}{1 - \chi_H(s) t + t^2}. \tag{3.3}$$

The same statements hold over every field of characteristic zero.

Proof. The universal property of the enveloping algebra gives

$$A = T(H) \Big/ \left(\sum_{i=1}^{g} (x_i y_i - y_i x_i) \right).$$

Order words first by length and then lexicographically, with x_1 larger than every other letter. Orient the relation as the reduction

$$x_1 y_1 \longmapsto y_1 x_1 - \sum_{i=2}^{g} (x_i y_i - y_i x_i).$$

Every word on the right is smaller than $x_1 y_1$. Thus reductions terminate in each fixed degree. The leading word $x_1 y_1$ has no proper self-overlap, since its last letter differs from its first. Two simultaneous reducible occurrences therefore occupy disjoint pairs of positions; reducing those two occurrences in either order gives the same linear combination. Induction in the word order now gives confluence. Equivalently, this is the elementary one-relation instance

[1] For the classical one-relator identification, see J. P. Labute, *On the descending central series of groups with a single defining relation*, Journal of Algebra **14** (1970), 16–23, doi:10.1016/0021-8693(70)90130-4. The character and concentration calculations below do not require a further theorem from that paper.

of the noncommutative diamond lemma. The words not containing x_1y_1 consequently form a basis of A.

Let $a_n = \dim A_n$, with $a_0 = 1$ and $a_1 = 2g$. Appending a letter to a normal word produces every normal word of the next degree, except that one cannot append y_1 to a word ending in x_1. Appending x_1 is always allowed. Thus there are exactly a_{n-2} normal words of degree $n-1$ ending in x_1, and

$$a_n = 2ga_{n-1} - a_{n-2} \qquad (n \geq 2).$$

Define

$$d_1(a \otimes h) = ah, \qquad d_2(a) = \sum_{i=1}^{g} (ax_i \otimes y_i - ay_i \otimes x_i).$$

The composite d_1d_2 vanishes. The coefficient of y_1 in $d_2(a)$ is ax_1. Right multiplication by x_1 takes distinct normal words to distinct normal words, so it is injective. Hence d_2 is injective. Multiplication $A_{n-1} \otimes H \longrightarrow A_n$ is surjective for $n \geq 1$. Its kernel has dimension $2ga_{n-1} - a_n = a_{n-2}$ for $n \geq 2$, exactly the dimension of the image of d_2. In degree one the kernel is zero. This proves exactness.

The tensor $\sum_i (x_i \otimes y_i - y_i \otimes x_i)$ is S-invariant. The maps above, including the copy of the trivial representation in degree two, are therefore S-equivariant. Taking graded traces in the exact sequence proves (3.3). The normal-word and exactness arguments work over the original characteristic-zero field as well. For GSp_{2g}, the same resolution is equivariant after replacing its first term by $A(-2) \otimes \nu$, where ν is the multiplier character. Thus for a similitude s the denominator becomes $1 - \chi_H(s)t + \nu(s)t^2$. We will use this form for complex conjugation. □

For an indeterminate x, define polynomials $B_n(x)$ by

$$B_0(x) = 2, \qquad B_1(x) = x, \qquad B_n(x) = xB_{n-1}(x) - B_{n-2}(x).$$

If $a + a^{-1} = x$, then $B_n(x) = a^n + a^{-n}$, including at repeated roots. Write $B_n(s) = B_n(\chi_H(s))$. This notation does not involve choosing one of the two roots.

Lemma 3.2 (Equivariant Witt identities)**.** *For $n \geq 1$ and $s \in S(\mathbb{C})$ one has*

$$B_n(s) = \sum_{d|n} d\, \chi_{V_d}(s^{n/d}), \tag{3.4}$$

$$n\chi_{V_n}(s) = \sum_{r|n} \mu(r) B_{n/r}(s^r), \tag{3.5}$$

where μ is the Möbius function. In particular,

$$\begin{aligned} d_n &= \frac{1}{n} \sum_{r|n} \mu(r) \left(\lambda^{n/r} + \lambda^{-n/r} \right) \\ &= \frac{\lambda^n}{n} + O_g\left(\frac{\tau(n)\lambda^{n/2}}{n} \right), \end{aligned} \tag{3.6}$$

where $\tau(n)$ is the number of positive divisors of n. There are positive constants c_g, C_g such that

$$c_g \frac{\lambda^n}{n} \leq d_n \leq C_g \frac{\lambda^n}{n} \qquad (n \geq 1). \tag{3.7}$$

Proof. The equivariant Poincaré–Birkhoff–Witt symmetrization identifies $U(L)$ with $\mathrm{Sym}(L)$ as graded S-representations. Taking the formal logarithm of its graded character gives

$$\log\left(\sum_{n\geq 0}\operatorname{tr}(s \mid A_n)t^n\right) = \sum_{k\geq 1}\frac{1}{k}\sum_{d\geq 1}\chi_{V_d}(s^k)t^{dk}.$$

Each coefficient involves only finitely many terms. On the other hand, (3.3) and the recurrence defining B_n give

$$-\log(1-\chi_H(s)t+t^2) = \sum_{n\geq 1}\frac{B_n(s)}{n}t^n.$$

Comparison of coefficients proves (3.4). For completeness, applying the sum on the right of (3.5) to (3.4) yields

$$\sum_{d\mid n} d\,\chi_{V_d}(s^{n/d})\sum_{r\mid n/d}\mu(r) = n\chi_{V_n}(s),$$

which proves the inversion formula with the power operations in the correct places.

At $s=1$, the two roots are λ and λ^{-1}. Each proper divisor term, indexed by $r\geq 2$, has absolute value at most $2\lambda^{n/2}$. The remaining term λ^{-n} is absorbed by the displayed error. This proves (3.6). The error is $o(\lambda^n/n)$, since $\tau(n)\leq n$. Also $d_n>0$ in every degree: sending x_1, x_2 to two free Lie generators and every y_i and remaining x_i to zero gives a surjection from L onto the free Lie algebra on two generators. The element $\mathrm{ad}(x_1)^{n-1}x_2$ there is nonzero, as its tensor expansion contains $x_1^{n-1}x_2$ with coefficient one. Enlarging the constants for the finitely many small degrees proves (3.7). □

3.2 A compact-group estimate in the full group dimension

Let $S_c = \mathrm{USp}(2g)$, equipped with Haar probability measure. It has real dimension $D = g(2g+1)$. Its defining matrices are unitary, and their eigenvalues occur in reciprocal pairs; hence χ_H is real-valued on S_c.

Lemma 3.3 (Isolated spectral maxima and a Gaussian integral)**.** *There is a constant C_g such that, for every $n\geq 1$,*

$$\int_{S_c}|B_n(s)|\,ds \leq C_g\lambda^n n^{-D/2}. \tag{3.8}$$

Consequently,

$$\int_{S_c}|\chi_{V_n}(s)|\,ds \leq C_g d_n n^{-D/2}. \tag{3.9}$$

Proof. Since all eigenvalues of s have modulus one, $|\chi_H(s)|\leq 2g$. Equality in the triangle inequality forces all eigenvalues to be the same. A unitary matrix is diagonalizable, so s is scalar. A scalar preserving the symplectic form is 1 or -1. The equality set is therefore exactly $\{I,-I\}$.

On the compact Lie algebra $\mathfrak{s}_c$ use the positive definite norm

$$\|X\|^2 = -\operatorname{tr}(X^2).$$

It is positive definite because X is skew-Hermitian; the defining representation is faithful. Also $\operatorname{tr}(X)=0$. Taylor expansion in exponential coordinates gives

$$\chi_H(\exp X) = 2g - \tfrac{1}{2}\|X\|^2 + O_g(\|X\|^3).$$

For sufficiently small $\|X\|$, this trace is positive, and hence

$$\tfrac{1}{4}\|X\|^2 \le 2g - |\chi_H(\pm \exp X)| \le \tfrac{3}{4}\|X\|^2. \tag{3.10}$$

This verifies the nondegenerate quadratic maximum in all D real tangent directions of the group. Exponential coordinates are coordinates on the group near the two central points; no quotient by conjugation is taken.

For real x, $|x| \le 2g$, put

$$r(x) = \begin{cases} 1, & |x| \le 2, \\ (|x| + \sqrt{x^2-4})/2, & 2 < |x| \le 2g. \end{cases}$$

The two roots of $a + a^{-1} = x$ have modulus at most $r(x)$, so $|B_n(x)| \le 2r(x)^n$. This bound includes $x = \pm 2$. The function r is continuous, $r(\pm 2g) = \lambda$, and near $y = 2g$

$$\frac{d}{dy}\log\left(\frac{y+\sqrt{y^2-4}}{2}\right) = \frac{1}{\sqrt{y^2-4}} > 0.$$

It follows from (3.10), after shrinking the two coordinate neighborhoods, that

$$r(\chi_H(\pm \exp X)) \le \lambda \exp(-c_g\|X\|^2)$$

for a fixed $c_g > 0$. The smooth Haar densities in those coordinates are bounded above. Integration there is consequently bounded by

$$C_g\lambda^n \int_{\mathbb{R}^D} \exp(-c_g n\|X\|^2)\, dX \le C_g\lambda^n n^{-D/2}.$$

On the complement of these neighborhoods, compactness and the isolated equality set give $r(\chi_H(s)) \le \rho_g < \lambda$. Its contribution is at most $2\rho_g^n$, which is absorbed by the same bound. This proves (3.8).

For the proper-divisor terms in (3.5), it is enough to use the pointwise estimate $|B_{n/r}(s^r)| \le 2\lambda^{n/r} \le 2\lambda^{n/2}$ for $r \ge 2$. There is no need to study the possibly nonisolated preimages of $\{I, -I\}$ under $s \mapsto s^r$. Thus

$$\int_{S_c} |\chi_{V_n}(s)|\, ds \le C_g \frac{\lambda^n}{n} n^{-D/2} + \frac{2\tau(n)}{n}\lambda^{n/2}.$$

The second term is absorbed into the first uniformly in n, since $\tau(n)n^{D/2}\lambda^{-n/2}$ is bounded. Finally use (3.7). □

Proposition 3.4 (Uniform multiplicity bound)**.** *For all $n \ge 1$ and all dominant weights α,*

$$\frac{m_{n,\alpha}}{d_n \dim E_\alpha} \le C_g n^{-D/2}, \qquad D = g(2g+1). \tag{3.11}$$

Proof. Character orthogonality on S_c gives

$$m_{n,\alpha} = \int_{S_c} \chi_{V_n}(s)\overline{\chi_{E_\alpha}(s)}\, ds.$$

The restriction of E_α to S_c is unitary after choosing an invariant Hermitian form, so $|\chi_{E_\alpha}(s)| \le \dim E_\alpha$. Apply (3.9). □

3.3 A Gaussian tail for highest weights

Use the standard type-C_g coordinates, in which the weights of H are $\pm e_1, \dots, \pm e_g$. A dominant highest weight is an integer vector

$$\alpha = (\alpha_1, \dots, \alpha_g), \qquad \alpha_1 \geq \cdots \geq \alpha_g \geq 0.$$

In the following statements $\|\alpha\|$ is the Euclidean norm in these coordinates. Changing this fixed norm only changes the constants.

Proposition 3.5 (Dimension-weighted highest-weight tail)**.** *For every $n \geq 1$ and every real $K \geq 0$,*

$$\sum_{\|\alpha\|>K} \frac{m_{n,\alpha} \dim E_\alpha}{d_n} \leq C_g (1+n)^{g^2} \exp\left(-\frac{K^2}{4gn}\right). \tag{3.12}$$

The sum is zero when $K \geq n$. Thus the Gaussian assertion uses a moment parameter in a fixed bounded interval throughout its entire nonempty range.

Proof. The free Lie algebra embeds into the tensor algebra, and V_n is a quotient of its degree-n part. Every weight of V_n is therefore a weight of $H^{\otimes n}$. Such a weight is a sum of n vectors from $\{\pm e_i\}$ and has ℓ^1 norm at most n. In particular, a highest weight occurring in V_n satisfies

$$\sum_{i=1}^{g} \alpha_i \leq n, \qquad \|\alpha\| \leq n, \qquad \alpha_1 \geq \|\alpha\|/\sqrt{g}. \tag{3.13}$$

Weyl's dimension formula in these coordinates is

$$\dim E_\alpha = \prod_{i=1}^{g} \frac{\alpha_i + g - i + 1}{g - i + 1} \cdot \prod_{1 \leq i < j \leq g} \frac{(\alpha_i - \alpha_j + j - i)(\alpha_i + \alpha_j + 2g - i - j + 2)}{(j-i)(2g - i - j + 2)}.$$

It contains g^2 linear numerator factors. By (3.13), each is bounded by a constant depending on g times $1 + n$. Hence, for every constituent of V_n,

$$\dim E_\alpha \leq C_g (1+n)^{g^2}. \tag{3.14}$$

Let s_u be the positive real diagonal torus element acting on x_1, y_1 by e^u, e^{-u}, and trivially on the other basis vectors. Then

$$\chi_H(s_u) = 2g - 2 + 2\cosh u.$$

Every weight eigenvalue of s_u in a rational representation is positive. All terms in the non-inverted identity (3.4) are therefore nonnegative. Its term with $d = n$ gives

$$n\chi_{V_n}(s_u) \leq B_n(s_u).$$

The larger root is $a(u) > 1$, with

$$\log a(u) = \operatorname{arcosh}(g - 1 + \cosh u), \qquad a(0) = \lambda.$$

For $0 \leq u \leq 1$, direct differentiation gives

$$\frac{d}{du} \log a(u) = \frac{\sinh u}{\sqrt{(g - 1 + \cosh u)^2 - 1}} \leq \frac{u \cosh 1}{\sqrt{g^2 - 1}}.$$

Since $g \geq 2$, integration shows

$$0 \leq \log a(u) - \log \lambda \leq \frac{\cosh 1}{2\sqrt{g^2-1}} u^2 \leq u^2.$$

Together with $B_n(s_u) = a(u)^n + a(u)^{-n} \leq 2a(u)^n$ and (3.7), this proves the positive moment bound

$$\frac{\chi_{V_n}(s_u)}{d_n} \leq C_g \exp(nu^2) \qquad (0 \leq u \leq 1). \tag{3.15}$$

Using the non-inverted identity here avoids any cancellation from the Möbius function.

Suppose first that $0 \leq K \leq n$. For every highest weight in the tail, its highest-weight vector contributes the positive term $e^{u\alpha_1}$ to $\chi_{E_\alpha}(s_u)$. By (3.13), this term is at least $e^{uK/\sqrt{g}}$. Positivity of all other weight terms gives

$$\sum_{\|\alpha\|>K} m_{n,\alpha} \leq e^{-uK/\sqrt{g}} \chi_{V_n}(s_u).$$

Multiplying by the uniform dimension bound (3.14) and applying (3.15) therefore yields

$$\sum_{\|\alpha\|>K} \frac{m_{n,\alpha} \dim E_\alpha}{d_n} \leq C_g (1+n)^{g^2} \exp(nu^2 - uK/\sqrt{g}).$$

Choose $u = K/(2\sqrt{g}\, n)$. Because $K \leq n$, this lies in $[0, 1/(2\sqrt{g})] \subset [0,1]$, exactly the interval where the moment bound has been proved. The exponent becomes $-K^2/(4gn)$. If $K > n$, the sum is zero by (3.13); this also proves the stated inequality in that range. □

Corollary 3.6 (A dyadic cutoff)**.** *For integers $N \geq 2$, put*

$$K_N = \left\lceil \sqrt{N} \log N \right\rceil.$$

Uniformly for $N \leq n < 2N$,

$$\frac{m_{n,\alpha}}{d_n} \leq C_g N^{-D/2} \dim E_\alpha, \tag{3.16}$$

$$\sum_{\|\alpha\|>K_N} \frac{m_{n,\alpha} \dim E_\alpha}{d_n} \leq C_g N^{g^2} \exp\left(-\frac{(\log N)^2}{8g}\right). \tag{3.17}$$

In particular, the second expression is $o(N^{-A})$ for each fixed $A > 0$.

Proof. The first bound follows from Proposition 3.4 and $n \geq N$. For the second, use $(1+n)^{g^2} \leq (2N)^{g^2}$ and $K_N^2/n \geq (\log N)^2/2$ in Proposition 3.5. Finally, for every fixed A, the quadratic negative term $-(\log N)^2/(8g)$ dominates $(A+g^2)\log N$. □

Remark 3.7 (Duals, scalar weights, and the field of definition). The longest Weyl element of type C_g is -1, so $E_\alpha^* \simeq E_\alpha$. Thus V_n^* has precisely the same highest-weight multiplicities and the same estimates. The scalar homothety aI in GSp_{2g} acts on V_n by a^n. Writing ν for the multiplier character, it acts on $V_n^* \otimes \nu$ by a^{2-n}. This is the representation denoted $V_n^*(1)$ when the symplectic pairing has values in the Tate line.

All multiplicities, dimensions, and character identities above are identities for the split symplectic root datum in characteristic zero. One may therefore compute them on the complex group and apply the resulting numerical estimates to the corresponding algebraic representations over a p-adic coefficient field. This does not identify a p-adic compact group with S_c, and it does not invoke semisimplicity of a Galois representation. The separate passage from these estimates to Iwasawa-module bounds remains part of the proposed argument and is not a consequence of this section alone.

4 Averaging over depth

We now combine the two estimates before returning to arithmetic. Let

$$S = \mathrm{Sp}_{2g}, \qquad D = g(2g+1),$$

and write $\Gamma_0 = \mathcal{G}$ for the small ambient scalar–symplectic product from Section 2. Write $R_0 = \mathbb{Z}_p[[\Gamma_0]][1/p]$. Let P be a fixed finitely generated torsion R_0-module with r generators. For a dominant symplectic weight α and an integer j, put

$$a(\alpha, j) = \dim_{\mathbb{Q}_p} \mathrm{Hom}_{R_0}(P, E_\alpha \otimes \chi^j).$$

Evaluation on a fixed set of generators gives

$$a(\alpha, j) \le r \dim E_\alpha. \tag{4.1}$$

The relation line has scalar weight two, so the degree-n dual-twisted surface representation is

$$W_n = \bigoplus_\alpha (E_\alpha^* \otimes \chi^{2-n})^{m_{n,\alpha}}.$$

Use the Euclidean highest-weight norm of Section 3. It is equivalent to the fundamental-weight norm in Proposition 2.1; rescaling the cutoff by a fixed constant leaves that estimate unchanged. In the symplectic case all the irreducible representations are self-dual, but writing the dual explicitly keeps track of its origin.

Proposition 4.1 (Proposed averaged multiplicity consequence)**.** *The aggregate estimate and the surface-Lie character estimates imply*

$$\frac{1}{N} \sum_{N \le n < 2N} \frac{\dim \mathrm{Hom}_{R_0}(P, W_n)}{d_n} \longrightarrow 0. \tag{4.2}$$

If $\Gamma \subset \Gamma_0$ is closed and uniform and M is a fixed finitely generated torsion R_Γ-module, the transfer of annihilators gives

$$\frac{1}{N} \sum_{N \le n < 2N} \frac{\dim \mathrm{Hom}_{R_\Gamma}(M, W_n)}{d_n} \longrightarrow 0. \tag{4.3}$$

Proof. Take $K = \lceil \sqrt{N} \log N \rceil$. The low-weight part of the first average is at most a constant times

$$\begin{aligned} & N^{-1-D/2} \sum_{N \le n < 2N} \sum_{\|\alpha\| \le K} \dim E_\alpha \, a(\alpha^*, 2-n) \\ & \le N^{-1-D/2} \sum_{\|\alpha\| \le K} \sum_{|j| \le 2N} \dim E_\alpha \, a(\alpha^*, j) \\ & \ll_P N^{-1-D/2}(NK^{D-1} + K^D) \\ & \ll_P N^{-1/2}(\log N)^{D-1} + N^{-1}(\log N)^D. \end{aligned}$$

The first inequality uses the uniform multiplicity bound $m_{n,\alpha}/d_n \ll n^{-D/2} \dim E_\alpha$. The next uses the fact that $n \mapsto 2-n$ is injective: a pair (α, j) is counted at most once in the depth interval. Thus there is no further factor of N.

By (4.1), the high-weight part at degree n is at most

$$r \sum_{\|\alpha\|>K} \frac{m_{n,\alpha} \dim E_\alpha}{d_n} \ll N^A \exp(-c'(\log N)^2),$$

uniformly for $N \le n < 2N$. This tends to zero, proving (4.2).

For the actual module M, choose generators and nonzero annihilators δ_i, and use the fixed ambient torsion module $P = \bigoplus_i R_0/R_0\delta_i$. The injection of Iwasawa algebras and evaluation on generators give

$$\dim \mathrm{Hom}_{R_\Gamma}(M, W_n) \le \dim \mathrm{Hom}_{R_0}(P, W_n).$$

This proves (4.3). □

Remark 4.2 (Why averaging is essential). Torsion over the full Iwasawa algebra does not imply torsion after every scalar specialization. For example, a module killed by $z - 1$ for a fixed scalar group element z may specialize to a free module in the symplectic directions at the trivial scalar character. The proof sums over a long interval of distinct scalar characters; it makes no assertion that every specialization is torsion and no pointwise little-o assertion in depth.

5 The arithmetic reduction over the rational numbers

We now give the arithmetic deduction from the averaged estimate (4.3). In particular, the deduction in this section has that estimate as an explicit hypothesis. We use the actual division-field Galois group and the ambient transfer from Proposition 2.8. No semisimplicity theorem for the Tate-module Galois representation is used.

5.1 The actual Galois image and its ambient group

Let $C/\mathbb{Q}$ be a smooth, projective, geometrically connected curve of genus $g \ge 2$, and let J be its Jacobian. For a finiteness assertion the case $C(\mathbb{Q}) = \varnothing$ needs no argument. Otherwise fix $b \in C(\mathbb{Q})$. Choose a prime $p \ge 5$ of good reduction for C. Write

$$H = V_p J, \qquad U = \pi_1^{\mathrm{un}}(C_{\overline{\mathbb{Q}}}, b), \qquad U^1 = U, \quad U^{n+1} = [U, U^n], \quad U_N = U/U^{N+1},$$

and put

$$V_n = \mathrm{Lie}(U^n/U^{n+1}), \qquad W_n = V_n^\vee(1), \qquad d_n = \dim_{\mathbb{Q}_p} V_n.$$

All dimensions in this section are over $\mathbb{Q}_p$, unless another field is indicated. The canonical principal polarization gives a nondegenerate alternating pairing $H \otimes H \longrightarrow \mathbb{Q}_p(1)$. The multiplier character of the resulting Galois representation

$$\rho : G_{\mathbb{Q}} \longrightarrow \mathrm{GSp}(H)$$

is the cyclotomic character.

The equivariant surface-Lie presentation 3.1 identifies $\bigoplus_n V_n$ with the free Lie algebra on H modulo the ideal generated by the symplectic relation $\omega \in \bigwedge^2 H$; see also [10] and [8, pp. 123–125]. The line spanned by ω is stable under $\mathrm{GSp}(H)$. Consequently each V_n has a canonical algebraic $\mathrm{GSp}(H)$-action, and its actual Galois action is the restriction of this action. On the ambient group, the notation W_n means $V_n^\vee \otimes \mu$, where μ is the multiplier

character. Restriction to $G_{\mathbb{Q}}$ recovers $V_n^\vee(1)$. In particular, a scalar uI acts on V_n by u^n and on W_n by u^{2-n}.

Choose the ambient uniform compact open subgroup $G_0 = \mathcal{G}$ of Section 2.1, inside $\mathrm{GSp}(H)(\mathbb{Q}_p)$ after a symplectic choice of lattice and basis. Because p is odd, a sufficiently small such group is the product of its scalar subgroup and a principal congruence subgroup of $\mathrm{Sp}(H)$: the required scalar factor is the unique square root of the multiplier near 1.

The image of ρ is a compact p-adic analytic matrix group. The uniform-open-subgroup theorem for compact p-adic analytic groups [4] permits choosing a finite extension $F/\mathbb{Q}$ for which

$$F_\infty = F(J[p^\infty]), \qquad \Gamma = \mathrm{Gal}(F_\infty/F) = \rho(G_F)$$

is uniform and is contained in G_0. Here one first chooses a uniform open subgroup of $\rho(G_{\mathbb{Q}}) \cap G_0$ and then takes the finite extension defined by its inverse image. We do not assert that Γ is open in G_0, or that its algebraic monodromy group is reductive.

Bogomolov's homothety theorem [2] says that the Lie algebra of the original Tate-module image contains the scalar Lie algebra. It follows that the image contains an open subgroup of the scalar units near 1. Passing to an open subgroup preserves this property. Thus Γ contains a nontrivial scalar uI with $u \in 1 + p\mathbb{Z}_p$; this scalar has infinite order. Only this homothety assertion from Bogomolov is needed below.

Fix a finite set T of rational places containing p, ∞, the primes of bad reduction of C, and the primes ramifying in F. Write T_F for the places of F above T, and $T_{F,f}$ for its finite places. Let $G_{F,T}$ denote the Galois group of the maximal extension of F unramified outside T_F. The division tower is contained in this extension. These choices will remain fixed throughout the estimates in n.

5.2 The unramified module and induction

Let $\mathcal{H}/F_\infty$ be the maximal everywhere-unramified abelian pro-p extension and put

$$\mathcal{M} = \mathrm{Gal}(\mathcal{H}/F_\infty), \qquad M = \mathcal{M} \otimes_{\mathbb{Z}_p} \mathbb{Q}_p, \qquad R = \mathbb{Z}_p[[\Gamma]][1/p], \quad A = \mathbb{Z}_p[[G_0]][1/p].$$

Conjugation gives $\mathcal{M}$ its natural left Iwasawa-module structure. The algebras R and A are bounded rational Iwasawa algebras. In particular, neither notation denotes a ring of power series with arbitrary unbounded denominators.

Proposition 5.1. *The module M is finitely generated and torsion over R.*

Proof. The Weil pairing implies that F_∞ contains the cyclotomic $\mathbb{Z}_p$-extension of F. The finite-generation assertion immediately preceding Corollary 4.8 in [11, author's preprint, p. 29], and Corollary 4.8(ii) there, apply to the everywhere-unramified module $\mathcal{M}$. They give finite generation and torsion over $\mathbb{Z}_p[[\Gamma]]$. Inverting p gives the assertion. The cited corollary applies to a uniform p-adic Lie extension containing the cyclotomic extension; it does not require pseudo-nullity or finite generation over a subgroup of codimension one. □

Lemma 5.2. *The natural homomorphism $R \longrightarrow A$ is injective, and the algebraically induced module*

$$\widetilde{M} = A \otimes_R M$$

is finitely generated and torsion over A. For every finite-dimensional algebraic representation W of $\mathrm{GSp}(H)$ there is a natural identification

$$\mathrm{Hom}_A(\widetilde{M}, W) \simeq \mathrm{Hom}_R(M, W|_\Gamma). \tag{5.1}$$

Proof. For each open normal subgroup $N \subset G_0$, the inclusion $\Gamma/(\Gamma \cap N) \subset G_0/N$ gives an injective map of integral group algebras. The subgroups $\Gamma \cap N$ are cofinal among the open neighborhoods of the identity in Γ. Taking inverse limits proves injectivity of the integral Iwasawa-algebra map; inverting p proves injectivity of $R \to A$.

The Iwasawa algebras of the uniform groups in question are noetherian domains, and hence Ore domains; see [11, Section 2.1]. Choose finitely many R-generators $m_1, \dots, m_q$ of M. Since M is torsion, for each i there is a nonzero $f_i \in R$ such that $f_i m_i = 0$. The image of f_i in A is nonzero and annihilates $1 \otimes m_i$. In the skew field of fractions $Q(A)$, this image is invertible, so

$$1 \otimes (1 \otimes m_i) = f_i^{-1} \otimes f_i(1 \otimes m_i) = 0 \qquad \text{in } Q(A) \otimes_A \widetilde{M}.$$

These elements generate $\widetilde{M}$, proving that it has rank zero over A. Finite generation is immediate from the same generators. This argument requires neither flatness of A over R nor a common two-sided annihilator of M.

Equation (5.1) is ordinary extension-of-scalars adjunction. The algebraic Hom spaces here agree with the continuous Hom spaces used in Galois cohomology. Indeed, pull a module homomorphism back along a compact finite-generator presentation of $\mathcal{M}$. On the free module it has the form $(a_i) \mapsto \sum_i a_i w_i$ for fixed vectors w_i. An algebraic representation admits a stable lattice for the compact group, so its completed group-algebra action is continuous. The map on the free module, and therefore the descended map, is continuous. Inverting p does not change the Hom space with this finite-dimensional target. □

5.3 Strict localization kernels and global second cohomology

For a finite-dimensional representation W of $G_{F,T}$, define

$$\mathrm{Sha}^1_T(F, W) = \ker\left(H^1(G_{F,T}, W) \longrightarrow \bigoplus_{v \in T_{F,f}} H^1(G_{F_v}, W) \right).$$

For p odd, the positive-degree cohomology of the archimedean decomposition groups with rational p-adic coefficients is zero; omitting those places does not affect this kernel.

Lemma 5.3. *For $n \geq 3$, there is an injection*

$$\mathrm{Sha}^1_T(F, W_n) \hookrightarrow \mathrm{Hom}_R(M, W_n). \tag{5.2}$$

Proof. Take the infinite-order scalar $uI \in \Gamma$ chosen above. If $f : \Gamma \to W_n$ is a continuous 1-cocycle, centrality gives

$$f(u\gamma) = f(u) + u^{2-n} f(\gamma) = f(\gamma) + \gamma f(u) = f(\gamma u).$$

Consequently

$$(u^{2-n} - 1) f(\gamma) = (\gamma - 1) f(u).$$

For $n \geq 3$ the scalar on the left is nonzero, so f is the coboundary of $(u^{2-n} - 1)^{-1} f(u)$. Thus $H^1(\Gamma, W_n) = 0$.

Let $N = \mathrm{Gal}(F_T/F_\infty)$. The action of N on W_n is trivial, because the action on the graded surface Lie algebra factors through the action on H, and the cyclotomic character is also trivial on N. Inflation–restriction now gives an injection

$$H^1(G_{F,T}, W_n) \hookrightarrow H^1(N, W_n)^\Gamma = \mathrm{Hom}_{\mathrm{cont}}(N, W_n)^\Gamma.$$

A continuous homomorphism from the compact group N to the additive group of W_n has compact pro-p image, and therefore factors through the maximal abelian pro-p quotient of N.

If the original cohomology class is locally trivial at a place $v \in T_{F,f}$, its restriction to G_{F_v} is a coboundary. On the intersection of N with that decomposition group, this coboundary vanishes since N acts trivially on W_n. The resulting homomorphism on N thus vanishes on every decomposition group over the places in $T_{F,f}$. Outside T_F the extension already is unramified. It follows that the homomorphism factors through $\mathcal{M}$, in fact through the quotient of $\mathcal{M}$ in which the places above $T_{F,f}$ split completely. Equivariance and rationalization give (5.2). This is the strict-localization argument used in [3, preprint, pp. 12–13]; the scalar calculation above supplies its inflation–restriction hypothesis in the present setting. □

Proposition 5.4. *For $n \geq 3$, put $h_n = \dim H^2(G_{\mathbb{Q},T}, V_n)$. Then*

$$h_n \leq \dim \mathrm{Hom}_A(\widetilde{M}, W_n) + \sum_{v \in T_{F,f}} \dim H^0(G_{F_v}, W_n). \tag{5.3}$$

Proof. Poitou–Tate duality identifies the dimension of the kernel of localization on $H^2(G_{F,T}, V_n)$ with $\dim \mathrm{Sha}^1_T(F, W_n)$. Local Tate duality identifies the dimension of $H^2(G_{F_v}, V_n)$ with that of $H^0(G_{F_v}, W_n)$. These standard duality statements give

$$\dim H^2(G_{F,T}, V_n) \leq \dim \mathrm{Sha}^1_T(F, W_n) + \sum_{v \in T_{F,f}} \dim H^0(G_{F_v}, W_n).$$

The same use of global and local duality is detailed in [3, preprint, pp. 12–13]. Apply Lemma 5.3 and (5.1). Finally, restriction from $\mathbb{Q}$ to F is injective on rational p-adic cohomology: its composition with corestriction is multiplication by the nonzero integer $[F : \mathbb{Q}]$. This proves (5.3). Notice that this step changes only an upper bound on H^2; it does not change the base field in the global Euler characteristic used below. □

5.4 The local terms, including bad reduction

Lemma 5.5. *For the fixed F and T, there is a constant B such that, for $n \geq 3$,*

$$\sum_{v \in T_{F,f}} \dim H^0(G_{F_v}, W_n) \leq B n^2 g^n. \tag{5.4}$$

No semisimplicity assertion for the Tate-module representation or for a Frobenius operator is required.

Proof. The iterated bracket map is a Galois-equivariant surjection $H^{\otimes n} \twoheadrightarrow V_n$. Dualizing and twisting gives

$$W_n \hookrightarrow H^1_{\mathrm{et}}(C_{\overline{\mathbb{Q}}}, \mathbb{Q}_p)^{\otimes n}(1). \tag{5.5}$$

It suffices to bound the invariants of the representation on the right.

First take $v \nmid p$. Pass to a fixed finite extension of F_v over which J has semistable reduction; taking invariants over this extension can only increase their dimension. If t is the toric rank of the semistable reduction, Raynaud's description gives a Frobenius-stable filtration of $H^1_{\mathrm{et}}(C_{\overline{F_v}}, \mathbb{Q}_p)$ whose graded parts have weights $0, 1, 2$ and dimensions

$$t, \qquad 2g - 2t, \qquad t,$$

respectively, where $0 \le t \le g$. This is the degree-one weight description used in the local calculation of [8, Lemma 6, pp. 124–125]. We use geometric Frobenius, so that its eigenvalues on a weight-w part have complex absolute value $q^{w/2}$ and the twist (1) multiplies them by q^{-1}.

An eigenvalue equal to 1 in the tensor representation in (5.5) can consequently occur only in tensor terms of total weight 2 before twisting. These terms have either one weight-2 factor and all other factors of weight 0, or two weight-1 factors and all other factors of weight 0. Their total dimension is

$$nt^n + \binom{n}{2}(2g-2t)^2 t^{n-2} = O_g(n^2 g^n).$$

For $t = 0$ and $n \ge 3$, both terms are zero. The dimension of an invariant space is bounded by the algebraic multiplicity of the Frobenius eigenvalue 1. This argument therefore does not assume that Frobenius is semisimple.

For $v \mid p$, let H^1_p denote the degree-one étale cohomology in (5.5). Its Hodge–Tate decomposition has the form

$$H^1_p \otimes_{\mathbb{Q}_p} \mathbb{C}_p \simeq \mathbb{C}_p^{\oplus g} \oplus \mathbb{C}_p(-1)^{\oplus g}.$$

After taking the nth tensor power and twisting by (1), the multiplicity of $\mathbb{C}_p(0)$ is ng^n: exactly one factor must come from $\mathbb{C}_p(-1)^{\oplus g}$. There is an injection

$$H^0(G_{F_v}, (H^1_p)^{\otimes n}(1)) \otimes_{\mathbb{Q}_p} F_v \hookrightarrow \left((H^1_p)^{\otimes n}(1) \otimes_{\mathbb{Q}_p} \mathbb{C}_p\right)^{G_{F_v}}.$$

The space on the right has dimension ng^n over F_v, since $\mathbb{C}_p(r)^{G_{F_v}} = 0$ for nonzero r. This gives the required ng^n bound over $\mathbb{Q}_p$. Summing over the fixed finite set $T_{F,f}$ proves (5.4). □

Corollary 5.6. *Assume the averaged estimate (4.3) for the arithmetic module M. Equivalently, by (5.1), its conclusion holds for the ambient module $\widetilde{M}$. Then*

$$\frac{1}{N} \sum_{N \le n < 2N} \frac{h_n}{d_n} \longrightarrow 0. \tag{5.6}$$

Moreover, the sequence h_n/d_n is bounded.

Proof. By Proposition 5.4, Lemma 5.5, and the dimension asymptotic 3.6, with $\lambda = g + \sqrt{g^2-1} > g$, we have

$$\frac{h_n}{d_n} \le \frac{\dim \mathrm{Hom}_A(\widetilde{M}, W_n)}{d_n} + O_{F,T,C}(n^3 (g/\lambda)^n).$$

The first term has dyadic average tending to zero by (4.3) and (5.1); the second tends to zero exponentially.

For boundedness, choose q generators of $\widetilde{M}$ over A. Evaluation on those generators injects $\mathrm{Hom}_A(\widetilde{M}, W_n)$ into $W_n^{\oplus q}$, so its dimension is at most qd_n. Each local invariant space has dimension at most d_n. Equation (5.3) therefore bounds h_n/d_n uniformly for $n \ge 3$; the finitely many initial terms can be included by enlarging the constant. □

5.5 From a degree average to the cumulative dimensions

Lemma 5.7. *Let $d_n > 0$ satisfy $d_n \sim \lambda^n/n$ for some $\lambda > 1$, and let $0 \le b_n \le B$. If*

$$\frac{1}{N} \sum_{N \le n < 2N} b_n \longrightarrow 0,$$

then, writing $D_N = \sum_{n \le N} d_n$,

$$\liminf_{N \to \infty} \frac{\sum_{n \le N} d_n b_n}{D_N} = 0.$$

In fact the dyadic averages of this last quotient tend to zero.

Proof. The dimension asymptotic implies

$$D_N \sim \frac{\lambda}{\lambda - 1} \frac{\lambda^N}{N}.$$

For example, this follows by splitting the sum into its last fixed number of terms and an exponentially small earlier part, and then letting the number of retained terms increase. In particular, for every fixed integer $R \ge 1$, $D_{N-R}/D_N \longrightarrow \lambda^{-R}$.

Put $r_N = D_N^{-1} \sum_{n \le N} d_n b_n$. The first $N - R$ terms contribute at most $B D_{N-R}/D_N$, while each of the last R coefficients d_{N-r}/D_N is at most 1. Hence

$$r_N \le B \frac{D_{N-R}}{D_N} + \sum_{r=0}^{R-1} b_{N-r}.$$

Average this inequality over $M \le N < 2M$, with R fixed. Each shifted average of b_{N-r} tends to zero: its difference from the corresponding unshifted dyadic average is bounded by $2rB/M$. Therefore

$$\limsup_{M \to \infty} \frac{1}{M} \sum_{M \le N < 2M} r_N \le B \lambda^{-R}.$$

Letting R tend to infinity proves the assertion, since $r_N \ge 0$. □

Applying this lemma with $b_n = h_n/d_n$, Corollary 5.6 gives a sequence $N_j \to \infty$ such that

$$\sum_{n \le N_j} h_n = o(D_{N_j}). \tag{5.7}$$

This conclusion is a subsequence statement. Pointwise decay of h_n/d_n has not been asserted or used.

5.6 Complex conjugation and the final dimension comparison

Lemma 5.8. *Let c be complex conjugation, and let V_n^- be its (-1)-eigenspace. Then*

$$\dim V_n^- = (1/2 + o(1)) d_n \tag{5.8}$$

as $n \to \infty$, in all degrees.

Proof. Comparison with complex cohomology gives $\operatorname{tr}(c \mid H) = 0$: complex conjugation exchanges the two Hodge summands of dimension g. It acts by -1 on $\omega \simeq \mathbb{Q}_p(1)$. Let $\mathfrak{l} = \bigoplus_n V_n$. The equivariant quadratic resolution of its enveloping algebra, as in the surface-Lie presentation, therefore gives

$$\sum_{n\geq 0} \operatorname{tr}(c \mid U(\mathfrak{l})_n)t^n = \frac{1}{1-t^2}.$$

Define $b_m(x)$ for an operator x preserving the presentation by

$$\log\left(\sum_{n\geq 0} \operatorname{tr}(x \mid U(\mathfrak{l})_n)t^n\right) = \sum_{m\geq 1} \frac{b_m(x)}{m} t^m.$$

The preceding identity gives $b_m(c) = 1 + (-1)^m$, while evaluation at the identity gives $b_m(1) = \lambda^m + \lambda^{-m}$. Equivariant PBW and Möbius inversion yield

$$n \operatorname{tr}(c \mid V_n) = \sum_{r\mid n} \mu(r) b_{n/r}(c^r).$$

For odd r, the summand is bounded in absolute value by 2. For even r, we have $r \geq 2$ and the summand is bounded by $2\lambda^{n/2}$. Consequently

$$\operatorname{tr}(c \mid V_n) = O_g\left(\frac{\tau(n)\lambda^{n/2}}{n}\right) = o(d_n).$$

Since c is an involution in characteristic zero, $\dim V_n^- = (d_n - \operatorname{tr}(c \mid V_n))/2$, proving (5.8). For odd n the displayed inversion in fact gives trace zero exactly. □

Proposition 5.9. *Under the averaged estimate (4.3), there are arbitrarily large depths N for which*

$$\dim H^1_f(G_{\mathbb{Q},T}, U_N) < \dim U_N^{\mathrm{dR}}/F^0. \tag{5.9}$$

Here the source imposes the usual crystalline condition at p and unramifiedness outside T.

Proof. At a good prime away from p, the eigenvalues on V_n have weight $-n$, because V_n is a quotient of $H^{\otimes n}$. In particular, $H^0(G_{\mathbb{Q},T}, V_n) = 0$. The global Euler characteristic over $\mathbb{Q}$ is therefore

$$\dim H^1(G_{\mathbb{Q},T}, V_n) = \dim V_n^- + h_n. \tag{5.10}$$

This is the Euler-characteristic formula used in [8, pp. 123–124] and [3, preprint, p. 12].

Successive central extensions of U_N have kernels V_n. The fibers of the corresponding maps in nonabelian cohomology have dimension at most $\dim H^1(G_{\mathbb{Q},T}, V_n)$. Induction on N, and the fact that additional local Selmer conditions only decrease the dimension, give

$$\dim H^1_f(G_{\mathbb{Q},T}, U_N) \leq \sum_{n\leq N} \dim H^1(G_{\mathbb{Q},T}, V_n).$$

Combining (5.10), Lemma 5.8, and (5.7), we obtain along the selected subsequence

$$\dim H^1_f(G_{\mathbb{Q},T}, U_{N_j}) \leq (1/2 + o(1))D_{N_j}. \tag{5.11}$$

On the de Rham side, the degree-one homology has Hodge types $(-1,0)$ and $(0,-1)$, each with multiplicity g. Thus $\dim F^0(H^{\mathrm{dR}})^{\otimes n} = g^n$. The iterated-bracket surjection onto V_n^{dR} is strict for the Hodge filtration, as a morphism of the corresponding mixed Hodge

structures; the resulting dimension statement descends to the algebraic de Rham realization. Consequently

$$\dim F^0 V_n^{\mathrm{dR}} \le g^n.$$

This is also the explicit Hodge estimate in [8, pp. 123–124]. Strictness on the successive lower-central quotients gives

$$\dim U_N^{\mathrm{dR}}/F^0 \ge \sum_{n\le N} (d_n - g^n) = (1 - o(1))D_N, \tag{5.12}$$

since $g < \lambda$. Equations (5.11) and (5.12) prove (5.9) for all sufficiently large members of the subsequence. □

Corollary 5.10. *Assume the averaged estimate* (4.3)*. Then the Chabauty–Kim set of $C/\mathbb{Q}$ is finite at some finite depth, and hence at every sufficiently large depth. In particular, $C(\mathbb{Q})$ is finite.*

Proof. Choose a depth satisfying Proposition 5.9. The localization image of the global Selmer variety lies in a proper Zariski-closed subset of the affine de Rham quotient U_N^{dR}/F^0. There is consequently a nonzero regular function on that quotient vanishing on the localization image. Kim's unipotent Albanese density theorem says that its pullback to any residue disk of $C(\mathbb{Q}_p)$ is not identically zero; the pullback is a Coleman function. These statements, and the resulting dimension criterion, are [8, Theorem 1 and Section 3, pp. 120–121].

On the compact $\mathbb{Q}_p$-points of a residue disk, a nonzero such analytic function has finitely many zeros. There are finitely many residue disks, so the inverse image of the localization image is finite. It contains $C(\mathbb{Q})$. The Chabauty–Kim sets are nested as the depth increases; thus finiteness at this one depth implies finiteness at every greater depth. □

Remark 5.11. All uses of the auxiliary field F occur in upper bounds on the second cohomology. The Euler characteristic (5.10) and the final local dimension comparison are over $\mathbb{Q}$. In particular, no factor $[F : \mathbb{Q}]$ has been introduced into the final one-half gap. The proof does not automatically extend the same conclusion to a curve over an arbitrary number field: after restriction of scalars the local geometric space has higher dimension, and an ambient dimension deficit alone does not give the one-variable analytic finiteness argument used above.

Remark 5.12. The corollary is a deduction from (4.3), not an independent proof of that estimate. The arithmetic steps require Bogomolov's homotheties, the unramified Iwasawa torsion theorem, standard local and global duality and weight comparisons, and Kim's dimension criterion. They require neither semisimplicity of the actual Tate-module representation nor a finiteness theorem for families of pointed curves.

6 Scope, dependencies and review obligations

6.1 What semisimplicity has been replaced by

The ambient representation-theoretic decomposition uses complete reducibility for the algebraic group Sp_{2g} in characteristic zero. It does not use complete reducibility of its restriction to the actual Galois image. The surface-Lie construction extends to the full ambient group, so this restriction is precisely the original representation, with its original extension data.

The unramified Iwasawa module remains attached to the actual division-field tower. The ambient torsion module is an auxiliary algebraic object formed from its annihilators; no field extension with full symplectic Galois image is postulated. Only closedness of the actual compact image, and the resulting injection of completed group algebras, are needed. Bogomolov's homothety assertion remains necessary in the stated argument for the vanishing of the inflation term. His algebraicity assertion and the Xie–Yang semisimplicity theorem are not used.

One must not instead take coinvariants of the same arithmetic module under a putative unipotent radical and assume torsion persists. For instance, with $\Gamma = \mathbb{Z}_p \times \Gamma'$ and u a generator of the first factor, the torsion module

$$R_\Gamma / R_\Gamma(u-1)$$

has coinvariants isomorphic to $R_{\Gamma'}$, which has rank one. The ambient construction avoids this problem by preserving the original nonzero annihilators.

6.2 Why the argument currently stops at $\mathbb{Q}$

For a curve over a general number field K, the relevant local space obtained from restriction of scalars has dimension $[K : \mathbb{Q}]$. A nonzero analytic equation on this space need not have a finite zero set. Dogra's method [5] controls intersections on subvarieties of the restriction of scalars, together with kernels of localization. These relative estimates are not supplied merely by a normalized bound on the total H^2 obstruction.

Schematically, a required comparison has the form

$$\sum_{i \le N} c_i(Z) - \sum_{i<N} \kappa_i > \dim Z,$$

where $c_i(Z)$ is a graded intersection codimension for a relevant subvariety Z, and κ_i measures a localization kernel, with appropriate Selmer conditions. An estimate $o(D_N)$ for a total error does not control it relative to a possibly much smaller codimension. No assertion that the present averaging proves this comparison is made. Finiteness of $C(K)$ for general K is therefore outside the conclusion of this draft.

6.3 The points on which the candidate should be tested

The proposed estimates are written as complete arguments so that they can be checked or refuted precisely. The central obligations are:

1. The injection of bounded Iwasawa measures into the stated p-adically completed enveloping algebra, with action compatible with all algebraic representations in the weight window.

2. Integral PBW spanning for the *actual image order*, and the reduction of its cokernel. Replacing this order by a saturated matrix lattice, or by its image on reduced representation lattices, would invalidate the proof.

3. Uniformity of the character estimates in the highest weight, together with the dimension-weighted tail bound needed for $K = \sqrt{N}\log N$.

4. The reduction to one fixed everywhere-unramified torsion module and the finite-place local estimates. The much larger module allowing arbitrary ramification at T must not be substituted for the unramified module.

5. The passage from graded averages to cumulative Selmer dimensions. Smallness at isolated individual degrees is insufficient, which is why the shifted-window argument is included.

The manuscript provides no effective value of the first successful depth and no explicit algorithm for computing the finite set of rational points. It also makes no claim of priority for the aggregate estimate or for the combination of these ideas. The dependency simplification is a formal result within the proposed framework; it is not a substitute for independent verification of the framework's new estimates.

References

[1] K. Ardakov and S. Wadsley, *On irreducible representations of compact p-adic analytic groups*, Ann. of Math. (2) **178** (2013), 453–557. doi:10.4007/annals.2013.178.2.3. The completion comparison used here is in Section 10.4; see also arXiv:1102.2606.

[2] F. A. Bogomolov, *Points of finite order on an Abelian variety*, Math. USSR-Izv. **17** (1981), 55–72. doi:10.1070/IM1981v017n01ABEH001329. The scalar conclusion of the Tate-image theorem is the input used here.

[3] J. Coates and M. Kim, *Selmer varieties for curves with CM Jacobians*, Kyoto J. Math. **50** (2010), 827–852. doi:10.1215/0023608X-2010-015; arXiv:0810.3354.

[4] J. D. Dixon, M. P. F. du Sautoy, A. Mann and D. Segal, *Analytic pro-p groups*, second edition, Cambridge Studies in Advanced Mathematics 61, Cambridge University Press, 1999. doi:10.1017/CBO9780511470882.

[5] N. Dogra, *Unlikely intersections and the Chabauty–Kim method over number fields*, Math. Ann. **389** (2024), 1–62. doi:10.1007/s00208-023-02638-2.

[6] G. Faltings, *Diophantine approximation on abelian varieties*, Ann. of Math. (2) **133** (1991), 549–576. doi:10.2307/2944319.

[7] W. Fu, *Sharp bounds for multiplicities of Bianchi modular forms*, Ann. of Math. (2) **200** (2024), 123–152. Journal version; arXiv:2201.11190.

[8] M. Kim, *The unipotent Albanese map and Selmer varieties for curves*, Publ. Res. Inst. Math. Sci. **45** (2009), 89–133. doi:10.2977/PRIMS/1234361156. The dimension criterion and the local cohomology estimates used here are in Section 3, especially pp. 120–125.

[9] M. Kim, *Galois theory and Diophantine geometry*, 2009, arXiv:0908.0533.

[10] J. P. Labute, *On the descending central series of groups with a single defining relation*, J. Algebra **14** (1970), 16–23. doi:10.1016/0021-8693(70)90130-4.

[11] O. Venjakob, *On the Iwasawa theory of p-adic Lie extensions*, Compos. Math. **138** (2003), 1–54. doi:10.1023/A:1025413030203; arXiv:math/0106270. The relevant statement is Corollary 4.8(ii), with finite generation immediately preceding it (p. 29 in the author's preprint).

[12] J. Xie and Z. Yang, *Faltings' Isogeny Theorem via Equidistribution*, 2026, arXiv:2609.10302v2. Cited for the origin of the investigation; its semisimplicity theorem is not an input to the ambient-group version of the present candidate.