Algebra & Number Theory

Volume 19
2025
No. 1

Ranks of abelian varieties in cyclotomic twist families

Ari Shnidman and Ariel Weiss

msp

ALGEBRA AND NUMBER THEORY 19:1 (2025)
https://doi.org/10.2140/ant.2025.19.39

msp

Ranks of abelian varieties in cyclotomic twist families

Ari Shnidman and Ariel Weiss

Let A be an abelian variety over a number field F, and suppose that $\mathbb{Z}[\zeta_n]$ embeds in $\mathrm{End}_{\bar{F}} A$, for some root of unity ζ_n of order $n = 3^m$. Assuming that the Galois action on the finite group $A[1-\zeta_n]$ is sufficiently reducible, we bound the average rank of the Mordell–Weil groups $A_d(F)$, as A_d varies through the family of μ_{2n}-twists of A. Combining this result with the recently proved uniform Mordell–Lang conjecture, we prove near-uniform bounds for the number of rational points in twist families of bicyclic trigonal curves $y^3 = f(x^2)$, as well as in twist families of theta divisors of cyclic trigonal curves $y^3 = f(x)$. Our main technical result is the determination of the average size of a 3-isogeny Selmer group in a family of μ_{2n}-twists.

1. Introduction

Let A be an abelian variety over a number field F and let $G_F = \mathrm{Gal}(\bar{F}/F)$. Any G_F-stable subgroup $H \subset \mathrm{Aut}_{\bar{F}} A$ gives rise to a twist family of abelian varieties A_ξ over F, indexed by the elements ξ of the Galois cohomology set $H^1(G_F, H)$. The base change of A_ξ to $\bar{F}$ is isomorphic to $A_{\bar{F}}$, but with G_F-action twisted by ξ. Our goal is to study the distributions of the ranks of the Mordell–Weil groups $A_\xi(F)$ in such twist families, and to give some applications.

Every abelian variety A has the automorphism -1, and since $H^1(G_F, \{\pm 1\}) \simeq F^\times/F^{\times 2}$, we obtain the *quadratic twist family* of A. The average rank of $A_\xi(F)$ in quadratic twist families has been extensively studied in the case of elliptic curves [Brumer 1992; Heath-Brown 1994; Katz and Sarnak 1999; Smith 2017; Bhargava et al. 2019], with [Bhargava et al. 2019] addressing many cases in higher dimension as well.

In this paper, we consider the case $H = \mu_{2n}$, the group of $2n$-th roots of unity, where $n = 3^m$ for some $m \geq 1$. More precisely, we assume that there is a G_F-equivariant ring embedding $\mathbb{Z}[\zeta] \hookrightarrow \mathrm{End}_{\bar{F}} A$, where $\zeta = \zeta_n \in \bar{F}$ is a root of unity of order n. We say that such an A has *ζ-multiplication*. For example, the Jacobian J of a curve of the form $y^3 = xf(x^{3^{m-1}})$ has ζ-multiplication induced from the order n automorphism $(x, y) \mapsto (\zeta^3 x, \zeta y)$.

Since $\mu_{2n} = \langle -\zeta \rangle \subset \mathrm{Aut}_{\bar{F}} A$, there is a twist A_d for each $d \in F^\times/F^{\times 2n} \simeq H^1(G_F, \mu_{2n})$. In the example above, J_d is the d-th quadratic twist of the Jacobian of $y^3 = xf\left(\frac{1}{d}x^{3^{m-1}}\right)$. In Section 5B, we recall a height function $h : F^\times/F^{\times 2n} \to \mathbb{R}$ with the property that the sets $\Sigma_X := \{d \in F^\times/F^{\times 2n} : h(d) < X\}$ are finite. When $F = \mathbb{Q}$, the height $h(d)$ is the absolute value of the smallest integer representing d. The average rank of $A_d(F)$ is then, by definition,

$$\mathrm{avg}_d \,\mathrm{rk}\, A_d(F) = \lim_{X \to \infty} \operatorname*{avg}_{d \in \Sigma_X} \mathrm{rk}\, A_d(F).$$

MSC2020: primary 11G10; secondary 11E76, 11S25, 14G05.
Keywords: arithmetic statistics, twist families, rational points on curves, ranks of abelian varieties.

In general, it is not known whether this limit exists or even if the limsup is finite. In the latter case, we say that the *average rank of $A_d(F)$ is bounded*.

1A. ***Mordell–Weil ranks.*** If A has ζ-multiplication, the endomorphism $1-\zeta \in \mathrm{End}_{\bar{F}} A$ descends to an isogeny $\pi : A \to A'$ over F (see Section 2). The kernel $A[\pi]$ is a G_F-stable subgroup of the 3-torsion group $A[3]$, and hence is a finite-dimensional $\mathbb{F}_3$-vector space. We show that the average rank of $A_d(F)$ is bounded, assuming that the G_F-action on $A[\pi]$ is sufficiently reducible.

Theorem 1.1. *Let A be an abelian variety with ζ_{3^m}-multiplication over F.*

(i) *If $A[\pi]$ is a direct sum of characters, then* $\mathrm{avg}_d \,\mathrm{rk}\, A_d(F)$ *is bounded.*

(ii) *If $A[\pi]$ has a full flag, then* $\mathrm{avg}_d \,\mathrm{rk}\, A_d(F)$*, over squarefree $d \in F^\times/F^{\times 2n}$, is bounded.*

Here, we say that $A[\pi]$ *has a full flag* if there are G_F-modules $0 = H_0 \subset H_1 \subset \cdots \subset H_k = A[\pi]$ such that $\dim_{\mathbb{F}_3} H_{i+1}/H_i = 1$. We say that $d \in F^\times/F^{\times 2n}$ is *squarefree* if $v(d) \equiv 0$ or $1 \pmod{2n}$, for all finite places v of F. Theorem 1.1 is a simultaneous generalization of [Bhargava et al. 2019, Theorem 2.2] and [Bhargava et al. 2020, Theorem 5] to a larger class of twist families of abelian varieties.

If J is the Jacobian of $y^3 = xf(x^{3^{m-1}})$, then the representation theoretic conditions on $J[\pi]$ translate into conditions on the Galois group $\mathrm{Gal}(f)$ of the splitting field of $f(x)$ over F. More generally, we deduce the following result from Theorem 1.1:

Corollary 1.2. *Let $f(x) \in F[x]$ be separable and nonconstant, and let J be the Jacobian of either $y^{3^m} = f(x)$ or $y^3 = xf(x^{3^{m-1}})$.*

(i) *If* $\mathrm{Gal}(f) \simeq (\mathbb{Z}/2\mathbb{Z})^k$*, for some $k \geq 0$, then* $\mathrm{avg}_d \,\mathrm{rk}\, J_d(F)$ *is bounded.*

(ii) *If* $\mathrm{Gal}(f)$ *is an extension of $(\mathbb{Z}/2\mathbb{Z})^k$ by a 3-group, then* $\mathrm{avg}_d \,\mathrm{rk}\, J_d(F)$*, over squarefree $d \in F^\times/F^{\times 2n}$, is bounded.*

Our proof of Theorem 1.1 gives an explicit upper bound on the average rank, however, the bound depends on subtle arithmetic properties of A. The following crude upper bound has the virtue that it applies to a large class of abelian varieties and depends only mildly on A.

Theorem 1.3. *Suppose that $A[\pi]$ has a full flag and that A admits a ζ_n-stable principal polarization. Let S be the set of places of F dividing $3\mathfrak{f}_A\infty$, where $\mathfrak{f}_A$ is the conductor of A. Then the average rank of $A_d(F)$, for squarefree $d \in F^\times/F^{\times 2n}$, is at most* $\dim A \cdot (\#S + 3^{-\#S})$.

For most A, this bound is significantly weaker than what our method actually gives. An interesting case is when A has complex multiplication (CM), i.e., $\dim A = 3^{m-1}$, in which case $\dim_{\mathbb{F}_3} A[\pi] = 1$ and the reducibility hypotheses are automatically satisfied. When the complex multiplication is defined over F, we obtain especially strong results:

Theorem 1.4. *Suppose that* $\dim A = 3^{m-1}$*, so that A has complex multiplication by $\mathbb{Z}[\zeta_{3^m}]$. Assume moreover that $\zeta_{3^m} \in F$, so that the complex multiplication is defined over F. Then the average $\mathbb{Z}[\zeta_{3^m}]$-module rank of $A_d(F)$ is at most $\frac{1}{2}$, and at least 50% of twists A_d have rank 0.*

In the CM case, we expect that 100% of twists A_d have rank 0, in which case our result is halfway towards the analogue of Goldfeld's conjecture in this context.

Such A arise as factors of the Jacobians of the curves $y^{3^m} = x^a(1-x)^b$, which have good reduction away from 3. Even when the CM is not defined over F, we obtain average rank bounds that depend only on $\dim A$. Over $\mathbb{Q}$, for example, the average rank of $A_d(\mathbb{Q})$ is at most $\frac{19}{9}\dim A$ by Theorem 1.3, a bound which can be improved to $\frac{13}{9}\dim A$ with a more refined analysis.

1B. ***Rational points on curves.*** Theorem 1.1 has concrete consequences for the arithmetic of curves C/F of genus $g \geq 2$. It was nearly 40 years ago that Faltings proved that $C(F)$ is a finite set, but very recently, there has been significant progress towards a uniform upper bound on $\#C(F)$. Building on work of Dimitrov, Gao, and Habegger [Dimitrov et al. 2021], Kühne [2021] has shown that

$$\#C(F) \leq c_g^{1+\operatorname{rk}\operatorname{Jac}(F)},$$

where c_g is a constant depending only g. Building on this work, Gao, Ge, and Kühne [Gao et al. 2021] proved the more general uniform Mordell–Lang conjecture for closed subvarieties of abelian varieties. These results reduce the question of uniform bounds for rational points on a large class of varieties to a question about ranks of abelian varieties.

By combining these results with Theorem 1.1, we show that "near-uniformity" holds for twists of bicyclic trigonal curves:

Theorem 1.5. *Let $f(x) \in F[x]$ be separable, of degree at least two, and with all of its roots nonzero elements of F. Consider the bicyclic trigonal curve $C : y^3 = f(x^2)$, and let $C_d : dy^3 = f(dx^2)$ be the corresponding sextic twist family. Then for every $\varepsilon > 0$, there is a constant N_ε such that the lower density of classes $d \in F^\times/F^{\times 6}$ for which*

$$\#C_d(F) \leq N_\varepsilon$$

is at least $1-\varepsilon$.

To prove Theorem 1.5, we apply Theorem 1.1 not to the Jacobian of C, but to the Prym variety for the double cover $C \to C'$, where $C' : y^3 = f(x)$; see Section 9. We remark that the constant N_ε depends only on ε, $\deg(f)$, and $\#S$ (using the notation of Theorem 1.3).

Unlike the curves in Theorem 1.5, a general cyclic trigonal curve $C : y^3 = f(x)$ has no sextic twists, so it may seem that Theorem 1.1 says nothing about rational points on the twists of C itself. However, for these curves, we can consider sextic twists of a theta divisor $\Theta \subset J = \operatorname{Jac}(C)$. Recall that Θ is birational to the symmetric power $C^{(g-1)}$, so its rational points parametrize low-degree points on C. We can choose Θ so that it is preserved by the μ_{2n}-action (see Section 9), which allows us to consider the twist $\Theta_d \subset J_d$, for each $d \in F^\times$.

Theorem 1.6. *Let $f(x) \in F[x]$ be separable and suppose that $\operatorname{Gal}(f) \simeq (\mathbb{Z}/2\mathbb{Z})^k$ for some $k \geq 0$. Let $C : y^3 = f(x)$, and suppose that $\operatorname{Jac}(C)$ is geometrically simple. Let $\Theta \subset \operatorname{Jac}(C)$ be a symmetric theta*

divisor. Then for every $\varepsilon > 0$, there is a constant N_ε such that the lower density of classes $d \in F^\times/F^{\times 6}$ for which

$$\#\Theta_d(F) \leq N_\varepsilon$$

is at least $1-\varepsilon$.

This result again follows from Theorem 1.1 and [Gao et al. 2021], and N_ε depends only on ε, $\deg(f)$, and $\#S$. Since the results of [Gao et al. 2021] are ineffective, we cannot say anything explicit about the constant N_ε in general. However, one can prove explicit results in this direction by instead combining our work with the Chabauty method. We illustrate this by way of an example.

Theorem 1.7. *Consider the sextic twist family $C_d : y^3 = (x^2-d)(x^2-4d)$ of genus-3 curves. For at least $\frac{1}{3}$ of squarefree $d \in \mathbb{Z}$ such that $d \equiv 2$ or $11 \pmod{36}$, we have $\#C_d(\mathbb{Q}) \leq 5$.*

The curve C_d admits a double cover $p : C_d \to E_d$ to the elliptic curve $E_d : y^3 = (x-d)(x-4d)$. Moreover, C_d embeds in the abelian surface $P_d = \mathrm{Jac}(C_d)/p^*\,\mathrm{Pic}^0(E_d)$. By making the rank bound in Theorem 1.1 explicit, we show that $\mathrm{rk}\, P_d(\mathbb{Q}) \leq 1$ for at least $\frac{1}{3}$ of twists d. Then we invoke, and generalize slightly, Stoll's uniform Chabauty result for twist families [2006].

The same method works for sextic twist families of the form $C_{a,d} : y^3 = (x^2-d)(x^2-ad)$. To prove the existence of twists with $\mathrm{rk}\, P_{a,d}(\mathbb{Q}) \leq 1$, we must check that a certain local 3-adic root number takes the value -1 for some twist d. We can verify this condition in Magma for seemingly any given curve $C_{a,d}$, but it would be nice to have a proof for all or most values of a.[1] It would also be interesting to prove explicit results for symmetric squares of trigonal plane quartics, as in Theorem 1.6, by using [Caro and Pasten 2023].

1C. ***3-isogeny Selmer groups.*** Having discussed some applications of Theorem 1.1, let us discuss its proof. Theorem 1.1 follows from a more precise result about Selmer groups. Let A/F be an abelian variety with ζ-multiplication and admitting a 3-isogeny $\phi : A \to B$. If $A[\phi] \subset A[\pi]$, or equivalently, if ϕ is ζ-linear, then each twist A_d is endowed with its own 3-isogeny $\phi_d : A_d \to B_d$. For each d, we consider the ϕ_d-Selmer group $\mathrm{Sel}_{\phi_d}(A_d)$, which sits in the exact sequence

$$0 \to B_d(F)/\phi_d A_d(F) \to \mathrm{Sel}_{\phi_d}(A_d) \to Ш(A_d)[\phi_d] \to 0.$$

The main technical result of this paper is the exact computation of $\mathrm{avg}_d \,\#\mathrm{Sel}_{\phi_d}(A_d)$.

To state the precise result, we recall the global Selmer ratio $c(\phi_d) = \prod_v c_v(\phi_d)$, where for each place v of F, we define

$$c_v(\phi_d) = \frac{\#\mathrm{coker}(A_d(F_v) \to B_d(F_v))}{\#\ker(A_d(F_v) \to B_d(F_v))}.$$

For $v \nmid 3\infty$, we have $c_v(\phi_d) = c_v(B_d)/c_v(A_d)$, where $c_v(A)$ is the Tamagawa number of A over F_v. Thus, up to some subtle factors at places v above 3 and ∞, the number $c(\phi_d)$ is the ratio of the global Tamagawa

[1] In [Shnidman and Weiss 2023, Theorem 1.4], we prove that a positive proportion of $P_{a,d}$ have rank at most 1 in the case that a is a square, using a different argument which sidesteps the root number question.

numbers $c(B_d)/c(A_d)$. In particular, we have $c_v(\phi_d) \in 3^{\mathbb{Z}}$, and $c_v(\phi_d) = 1$ for all but finitely many v (having fixed d).

We say $A[\phi]$ is *almost everywhere locally a direct summand* of $A[\pi]$ if, for almost all places v of F, the G_{F_v}-module $A[\phi]$ is a direct summand of $A[\pi]$.

Theorem 1.8. *Assume that $A[\phi]$ is almost everywhere locally a direct summand of $A[\pi]$. Then $\operatorname{avg}_d \#\operatorname{Sel}_{\phi_d}(A_d) = 1 + \operatorname{avg}_d c(\phi_d)$, where both averages are finite and taken over $d \in F^\times/F^{\times 2n}$, ordered by height.*

This result is a simultaneous generalization of [Bhargava et al. 2019, Theorem 2.1] and [Bhargava et al. 2020, Theorem 1]. Interestingly, the condition of being everywhere locally a direct summand, which is automatically satisfied for the families considered in [Bhargava et al. 2019; 2020], seems to be an obstruction to computing the average size of $\#\operatorname{Sel}_{\phi_d}(A_d)$ in the entire family of twists, at least using our methods. In any case, if we only consider squarefree twists, then this obstruction goes away:

Theorem 1.9. *Let $\phi : A \to B$ be a ζ-linear 3-isogeny. Then the average size of $\#\operatorname{Sel}_{\phi_d}(A_d)$ over squarefree $d \in F^\times/F^{\times 2n}$ is finite and equal to $1 + \operatorname{avg}_d c(\phi_d)$.*

The quantity $\operatorname{avg}_d c(\phi_d)$ is governed by local arithmetic data which can be made explicit in certain cases. For example, in Theorem 1.4 we have $c(\pi_d) = c(\phi_d) = 1$ for all d. However, in general, computing the exact value of $\operatorname{avg}_d c(\phi_d)$ is hard. Nonetheless, one can give an explicit upper bound on $\operatorname{avg}_d \#\operatorname{Sel}_{\phi_d}(A_d)$ depending only on F, $\dim A$, and the number of primes dividing the conductor of A (Proposition 5.5).

In Section 6 we show how to deduce Theorem 1.1 from Theorems 1.8 and 1.9. In the remainder of the introduction, we discuss the proofs of the latter two results.

1D. ***Methods.*** We prove Theorems 1.8 and 1.9 using geometry-of-numbers methods. As in the previous works [Bhargava et al. 2019; 2020] of the first author and his collaborators, we first identify the elements of $\operatorname{Sel}_{\phi_d}(A_d)$ with $\mathrm{SL}_2(F)$-orbits of binary cubic forms of discriminant d. We then wish to use lattice-point counting techniques, which have been extended to global fields in [Bhargava et al. 2015], to count the number of such orbits of bounded discriminant. However, it is not at all clear (and indeed, it is not always true) that the $\mathrm{SL}_2(F)$-orbits corresponding to Selmer elements are *integral*, i.e., that they contain cubic forms whose coefficients are algebraic integers. This integrality is of course essential for lattice-point counting.

For the quadratic twist families considered in [Bhargava et al. 2019], integrality follows quickly once it is realized that the local Selmer conditions are very mild outside the finitely many primes dividing the conductor of A. One needs only to "clear denominators" at those finitely many primes, and then the Selmer orbits become integral. For the families considered in this paper, the question of integrality is more subtle. The first interesting case is the family of sextic twists $E : y^2 = x^3 + d$ considered in [Bhargava et al. 2020], which is the unique twist family of elliptic curves with 3-power ζ-multiplication. In that special family, the authors give an *explicit* bijection between Selmer elements and binary cubic forms. Integrality is then proven using a direct connection with Bhargava's higher composition laws.

In the more general setup of this paper, we cannot rely on such an explicit parametrization, nor do we expect one to exist. Instead, our method is more abstract and involves two independent steps. First, we give a complete analysis of the *integral* arithmetic invariant theory for the representation $\mathrm{Sym}^3 \mathbb{Z}^2$ of SL_2, in the sense that we identify precisely which $\mathrm{SL}_2(F_v)$-orbits of binary cubic forms, over a local field F_v, have integral representatives. We show that once the valuation of the discriminant $v(d)$ is at least 3, all orbits have integral representatives, and we determine what happens for $v(d) \leq 2$ as well. Our strategy is to translate the question into one about cubic rings, whose local structure we understand well. Second, we study the Selmer groups $\mathrm{Sel}_{\phi_d}(A_d)$ from a purely cohomological point of view, in the spirit of Mazur and Rubin [2007]; see also [Klagsbrun et al. 2013]. Upon comparing the results, we find that for all but finitely many primes v, the Selmer orbits of discriminant d are v-integral, except possibly when $v(d) = 2$. In particular, the Selmer orbits are integral when d is squarefree. When $v(d) = 2$, we find that the local direct summand condition on $A[\phi] \subset A[\pi]$ exactly matches up with the local integrality condition.

1E. ***Future directions.*** In many situations, the integral orbits of a reductive group G acting on a representation V have been shown to parametrize Selmer elements in a certain explicit family of abelian varieties [Bhargava and Shankar 2015a; 2015b; Bhargava and Ho 2016; Bhargava and Gross 2014; Thorne 2013; Laga 2023]. The results of this paper show that there is a tremendous amount of flexibility in these constructions, in the sense that (G, V) can be used to parametrize Selmer elements in very different looking families over the same space of invariants $V /\!\!/ G$ (which is $\mathbb{A}^1$ in our case). Our analysis of the integral arithmetic invariant theory of $(\mathrm{SL}_2, \mathrm{Sym}^3 \mathbb{Z}^2)$ can be adapted to some of these other representations (G, V), so it would be interesting to understand the following (vaguely formulated) question. For which families $\phi : \mathcal{A} \to \mathcal{B}$ of isogenies of abelian varieties over $S = V /\!\!/ G$ can the elements of $\mathrm{Sel}_{\phi_s}(\mathcal{A}_s)$, for $s \in S(F)$, be parametrized by orbits of $G(F)$ on $V(F)$?

1F. ***Outline.*** We begin in Section 2 with basics on abelian varieties with ζ-multiplication. Sections 3–6 are the technical heart of the paper. In Section 3, we give a complete analysis of the integral arithmetic theory for the representation $\mathrm{Sym}^3 \mathbb{Z}^2$ of SL_2. In Section 4, we give a parallel, but independent analysis of the Selmer groups $\mathrm{Sel}_{\phi_d}(A_d)$ of the twists ϕ_d of a general ζ-linear 3-isogeny. In Section 5, we combine these two sections and prove Theorems 1.8 and 1.9. In Section 6, we apply the results of Section 5 to prove Theorems 1.1 and 1.3.

The remainder of the paper is devoted to applications of our main results. In Section 7, we study the average ranks of the Jacobians of the curves $y^3 = xf(x^{3^{m-1}})$ and $y^{3^m} = f(x)$ and prove Corollary 1.2. In Section 8, we give explicit results for abelian varieties with CM and prove Theorem 1.4. In Section 9, we study rational points in twist families of curves as in Section 1B, and prove Theorems 1.5 and 1.6. Finally, in Section 10, we study twist families of genus-3 curves and prove Theorem 1.7.

2. Abelian varieties with ζ-multiplication

Let F be a field of characteristic 0. Fix an odd prime p, an integer $n = p^m$, and a primitive n-th root of unity $\zeta = \zeta_n$. In this section, φ denotes Euler's totient function.

Definition 2.1. An abelian variety with ζ-multiplication is a pair (A, ι_A), where A is an abelian variety over F and $\iota_A : \mathbb{Z}[\zeta] \hookrightarrow \mathrm{End}_{\bar{F}} A$ is a G_F-equivariant injective ring homomorphism.

We usually suppress any mention of ι_A and view $\mathbb{Z}[\zeta]$ as a subring of $\mathrm{End}_{\bar{F}} A$. In this section, we collect some basic facts and constructions relating to abelian varieties with ζ-multiplication.

2A. ***The isogeny π.*** If A is an abelian variety with ζ-multiplication, then since $1-\zeta$ divides p in $\mathbb{Z}[\zeta]$, the map $1-\zeta \in \mathrm{End}_{\bar{F}} A$ is an isogeny whose degree is a power of p.

Lemma 2.2. *The kernel of $1-\zeta$ is G_F-stable, and hence is an F-subgroup of $A[p]$. In particular, there is an abelian variety $A^{(1)}$ over F, such that the endomorphism $1-\zeta$ of A over $\bar{F}$ descends to an isogeny $\pi : A \to A^{(1)}$ over F.*

Proof. If $P \in A_{\bar{F}}[1-\zeta]$ and $\sigma \in G_F$, then $\zeta^{\sigma^{-1}} = \zeta^i$ for some $i \in (\mathbb{Z}/n\mathbb{Z})^\times$ and

$$\zeta(P^\sigma) = (\zeta^{\sigma^{-1}}(P))^\sigma = (\zeta^i P)^\sigma = P^\sigma,$$

which shows that $P^\sigma \in A_{\bar{F}}[1-\zeta]$. Hence, $A_{\bar{F}}[1-\zeta]$ descends to an F-subgroup H of A. Thus, we obtain a an isogeny $\pi : A \to A/H =: A^{(1)}$ over F.

The equality of ideals $(1-\zeta)^{\varphi(n)} = (p)$ in $\mathbb{Z}[\zeta]$ shows that $A_{\bar{F}}[1-\zeta] \subset A_{\bar{F}}[p]$. □

If $\zeta \in F$, then $A^{(1)} = A/A[1-\zeta] \simeq A$, and $\pi : A \to A^{(1)}$ can be identified with the endomorphism $1-\zeta$. If $\zeta \notin F$, then $A^{(1)} = A/A[\pi]$ is a twist of A which we now identify:

Lemma 2.3. *$A^{(1)}$ is the twist of A corresponding to the cocycle $\sigma \mapsto \frac{1-\zeta^\sigma}{1-\zeta} \in H^1(F, \mathbb{Z}[\zeta]^\times)$.*

Proof. Over $F(\zeta)$, the map $\eta : A/A[1-\zeta] \to A$ given by $\bar{x} \mapsto (1-\zeta)x$ defines an isomorphism. Hence, $A^{(1)}$ is the twist corresponding to the cocycle $\sigma \mapsto \eta^\sigma \eta^{-1}$. □

The abelian variety $A^{(1)}$ also has ζ-multiplication. Iterating Lemma 2.2, for each integer s, we obtain an abelian variety $A^{(s)} = A/A[(1-\zeta)^s]$. As in Lemma 2.3, $A^{(s)}$ is isomorphic to the twist of A corresponding to the cocycle

$$\sigma \mapsto \frac{(1-\zeta^\sigma)^s}{(1-\zeta)^s} \in H^1(F, \mathbb{Z}[\zeta]^\times).$$

We define π^s to be the corresponding isogeny $A = A^{(0)} \to A^{(s)}$.

Note that $A^{(\varphi(n))} \simeq A$ and that $\pi^{\varphi(n)} : A \to A$ is multiplication by pu, for some unit $u \in \mathrm{Aut}\, A$. In particular, when writing $A^{(s)}$, we can always consider s modulo $\varphi(n) = p^{m-1}(p-1)$, and we have inclusions

$$A[\pi] \subset A[\pi^2] \subset \cdots \subset A[\pi^{\varphi(n)-1}] \subset A[\pi^{\varphi(n)}] = A[p].$$

In general, the Galois action on $A^{(s)}$ is related to the Galois action on A in a convoluted way. However, on a subset of the torsion of A, the action is especially simple.

Lemma 2.4. *Let $s = p^r$ for some $0 \le r < m$, and let $i \in \mathbb{Z}$. Then, as G_F-modules,*

$$A^{(is)}[\pi^s] \simeq A[\pi^s] \otimes \chi_p^i,$$

where χ_p is the mod p cyclotomic character.

Proof. By Lemma 2.3, the abelian variety $A^{(is)}$ is the twist of A corresponding to the cocycle

$$\sigma \mapsto \frac{(1-\zeta^\sigma)^{is}}{(1-\zeta)^{is}} \in H^1(F, \mathbb{Z}[\zeta]^\times).$$

Equivalently, there is an isomorphism $\phi : A \to A^{(is)}$ over $F(\zeta)$ such that for all $\sigma \in G_F$ and $P \in A(\bar{F})$, we have

$$(\phi^{-1})^\sigma \circ \phi(P) = \frac{(1-\zeta^\sigma)^{is}}{(1-\zeta)^{is}} P.$$

Hence, if $P \in A[\pi^s]$ and $\sigma \in G_F$, then in $A^{(is)}[\pi^s]$, we have

$$(\phi(P))^\sigma = \frac{(1-\zeta^\sigma)^{is}}{(1-\zeta)^{is}} \phi(P^\sigma).$$

Suppose that $\sigma : \zeta \mapsto \zeta^j$ for some $j \in (\mathbb{Z}/n\mathbb{Z})^\times$. Then, since $(1-\zeta^{p^r})\phi(P^\sigma) = (1-\zeta)^{p^r}\phi(P^\sigma) = 0$, we have

$$\begin{aligned}\frac{(1-\zeta^\sigma)^{ip^r}}{(1-\zeta)^{ip^r}}\phi(P^\sigma) &= (1+\zeta+\zeta^2+\cdots+\zeta^{s-1})^{ip^r}\phi(P^\sigma)\\ &= (1+\zeta^{p^r}+\zeta^{2p^r}+\cdots+\zeta^{p^r(j-1)})^i\phi(P^\sigma)\\ &= j^i\phi(P^\sigma).\end{aligned}$$

Since, by definition, $j \pmod p = \chi_p(\sigma)$, we see that $(\phi(P))^\sigma = \chi_p(\sigma)^i\phi(P^\sigma)$, as claimed. □

2B. *ζ-linear isogenies.* We keep the notation $n = p^m$ and $\zeta = \zeta_n$.

Definition 2.5. Let (A, ι_A) and (B, ι_B) be abelian varieties over F with ζ-multiplication, and let $\phi : A \to B$ be an isogeny. We say that ϕ is *ζ-linear* if $\iota_B(\alpha) \circ \phi = \phi \circ \iota_A(\alpha)$ for all $\alpha \in \mathbb{Z}[\zeta]$.

Lemma 2.6. *If $\phi : A \to B$ is a ζ-linear p-isogeny, then $A[\phi] \subset A[\pi]$. Conversely, if $H \subset A[\pi]$ is a G_F-stable subgroup of order p, then the quotient $B := A/H$ inherits a ζ-multiplication from A, and the canonical p-isogeny $\phi : A \to B$ is ζ-linear.*

Proof. Since ϕ is ζ-linear, if $P \in A[\phi]$, then so is ζP. Hence, the action of ζ is given by a homomorphism $\mu_n = \mu_{p^m} \to \mathrm{Aut}_{\bar{F}} A[\phi] \simeq (\mathbb{Z}/p\mathbb{Z})^\times$, which must of course be trivial. Thus, ζ acts as the identity on $A[\phi]$, so $A[\phi] \subset A[1-\zeta] = A[\pi]$.

For the converse, since $\iota_A(\zeta)$ fixes H, we have $\ker(\phi) = \ker(\phi \circ \iota_A(\zeta))$, so that $\phi \circ \iota_A(\zeta)$ factors through ϕ. That is, there exists an automorphism $\zeta_B : B \to B$ such that $\phi \circ \iota_A(\zeta) = \zeta_B \circ \phi$. This automorphism ζ_B has order n and has the same minimal polynomial as $\iota_A(\zeta)$. Thus, the map $\iota_B : \mathbb{Z}[\zeta] \to \mathrm{End}_{\bar{F}} B$ given by $\zeta \mapsto \zeta_B$ is a ζ-multiplication on B, and the p-isogeny ϕ is ζ-linear by construction. □

2C. *Twists.* If (A, ι_A) has ζ_n-multiplication, then ι_A induces an inclusion $\mathbb{Z}[\zeta_{2n}]^\times \subset \mathrm{Aut}_{\bar{F}} A$ of G_F-modules. For each $d \in F^\times$, let A_d be the twist of A corresponding to the image of d under

$$F^\times \to F^\times/F^{\times 2n} \simeq H^1(F, \mu_{2n}) \to H^1(F, \mathrm{Aut}_{\bar{F}} A).$$

Then A_d is an abelian variety over F that becomes isomorphic to A over $F(d^{1/2n})$. Moreover, A_d also has ζ_n-multiplication.

Remark 2.7. If $\mathrm{Aut}_{\bar{F}} A = \mu_{2n}$, then distinct $2n$-th power classes d give nonisomorphic A_d. However, if $\mathrm{Aut}_{\bar{F}} A \supsetneq \mu_{2n}$, then the map $H^1(F, \mu_{2n}) \to H^1(F, \mathrm{Aut}_{\bar{F}} A)$ need not be injective, and hence the twists A_d need not be distinct.

Now let $\phi : A \to B$ be a ζ-linear p-isogeny over F. By Lemma 2.6, the automorphisms $\pm\zeta \in \mathrm{Aut}_{\bar{F}} A$ preserve the subgroup $A[\phi]$, giving an inclusion of G_F-modules $\mu_{2n} \hookrightarrow \mathrm{Aut}_{\bar{F}}(\phi)$, where $\mathrm{Aut}_{\bar{F}}(\phi)$ is the subgroup of $\mathrm{Aut}_{\bar{F}} A$ stabilizing $A[\phi]$. For $d \in F^\times$, let $\phi_d : A_d \to B_d$ be the twist of ϕ corresponding to the image of d under

$$F^\times \to F^\times / F^{\times 2n} \simeq H^1(F, \mu_{2n}) \to H^1(F, \mathrm{Aut}_{\bar{F}}(\phi)).$$

Then ϕ_d is a ζ-linear p-isogeny over F. Similarly, we may twist the isogeny $\pi : A \to A^{(1)}$ to obtain $\pi_d : A_d \to A_d^{(1)}$, and this is the canonical isogeny "π" associated to A_d (and its ζ-multiplication).

Remark 2.8. By Lemma 2.3, the abelian variety $A^{(1)}$ is the twist of A corresponding to the cocycle $\xi : \sigma \mapsto \frac{1-\zeta^\sigma}{1-\zeta}$ in $H^1(F, \mathbb{Z}[\zeta_n]^\times)$. When $n = 3$, we have $\mathbb{Z}[\zeta_{2n}]^\times = \mu_6$, and since $\frac{1-\zeta}{\sqrt[6]{-27}} \in \mu_6$, the cocycle ξ is in the same cohomology class as $-27 \in F^\times / F^{\times 6}$. It follows that $A^{(1)} = A_{-27}$, which is the quadratic twist of A by the mod 3 cyclotomic character. More generally, we have $A^{(\frac{1}{2}\varphi(n))} = A_{(p^*)^n}$, where $p^* = (-1)^{(p-1)/2} p$ is such that $\mathbb{Q}(\sqrt{p^*}) \subset \mathbb{Q}(\zeta_p)$. However, for general s and n, the twist $A^{(s)}$ need not be isomorphic to A_d, for any $d \in F^\times / F^{\times 2n}$.

3. Integral orbits of binary cubic forms

In this section, we classify the integral $\mathrm{SL}_2(F)$-orbits on the space of binary cubic forms over a local field F. We first recall some facts from [Bhargava et al. 2020, §2] and [Bhargava 2004, Theorem 13].

Let $V = \mathrm{Sym}^3 \mathbb{Z}^2$ be the space of binary cubic forms. The group SL_2 acts on V, and the ring of invariant functions is generated by the usual polynomial discriminant $\mathrm{Disc} : V \to \mathbb{Z}$. Let F be any field of characteristic not 2 or 3, and for any $d \neq 0$ in F, define $V(F)_d := \{f \in V(F) : \mathrm{Disc}(f) = d\}$. There is a unique *reducible* $\mathrm{SL}_2(F)$-orbit of cubic forms $f \in V(F)_d$. The stabilizer of such an f is a commutative F-group scheme C_d of order 3. The Galois action on $C_d(\bar{F})$ is by the quadratic character $\chi_d : \mathrm{Gal}(F(\sqrt{d})/F) \to \{\pm 1\}$.

Proposition 3.1. *The group $H^1(F, C_d)$ is in bijection with the $\mathrm{SL}_2(F)$-orbits on $V(F)_d$.*

Now let F be a local field of residue characteristic neither 2 nor 3, with surjective discrete valuation $v : F^\times \to \mathbb{Z}$, ring of integers $\mathcal{O}_F$, maximal ideal $\mathfrak{m}$ and residue field $\mathbb{F}$ of cardinality q.

We wish to determine which $\mathrm{SL}_2(F)$-orbits in $V(F)_d$ have representatives in $V(\mathcal{O}_F)_d$. We call these orbits the *integral orbits*, and we let $H^1_{\mathrm{int}}(C_d)$ be the subset of $H^1(F, C_d)$ that they correspond to under the bijection of Proposition 3.1. Of course, a necessary condition for there to be any integral orbits at all is that $d \in \mathcal{O}_F$. We see that even though the abstract group $H^1(F, C_d)$ depends only on the square-class of d, the notion of integrality depends on the actual value of d, and in particular its valuation.

We recall some facts about cubic rings over F and over $\mathcal{O}_F$ [Bhargava et al. 2013]. The action of SL_2 on V extends to the following action of GL_2: if $\gamma = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, then

$$(\gamma \cdot f)(x, y) = \frac{1}{\det \gamma} f(ax + cy, bx + dy).$$

If R is a principal ideal domain, then a *cubic ring over* R is an R-algebra S that is free of rank 3 as an R-module. The discriminant of S is a well-defined element of $R^\times / R^{\times 2}$.

Proposition 3.2 (Levi, Delone–Faddeev, Gan–Gross–Savin). *For any principal ideal domain R, there is a discriminant preserving bijection between $\mathrm{GL}_2(R)$-orbits on $V(R)$ and isomorphism classes of cubic rings over R. Moreover, this bijection is functorial in R.*

Proof. Building on [Levi 1914; Delone and Faddeev 1940; Gan et al. 2002], it is shown in [Bhargava et al. 2013] that the bijection sends a cubic R-ring S to any binary cubic form representing the cubic map $S/R \to \wedge^2_R(S/R)$, $s \mapsto s \wedge s^2$, which is functorial in R. □

If $\gamma \in \mathrm{GL}_2(F)$ and $f \in V(F)$, then $\mathrm{Disc}(\gamma f) = \det(\gamma)^2 \mathrm{Disc}(f)$. It follows that isomorphism classes of cubic F-algebras L of discriminant d are in bijection with $\mathrm{GL}_2(F)_{\pm 1}$-orbits on $V(F)_d$. Here, $\mathrm{GL}_2(F)_{\pm 1}$ is the subgroup of $\mathrm{GL}_2(F)$ consisting of elements with determinant ± 1. Since $\mathrm{SL}_2(F)$ has index 2 in $\mathrm{GL}_2(F)_{\pm 1}$, the $\mathrm{GL}_2(F)_{\pm 1}$-orbits break up into at most two $\mathrm{SL}_2(F)$-orbits. It is a fun exercise to show that there are exactly two orbits if and only if L is a field; the orbits are represented by $f(x, y)$ and $f(y, x)$.

Remark 3.3. The trivial class in $H^1(F, C_d)$ corresponds to the unique orbit of *reducible* forms of discriminant d. Hence, $\alpha \in H^1(F, C_d)$ is nontrivial if and only if the corresponding cubic algebra L is a field (if and only if L is generated over F by a root of $f(x, 1)$). The trivial class corresponds to $F \times E_d$, where $E_d = F[x]/(x^2 - d)$ is the quadratic F-algebra of discriminant d. Note that the trivial class is represented by $\frac{d}{4}x^3 + xy^2$, which is integral as long as d is.

From the functoriality in Proposition 3.2 applied to the base change $\mathcal{O}_F \hookrightarrow F$, we deduce:

Proposition 3.4. *Let $\alpha \in H^1(F, C_d)$, and let L be the corresponding cubic F-algebra. Then α is integral if and only if there is an $\mathcal{O}_F$-order $S \subset \mathcal{O}_L$ with $v(\mathrm{Disc}\ S) = v(d)$.*

Proof. The "only if" direction is clear. For the "if" direction, observe that any $\mathcal{O}_F$-order $S \subset \mathcal{O}_L$ has discriminant congruent to d modulo $F^{\times 2}$. So if $v(\mathrm{Disc}\ S) = v(d)$, then we see that $\mathrm{Disc}(S)$ is congruent to d modulo $\mathcal{O}_F^{\times 2}$. Thus we may choose bases so that S corresponds to a binary cubic form with coefficients in $\mathcal{O}_F$ of exact discriminant d. □

The following two facts about cubic orders will be useful [Bhargava et al. 2013, Propositions 15–16].

Proposition 3.5. *Let L be an étale cubic F-algebra. Suppose $f(x, y)$ corresponds to the maximal order $\mathcal{O}_L$ under the bijection of Proposition 3.2. Then the factorization type of $f(x, y)$ over the residue field $\mathbb{F}$ is the factorization type of the maximal ideal $\mathfrak{m}$ of $\mathcal{O}_F$ in the ring $\mathcal{O}_L$.*

Proposition 3.6. *Let $f(x, y) \in V(\mathcal{O}_F)_d$ correspond to a cubic ring S over $\mathcal{O}_F$. Then the sub-$\mathcal{O}_F$-rings $S' \subset S$ of index q correspond bijectively with the zeros of $f \pmod{\mathfrak{m}}$ in $\mathbb{P}^1(\mathbb{F})$.*

	$d \in F^{\times 2}$	$-3d \in F^{\times 2}$
$\zeta_3 \in F$	$\dim H^1(F, C_d) = 2$ $\dim H^1_{\mathrm{un}}(F, C_d) = 1$	
$\zeta_3 \notin F$	$\dim H^1(F, C_d) = 1$ $\dim H^1_{\mathrm{un}}(F, C_d) = 1$	$\dim H^1(F, C_d) = 1$ $\dim H^1_{\mathrm{un}}(F, C_d) = 0$

Table 1. Dimensions of $H^1(F, C_d)$ and $H^1_{\mathrm{un}}(F, C_d)$.

We also need the following result, which requires $\mathrm{char}\,\mathbb{F} \neq 3$, and which describes the subgroup $H^1_{\mathrm{un}}(F, C_d) \subset H^1(F, C_d)$ of unramified classes.

Proposition 3.7. *Suppose $0 \neq \alpha \in H^1(F, C_d)$ corresponds to the cubic extension L/F. Then $\alpha \in H^1_{\mathrm{un}}(F, C_d)$ if and only if L is unramified.*

Proof. Let $f_L \in V(F)$ be the corresponding binary cubic form. If L is unramified, then since f_L becomes reducible over L, the restriction of α to $H^1(L, C_d)$ is trivial. Thus, α is an unramified class. If L is ramified, then f_L remains irreducible over every unramified extension of F, and hence α is ramified. □

Lemma 3.8. *Assume the residue characteristic of F is not 3. Then*:

(i) $\dim H^1(F, C_d) = \dim H^0(F, C_d) + \dim H^0(F, C_{-3d})$.

(ii) $\dim H^1_{\mathrm{un}}(F, C_d) = \dim H^0(F, C_d)$.

When d has even valuation, these dimensions are computed in Table 1.

Proof. First note that C_d is Cartier dual to $C_{-3d} \simeq C_d \otimes \mu_3$. Since the residue characteristic is not 3, the Euler–Poincaré characteristic formula [Milne 1986, I.2.8] immediately gives (i). Let $I_F \subset G_F$ be the inertia group and let g be the Frobenius element of G_F/I_F. Then the groups $H^1_{\mathrm{un}}(F, C_d) \simeq H^0(I_F, C_d)/(g-1)H^0(I_F, C_d)$ and $H^0(I_F, C_d)[g-1] = H^0(F, C_d)$ have the same cardinality, which proves (ii). The table is computed using the fact that $\dim H^0(F, C_d) = 1$ if and only if $d \in F^{\times 2}$ and the dimension is 0 otherwise. □

The main result of this section is the following classification of the integral orbits in $V(F)_d$.

Theorem 3.9. *Let $\mathcal{O}_F$ be the ring of integers of a local field F with* $\mathrm{char}\,\mathcal{O}_F/\mathfrak{m} > 3$, *and let $d \in \mathcal{O}_F$ be nonzero.*

(a) *If $v(d) = 0$, then $H^1_{\mathrm{int}}(F, C_d) = H^1_{\mathrm{un}}(F, C_d)$.*

(b) *If $v(d)$ is odd, then $H^1_{\mathrm{int}}(F, C_d) = H^1(F, C_d) = 0$.*

(c) *If $v(d) = 2$, then the only nonintegral classes are the nontrivial unramified classes.*

(d) *If $v(d) > 2$, then all classes are integral.*

Proof. (a) This case follows from Propositions 3.4 and 3.7.

(b) We have $H^1(F, C_d) = 0$ by Lemma 3.8 (since $H^0(F, C_d) = 0$ whenever d has odd valuation). By Remark 3.3, the trivial class is integral.

(c) The ramified classes α correspond to totally ramified cubic extensions L/F. For such L we have $v(\mathrm{Disc}\, \mathcal{O}_L) = v(d)$, and hence these α are integral by Proposition 3.4. If $H^1(F, C_d)$ has a nontrivial unramified class α, then it corresponds to the unique unramified cubic extension L/F, which has unit discriminant. By Proposition 3.4, α is integral if and only if $\mathcal{O}_L$ has an order of index q. By Proposition 3.5, the binary form corresponding to $\mathcal{O}_L$ has no root over $\mathbb{F}$. So by Proposition 3.6, $\mathcal{O}_L$ has no order of index q. Hence the nontrivial unramified classes are indeed nonintegral.

(d) If $\alpha \in H^1(F, C_d)$ corresponds to a ramified cubic extension L/F, then $\mathcal{O}_L$ has discriminant of valuation 2. By Propositions 3.5 and 3.6, $\mathcal{O}_L$ has a unique order S of index q, and hence $v(\mathrm{Disc}\ S) = 4$. Note that if S_0 is a cubic $\mathcal{O}_F$-ring, then $S_0' = \mathcal{O}_F + \mathfrak{m}S_0$ is a subring of S of index q^2 and $\mathrm{Disc}(S_0') = q^4 \,\mathrm{Disc}(S_0)$. Thus, by considering the orders $\mathcal{O}_F + \mathfrak{m}^k\mathcal{O}_L$ and $\mathcal{O}_F + \mathfrak{m}^k S$, we can find an order $S' \subset \mathcal{O}_L$ with $v(\mathrm{Disc}\ S') = 2k$, for any $k \geq 1$.

Next let $\alpha \in H^1(F, C_d)$ be a nontrivial unramified class corresponding to the unramified cubic extension L/F. Then $v(\mathrm{Disc}(\mathcal{O}_F + \mathfrak{m}\mathcal{O}_L)) = 4$. By Proposition 3.6, there are $q + 1$ suborders $S' \subset \mathcal{O}_F + \mathfrak{m}\mathcal{O}_L$ of index q, so that $v(\mathrm{Disc}\ S') = 6$. As before, we deduce that there exists an order $S'' \subset \mathcal{O}_L$ with $v(\mathrm{Disc}\ S'') = 2k$, for any $k \geq 2$. □

4. Local Selmer conditions for ζ-linear isogenies

Let F be a finite extension of $\mathbb{Q}_p$ with surjective discrete valuation v, ring of integers $\mathcal{O}_F$, uniformizer ϖ, and residue field $\mathbb{F}$. Let $m \geq 1$, $n = 3^m$, and let $\zeta = \zeta_n$ be a primitive n-th root of unity. Let A, B be abelian varieties over F that admit ζ-multiplication.

Let $\phi : A \to B$ be a ζ-linear 3-isogeny over F, as defined in Section 2B. For each $d \in F^\times$, we consider the 3-isogeny $\phi_d : A_d \to B_d$, as in Section 2C. We will assume in this section that $A[\phi](F) \neq 0$, that is, that $A[\phi]$ is generated by a rational point. This can always be achieved by replacing ϕ with an appropriate twist, so there is no loss in generality. We also assume that $0 \leq v(d) < 2n$, again with no loss in generality.

The group $H^1(F, A_d[\phi_d])$ is a finite-dimensional $\mathbb{F}_3$-vector space. In fact, if $\chi_d : G_F \to \mathbb{F}_3^\times$ is the quadratic character cutting out $F(\sqrt{d})$, then $A_d[\phi_d] \simeq A[\phi] \otimes \chi_d$ is isomorphic to C_d from Section 3. Thus, $H^1(F, A_d[\phi_d]) \simeq H^1(F, C_d)$. The exact sequence

$$0 \to A_d[\phi_d] \to A_d \to B_d \to 0$$

induces a Kummer map

$$\partial_d : B_d(F) \to H^1(F, A_d[\phi_d]).$$

We call its image $\mathrm{im}(\partial_d) \subset H^1(F, A_d[\phi_d])$ the subgroup of *soluble classes*. The goal of this section is to prove the following theorem describing the soluble classes and to compute the local Selmer ratios of ϕ_d.

Theorem 4.1. *Assume that* $\operatorname{char}\mathbb{F} \neq 3$, *that* A *has good reduction, and that* $A[\phi](F) \neq 0$. *Suppose that* $0 \leq v(d) < 2n$.

(a) *If* $v(d) = 0$, *then* $\operatorname{im}(\partial_d) = H^1_{\mathrm{un}}(F, A_d[\phi_d])$.

(b) *If* $v(d)$ *is odd, or more generally, if* $F(\sqrt{d})/F$ *is ramified, then* $\operatorname{im}(\partial_d) = H^1(F, A_d[\phi_d]) = 0$.

(c) *If* $v(d) > 0$ *is even and* $d \notin F^{\times 2}$, *then* $H^1_{\mathrm{un}}(F, A_d[\phi_d]) = 0$.

(d) *If* $v(d) > 0$ *is even and* $d \in F^{\times 2}$, *then write* $d = \varpi^{v(d)}u$ *with* $u \in \mathcal{O}_F^\times$ *and let* $s = \gcd(3^m, v(d))$. *Then* $\operatorname{im}(\partial_d) \cap H^1_{\mathrm{un}}(F, A_d[\phi_d]) = 0$ *if and only if* $A_u[\phi]$ *is a direct summand of* $A_u[\pi^s]$.

We retain these assumptions on A and $\mathbb{F}$ for the remainder of this section.

4A. ***Nonsquare and unramified twists.*** We first prove parts (a)–(c).

Proof of Theorem 4.1(a)–(c). For (a), assume at first that $\operatorname{char}\mathbb{F} > 3$. Then $F(d^{1/2n})$ is unramified over F, since $v(d) = 0$ and $(\operatorname{char}\mathbb{F}, 2n) = 1$. Since A_d is isomorphic to A over $F(d^{1/2n})$, it has good reduction over an unramified extension, and hence has good reduction already over F. Since A_d has good reduction and $\operatorname{char}\mathbb{F} \nmid \deg(\phi_d)$, the image of the Kummer map ∂_d is exactly the unramified classes [Česnavičius 2016, Proposition 2.7(d)]. The proof just given works even when $\operatorname{char}\mathbb{F} = 2$, as long as $F(\sqrt{d})/F$ is unramified. When this extension is ramified, the result follows from (b). Part (b) itself follows from Theorem 3.9(b). The case $\operatorname{char}\mathbb{F} = 2$ was not dealt with there, but the proof is identical.

For (c), we have $H^1_{\mathrm{un}}(F, A_d[\phi_d]) = H^1_{\mathrm{un}}(F, C_d) = 0$ by Lemma 3.8, since $H^0(F, C_d) = 0$ whenever $d \notin F^{\times 2}$. □

4B. ***Twists of positive valuation.*** The proof of Theorem 4.1(d) will take more work. Indeed, we will compute more generally the size of $\operatorname{im}(\partial_d)$ for all d (including $d \notin F^{\times 2}$) such that $v(d)$ is even and positive. Let $s = \gcd(v(d), 3^m)$, and write $d = \varpi^{v(d)}u$ for $u \in \mathcal{O}_F^\times$. Recall the map $\pi^s : A \to A^{(s)}$ defined in Section 2, and let $\psi^s : B \to A^{(s)}$ be the isogeny such that $\psi^s \circ \phi = \pi^s$.

4B1. *Extension classes.* For each $t \in F^\times$, let κ_t^s be the extension class corresponding to the short exact sequence

$$0 \to A_t[\phi_t] \to A_t[\pi_t^s] \xrightarrow{\phi_t} B_t[\psi_t^s] \to 0. \tag{4-1}$$

Thus $\kappa_t^s = 0$ if and only if $A_t[\phi_t]$ is a direct summand of $A_t[\pi_t^s]$ as a G_F-module. Similarly, let $\hat{\kappa}_t^s$ be the class of the extension

$$0 \to B_t[\psi_t^s] \to B_t[\pi_t^s] \xrightarrow{\psi_t^s} A_t^{(s)}[\phi_t^{(s)}] \to 0. \tag{4-2}$$

Thus, $\hat{\kappa}_t^s = 0$ if and only if $B_t[\psi_t^s]$ is a direct summand of $B_t[\pi_t^s]$.

Remark 4.2. By Lemma 2.4, the cocycle $\hat{\kappa}_t^s$ is equal to the class of the extension

$$0 \to B_t^{(-s)}[\psi_t^{-s}] \to B_t^{(-s)}[\pi_t^s] \to A_t[\phi_t] \to 0.$$

	$d \in F^{\times 2}$	$-3d \in F^{\times 2}$	$d, -3d \notin F^{\times 2}$
$\zeta_3 \in F$	$\dim H^1(F, A_d[\phi_d]) = 2$ $\dim H^1_{\mathrm{un}}(F, A_d[\phi_d]) = 1$		$\dim H^1(F, A_d[\phi_d]) = 0$ $\dim H^1_{\mathrm{un}}(F, A_d[\phi_d]) = 0$
$\zeta_3 \notin F$	$\dim H^1(F, A_d[\phi_d]) = 1$ $\dim H^1_{\mathrm{un}}(F, A_d[\phi_d]) = 1$	$\dim H^1(F, A_d[\phi_d]) = 1$ $\dim H^1_{\mathrm{un}}(F, A_d[\phi_d]) = 0$	$\dim H^1(F, A_d[\phi_d]) = 0$ $\dim H^1_{\mathrm{un}}(F, A_d[\phi_d]) = 0$

Table 2. Dimensions of $H^1(F, A_d[\phi_d])$ and $H^1_{\mathrm{un}}(F, A_d[\phi_d])$.

Here, we have $\phi \circ \psi_t^{-s} = \pi^s : B_t^{(-s)} \to B_t$. By duality, $\hat{\kappa}_t^s$ is the class of the extension

$$0 \to \widehat{B}_t[\hat{\phi}_t] \to \widehat{B}_t[\hat{\pi}_t^s] \to \hat{A}_t[\hat{\psi}_t^{-s}] \to 0.$$

Thus $\hat{\kappa}_t^s = 0$ if and only if $\widehat{B}_t[\hat{\phi}_t]$ is a direct summand of $\widehat{B}_t[\hat{\pi}_t^s]$, which explains the notation.

Let $|\kappa_u^s|$ and $|\hat{\kappa}_u^s|$ denote the orders of the classes κ_u^s and $\hat{\kappa}_u^s$ in their respective Ext-groups.

4B2. *The image of ∂_d.* The following theorem relates the size of $\operatorname{im} \partial_d$ to $|\kappa_u^s|$ and $|\hat{\kappa}_u^s|$ and will finish the proof of Theorem 4.1(d).

Theorem 4.3. *Assume that $v(d)$ is even and positive, and write $d = \varpi^{v(d)} u$ with $u \in \mathcal{O}_F^\times$.*

(i) $$\#\operatorname{im} \partial_d \cap H^1_{\mathrm{un}}(F, A_d[\phi_d]) = \begin{cases} |\kappa_u^s| & \text{if } d \in F^{\times 2}, \\ 1 & \text{otherwise.} \end{cases}$$

(ii) $$\#\left(\frac{\operatorname{im} \partial_d}{\operatorname{im} \partial_d \cap H^1_{\mathrm{un}}(F, A_d[\phi_d])} \right) = \begin{cases} \dfrac{3}{|\hat{\kappa}_u^s|} & \text{if } -3d \in F^{\times 2}, \\ 1 & \text{otherwise.} \end{cases}$$

In particular, if $d \in F^{\times 2}$, then $\operatorname{im} \partial_d \cap H^1_{\mathrm{un}}(F, A_d[\phi_d]) = 0$ if and only if $A_u[\phi_u]$ is a direct summand of $A_u[\pi_u^s]$.

The proof of Theorem 4.3 requires several preliminary results.

Lemma 4.4. *The dimensions of $H^1(F, A_d[\phi_d])$ and $H^1_{\mathrm{un}}(F, A_d[\phi_d])$ are as in Table 2.*

Proof. Since $H^1(F, A_d[\phi_d]) \simeq H^1(F, C_d)$ and $H^1_{\mathrm{un}}(F, A_d[\phi_d]) \simeq H^1_{\mathrm{un}}(F, C_d)$, the dimensions in Table 2 are identical to those in Table 1. The bottom right cell is only relevant when $\operatorname{char} \mathbb{F} = 2$, and follows from Lemma 3.8(i). □

Lemma 4.5. *If $t \in F^{\times 2s}$, then $A_t[\pi_t^s] \simeq A[\pi^s]$ and $B_t[\pi_t^s] \simeq B[\pi^s]$ as G_F-modules.*

Proof. By the definition of A_t, there is an isomorphism $\phi : A \to A_t$ over $F(t^{1/2n})$ such that if $P \in A[\pi^s]$ and $\sigma \in G_F$, then in $A_t[\pi_t^s]$, we have

$$(\phi(P))^\sigma = \frac{\sigma(t^{1/2n})}{t^{1/2n}} \phi(P^\sigma).$$

If $t \in F^{\times 2s}$, then $\frac{\sigma(t^{1/2n})}{t^{1/2n}} \in \langle \zeta^s \rangle$. Since ζ^s acts as the identity on $A[\pi^s]$, we see that $A_t[\pi_t^s] \simeq A[\pi^s]$. The proof that $B_t[\pi_t^s] \simeq B[\pi^s]$ is identical. □

Remark 4.6. In particular, if $s = 1$, the condition in Theorem 4.1(d) is simply that $A[\phi]$ is a direct summand of $A[\pi]$, which is independent of u.

Corollary 4.7. (i) *There is an isomorphism between*

$$0 \to A_u[\phi_u] \to A_u[\pi_u^s] \xrightarrow{\phi_u} B_u[\psi_u^s] \to 0$$

and

$$0 \to A_d[\phi_d] \to A_d[\pi_d^s] \xrightarrow{\phi_d} B_d[\psi_d^s] \to 0$$

as short exact sequences of G_F-modules.

(ii) *There is an isomorphism between*

$$0 \to B_u[\psi_u^s] \to B_u[\pi_u^s] \xrightarrow{\psi_u^s} A_u^{(s)}[\phi_u^{(s)}] \to 0$$

and

$$0 \to B_d[\psi_d^s] \to B_d[\pi_d^s] \xrightarrow{\psi_d^s} A_d^{(s)}[\phi_d^{(s)}] \to 0$$

as short exact sequences of G_F-modules.

Proof. The first claim follows from Lemma 4.5 together with the observation that the isomorphism $A_u[\pi_u^s] \to A_d[\pi_d^s]$ restricts to an isomorphism $A_u[\phi_u] \to A_d[\phi_d]$. The second claim follows similarly. □

Lemma 4.8. *Suppose that $v(d) = 2a \cdot 3^r$ is even and positive, and let k be an unramified extension of F. For $X \in \{A, B\}$, we have $X_d[3](k) = X_d[\pi_d^{3^r}](k)$.*

Proof. Let F_{un} be the maximal unramified extension of F, let $L = F_{\mathrm{un}}(\sqrt{d})$, and let $M = L(d^{1/n}) = F_{\mathrm{un}}(d^{1/2n})$. Since $k \subset L$, it suffices to show that $X_d[3](L) = X_d[\pi_d^{3^r}](L)$.

The extension M/L is tamely ramified of order 3^{n-r}, and we can choose a generator τ of $\mathrm{Gal}(M/L)$ so that $\tau(d^{1/2n})/d^{1/2n} = \zeta^{3^r}$. Since X has good reduction, we have $X[3] \subset X(L)$. Since X and X_d are isomorphic over $L(d^{1/n})$, it follows that $X_d[3] \subset X(M)$.

Now, if $\phi : X_d \to X$ is an isomorphism over M and $P \in X_d(M)$, then by definition,

$$\phi(P)^\tau = \tau(d^{1/2n})/d^{1/2n}\phi(P^\tau) = \zeta^{3^r}\phi(P^\tau).$$

Hence, if $P \in X_d[3](L)$, then $\phi(P) \in X[3] \subset X[3](L)$, so $\phi(P) = \zeta^{3^r}\phi(P)$. It follows that

$$0 = (1 - \zeta^{3^r})\phi(P) = (1 - \zeta)^{3^r}\phi(P) = \pi^{3^r}\phi(P) = \phi(\pi_d^{3^r} P),$$

where the second equality uses the fact that P is a 3-torsion point. It follows that $P \in X_d[\pi_d^{3^r}]$, so $X_d[\pi_d^{3^r}](L) = X_d[3](L)$. □

From (4-1), we obtain a long exact sequence

$$0 \to A_d[\phi_d](F) \to A_d[\pi_d^s](F) \xrightarrow{\phi_d} B_d[\psi_d^s](F) \xrightarrow{\delta_d} H^1(F, A_d[\phi_d]).$$

Similarly, from (4-2), there is a long exact sequence

$$0 \to B_d[\psi_d^s](F) \to B_d[\pi_d^s](F) \xrightarrow{\psi_d^s} A_d^{(s)}[\phi_d^{(s)}](F) \xrightarrow{\hat{\delta}_d} H^1(F, B_d[\psi_d^s]). \tag{4-3}$$

Clearly $\operatorname{im}\delta_d \subset \operatorname{im}\partial_d$, and ∂_d induces an injective map

$$\frac{B_d[\pi_d^s](F)}{B_d[\psi_d^s](F)} \hookrightarrow \frac{\operatorname{im}\partial_d}{\operatorname{im}\delta_d}.$$

Proposition 4.9. *We have*

$$\operatorname{im}\delta_d = \operatorname{im}\partial_d \cap H^1_{\mathrm{un}}(F, A_d[\phi_d]).$$

Proof. Consider the Kummer exact sequence

$$0 \to A_d[\phi_d](F) \to A_d(F) \xrightarrow{\phi_d} B_d(F) \xrightarrow{\partial_d} H^1(F, A_d[\phi_d]).$$

Since $\operatorname{char}\mathbb{F} \neq 3$, we have $B_d(F)/\phi_d A_d(F) = B_d[3^\infty](F)/\phi_d A_d[3^\infty](F)$, so ∂_d depends only on the exact sequence

$$0 \to A_d[\phi_d](F) \to A_d[3^\infty](F) \xrightarrow{\phi_d} B_d[3^\infty](F) \xrightarrow{\partial_d} H^1(F, A_d[\phi_d]).$$

By Lemma 4.8, this exact sequence is the same as

$$0 \to A_d[\phi_d](F) \to A_d[\pi_d^s](F) \xrightarrow{\phi_d} B_d[\pi_d^s](F) \xrightarrow{\partial_d} H^1(F, A_d[\phi_d]).$$

By Corollary 4.7, the long exact sequence

$$0 \to A_d[\phi_d](F) \to A_d[\pi_d^s](F) \xrightarrow{\phi_d} B_d[\psi_d^s](F) \xrightarrow{\delta_d} H^1(F, A_d[\phi_d])$$

is isomorphic to the long exact sequence

$$0 \to A_u[\phi_u](F) \to A_u[\pi_u^s](F) \xrightarrow{\phi_u} B_u[\psi_u^s](F) \xrightarrow{\delta_u} H^1(F, A_u[\phi_u]).$$

Since A_u has good reduction, the image of δ_u is contained in $H^1_{\mathrm{un}}(F, A_u[\phi_u])$ [Česnavičius 2016, Proposition 2.7(d)]. Thus the image of δ_d is contained in both $H^1_{\mathrm{un}}(F, A_d[\phi_d])$ and $\operatorname{im}\partial_d$.

Conversely, if $y \in B_d[\pi_d^s](F) \setminus B_d[\psi_d^s](F)$, then $\partial_d(y)$ is the cocycle $\sigma \mapsto x^\sigma - x$, where $\phi_d(x) = y$. But $x \in A_d[3] \setminus A_d[\pi_d^s]$, so by Lemma 4.8, x cannot be defined over an unramified extension of F. It follows that $\partial_d(y) \notin H^1_{\mathrm{un}}(F, A_d[\phi_d])$. Hence $\operatorname{im}\delta_d \supset \operatorname{im}\partial_d \cap H^1_{\mathrm{un}}(F, A_d[\phi_d])$. □

We next relate the sizes of the images of δ_d and $\hat{\delta}_d$ to the sizes of $|\kappa_u^s|$ and $|\hat{\kappa}_u^s|$.

Lemma 4.10. *If $d \in F^{\times 2}$, then $\#\operatorname{im}\delta_d = |\kappa_u^s|$. Similarly, if $-3d \in F^{\times 2}$, then $\#\operatorname{im}\hat{\delta}_d = |\hat{\kappa}_u^s|$.*

Proof. By duality, we have $A_u[\phi_u] \simeq \widehat{B}_u[\hat{\phi}_u] \otimes \chi_3$. Thus, the two claims of the lemma are in fact equivalent to each other; see Remark 4.2. We prove the second one. Since $-3d \in F^{\times 2}$, we have $A_u^{(s)}[\phi_u^{(s)}] = \mathbb{F}_3$ as a G_F-module. Moreover, by definition we have $\hat{\delta}_d(1) = \hat{\kappa}_d^s$. By Corollary 4.7, $|\hat{\kappa}_d^s| = |\hat{\kappa}_u^s|$. It follows that $|\hat{\kappa}_u^s| = \#\operatorname{im}\hat{\delta}_d$. □

Proof of Theorem 4.3. If $d \in F^{\times 2}$, then part (i) follows immediately from Proposition 4.9 and Lemma 4.10. If $d \notin F^{\times 2}$, then $H^1_{\mathrm{un}}(F, A_d[\phi_d]) = 0$ by Table 2.

	$d \in F^{\times 2}$	$-3d \in F^{\times 2}$	$d, -3d \notin F^{\times 2}$
$\zeta_3 \in F$	$\dfrac{\lvert\kappa_u^s\rvert}{\lvert\hat{\kappa}_u^s\rvert}$		1
$\zeta_3 \notin F$	$\dfrac{\lvert\kappa_u^s\rvert}{3}$	$\dfrac{3}{\lvert\hat{\kappa}_u^s\rvert}$	1

Table 3. Values of $c(\phi_d)$ over local fields F, when $v(d)$ is even and positive.

By Lemma 4.8 and Proposition 4.9,

$$\frac{\operatorname{im}\partial_d}{\operatorname{im}\partial_d \cap H^1_{\mathrm{un}}(F, A_d[\phi_d])}$$

is isomorphic to the image of the injective map

$$\frac{B_d[\pi_d^s](F)}{B_d[\psi_d^s](F)} \to \frac{\operatorname{im}\partial_d}{\operatorname{im}\delta_d}$$

induced by ∂_d. From (4-3), we have

$$\#\left(\frac{B_d[\pi_d^s](F)}{B_d[\psi_d^s](F)}\right) = \frac{\#A_d^{(s)}[\phi_d^{(s)}](F)}{\#\operatorname{im}\hat{\delta}_d}.$$

If $-3d \in F^{\times 2}$, this is $\frac{3}{|\hat{\kappa}_u^s|}$ by Lemma 4.10. If $-3d \notin F^{\times 2}$, then $H^1(F, A_d[\phi_d]) = H^1_{\mathrm{un}}(F, A_d[\phi_d])$ by Table 2, so the result follows from Proposition 4.9. □

4C. ***Local Selmer ratios.*** For applications, we record the local Selmer ratios

$$c(\phi_d) = \frac{\#\operatorname{coker}(A_d(F) \to B_d(F))}{\#\ker(A_d(F) \to B_d(F))} = \frac{\#\operatorname{im}(\partial_d)}{\#A_d[\phi_d](F)},$$

which have implicitly been computed in the previous subsection.

Theorem 4.11. *Assume that* $\operatorname{char}\mathbb{F} \neq 3$, *that A has good reduction, and that $A[\phi](F) = \mathbb{Z}/3\mathbb{Z}$. Then $c(\phi_d) = 1$ unless $v(d)$ is even and positive. If $v(d)$ is even and positive, write $d = \varpi^{v(d)}u$ with $u \in \mathcal{O}_F^\times$, and let $s = \gcd(3^m, v(d))$. Let κ_u^s and $\hat{\kappa}_u^s$ be the classes defined in the previous section, and write $|\kappa_u^s|$ and $|\hat{\kappa}_u^s|$ for their orders. Then $c(\phi_d)$ is as in Table 3.*

Proof. If neither d nor $-3d$ is a square in F, then $\#A_d[\phi_d](F) = 1$, and by Table 2, $\#\operatorname{im}(\partial_d) = 1$, so $c(\phi_d) = 1$ as claimed. Henceforth, we assume that either d or $-3d$ is a square in F. When $v(d) = 0$, then A_d has good reduction, so by [Shnidman 2021, Proposition 3.1]

$$c(\phi_d) = c(B_d)/c(A_d) = 1,$$

where $c(A_d)$ and $C(B_d)$ are the Tamagawa numbers of A_d and B_d. When $v(d)$ is odd, then $\operatorname{im}(\partial_d) = 0$ (Theorem 4.1) and $A_d[\phi_d](F) = 0$, so again $c(\phi_d) = 1$. So it remains to compute $c(\phi_d)$ when $v(d)$ is even and positive. This is done by combining the formula for $\#\operatorname{im}\partial_d$ in Theorem 4.3 with the fact that $\#A_d[\phi_d](F) = 3$ if d is a square and 1 otherwise. The result of this computation is Table 3. □

5. Selmer groups and integrality

We return to the global setting of the introduction, so that F is a number field and $n = 3^m$ for some $m \geq 1$. Recall that $\zeta \in \overline{F}$ is a primitive n-th root of unity. Let $\phi : A \to B$ be a ζ-linear 3-isogeny over F. Recall from Section 2 that there is a twist $A^{(1)}$ of A, and an isogeny $\pi : A \to A^{(1)}$, which becomes isomorphic to the endomorphism $1 - \zeta$ over $F(\zeta)$.

For any F-algebra K, define $\mathbb{B}(K) = K^\times / K^{\times 2n}$. The notation is meant to suggest that $\mathbb{B}$ is the classifying stack $B\mu_{2n}$. For $d \in \mathbb{B}(F)$, let $\phi_d : A_d \to B_d$ be a twist of ϕ corresponding to

$$d \in F^\times / F^{\times 2n} \simeq H^1(F, \mu_{2n}) \to H^1(F, \mathrm{Aut}_{\overline{F}}(\phi)),$$

as in Section 2C. The Selmer group $\mathrm{Sel}_{\phi_d}(A_d)$ is the subgroup of $H^1(F, A_d[\phi_d])$ consisting of classes whose restriction lies in the image of the Kummer map

$$\partial_{d,v} : B_d(F_v)/\phi_d(A_d(F_v)) \hookrightarrow H^1(F_v, A_d[\phi_d])$$

for all places v of F. Sometimes we use the notation $\mathrm{Sel}(\phi_d)$ instead of the more clunky $\mathrm{Sel}_{\phi_d}(A_d)$.

The goal of this section is to compute the average size of $\mathrm{Sel}_{\phi_d}(A_d)$ as d varies. The idea is to view Selmer elements as $\mathrm{SL}_2(F)$-orbits of binary cubic forms and then apply geometry-of-numbers counting techniques. To carry this out, we must show that the orbits corresponding to Selmer elements have representatives with bounded denominator. In fact, we will show that this boundedness only holds if $A[\phi]$ is almost everywhere locally a direct summand of $A[\pi]$.

5A. ***Integrality of Selmer elements.*** We assume for simplicity that $A[\phi] \simeq \mathbb{Z}/3\mathbb{Z}$ as group schemes. This is not really a constraint, since there is always a quadratic twist of A with this property. This assumption implies that $A_d[\phi_d] \simeq C_d$, where C_d is the order 3 group scheme cut out by the quadratic field of discriminant d, from Section 3.

Recall the space V of binary cubic forms from Section 3. Recall also the set $V(F)_d$ of cubic forms of discriminant d, whose $\mathrm{SL}_2(F)$-orbits are in bijection with $H^1(F, C_d) \simeq H^1(F, A_d[\phi_d])$. Similarly, for each place v of F, there is a bijection between $\mathrm{SL}_2(F_v)$-orbits on the set $V(F_v)_d$ and $H^1(F_v, A_d[\phi_d])$. Let $V(F_v)_d^{\mathrm{sol}}$ denote the subset of $V(F_v)_d$ corresponding to classes $\alpha \in H^1(F_v, A_d[\phi_d])$ in the image of $\partial_{d,v}$. Similarly, let $V(F)_d^{\mathrm{sel}}$ denote the subset of $V(F)_d$ corresponding to classes in $\mathrm{Sel}_{\phi_d}(A_d)$. Define

$$V(F)^{\mathrm{sel}} = \bigcup_{0 \neq d \in \mathcal{O}_F} V(F)_d^{\mathrm{sel}} \quad \text{and} \quad V(F_v)^{\mathrm{sol}} = \bigcup_{d \in \mathcal{O}_v(2n)} V(F_v)_d^{\mathrm{sol}},$$

where $\mathcal{O}_v$ is the ring of integers in F_v, and $\mathcal{O}_v(2n) = \{d \in \mathcal{O}_v : v(d) < 2n\}$. Similarly, define

$$V(F)_{\mathrm{sq.free}}^{\mathrm{sel}} = \bigcup_{0 \neq d \in \mathcal{O}_F \text{ sq.free}} V(F)_d^{\mathrm{sel}} \quad \text{and} \quad V(F_v)_{\mathrm{sq.free}}^{\mathrm{sol}} = \bigcup_{d \in \mathcal{O}_v(2)} V(F_v)_d^{\mathrm{sol}},$$

where $\mathcal{O}_v(2) = \{d \in \mathcal{O}_v : v(d) < 2\}$ and the union on the left runs over elements $d \in \mathcal{O}_F$ that are squarefree. Of course, these sets depend on the initial choices of A and ϕ.

In order to count the $\mathrm{SL}_2(F)$-orbits on $V(F)^{\mathrm{sel}}$ and $V(F)^{\mathrm{sel}}_{\mathrm{sq.free}}$ of bounded discriminant, we wish to prove that these orbits have representatives in $V(\mathcal{O}_F)$, or at least representatives with denominators that are uniformly bounded. Since SL_2 has class number 1, this is ultimately a local question, and it is enough to prove that for almost all v, each $\mathrm{SL}_2(F_v)$-orbit in $V(F_v)^{\mathrm{sol}}$ and $V(F_v)^{\mathrm{sol}}_{\mathrm{sq.free}}$ has a representative in $V(\mathcal{O}_v)$. For $V(F)^{\mathrm{sel}}_{\mathrm{sq.free}}$, this integrality holds without any conditions. However, for $V(F)^{\mathrm{sel}}$, we are forced to assume that $A[\phi]$ is almost everywhere locally a direct summand of $A[\pi]$, as defined in the introduction.

The following integrality result is crucial for the proofs of Theorems 5.2 and 5.3 below.

Theorem 5.1. *Let $v \nmid 6\infty$ be a place of F at which A has good reduction.*

(i) *Every element of $V(F_v)^{\mathrm{sol}}_{\mathrm{sq.free}}$ is $\mathrm{SL}_2(F_v)$-equivalent to an element of $V(\mathcal{O}_v)$.*

(ii) *If $A[\phi]$ is a direct summand of $A[\pi]$ as a G_{F_v}-module, then every element of $V(F_v)^{\mathrm{sol}}$ is $\mathrm{SL}_2(F_v)$-equivalent to an element of $V(\mathcal{O}_v)$.*

Proof. For each $d \in \mathcal{O}_v$, we must show that each class of $H^1(F_v, A_d[\phi_d]) \simeq H^1(F_v, C_d)$ that lies in the image of $\partial_{d,v}$ corresponds to an integral orbit of discriminant d. This follows from a comparison of Theorem 3.9 with Theorem 4.1 and Remark 4.6. □

5B. ***Average size of the Selmer group.*** There is a natural height function on $\mathbb{B}(F)$ defined as follows. Let M_∞ be the set of archimedean places of F. If $d \in \mathbb{B}(F)$ with lift $d_0 \in F^\times$, then define the ideal $I = \{a \in F : a^{2n} d_0 \in \mathcal{O}_F\}$. The height of d is then

$$h(d) = \mathrm{Nm}(I)^{2n} \prod_{v \in M_\infty} |d_0|_v.$$

This is independent of the lift d_0, by the product formula. If $F = \mathbb{Q}$, then $h(d) = |d_0|$, where d_0 is the unique $2n$-th power free integer representing d. For any $X > 0$, there are finitely many $d \in \mathbb{B}(F)$ with $h(d) < X$.

In order to state a robust version of Theorems 1.8 and 1.9, we recall from [Bhargava et al. 2020] the notion of functions on F that are defined by local conditions. Let $F_\infty = \prod_{v \in M_\infty} F_v$. We say a function $\psi : F \to [0, 1]$ is *defined by local congruence conditions* if there exist local functions $\psi_v : F_v \to [0, 1]$ for every finite place v of F, and a function $\psi_\infty : F_\infty \to [0, 1]$, such that the following two conditions hold:

(1) For all $w \in F$, the product $\psi_\infty(w) \prod_{v \notin M_\infty} \psi_v(w)$ converges to $\psi(w)$.

(2) For each finite place v, and for $v = \infty$, the function ψ_v is nonzero on some open set and locally constant outside some closed subset of F_v of measure 0.

A subset of F is said to be *defined by local congruence conditions* if its characteristic function is defined by local congruence conditions.

Let Σ_0 be a fundamental domain for the action of $F^\times$ on F defined by $\alpha \cdot \beta = \alpha^{2n}\beta$. We may take Σ_0 so that it is defined by local congruence conditions. For any $X > 0$, let F_X denote the set of $d \in F^\times$ such that $h(d) < X$. Then $\Sigma_0 \cap F_X$ is finite and we think of the abelian varieties A_d as elements of Σ_0, so that the set of all A_d, with $d \in \mathbb{B}(F)$ and $h(d) < X$, is naturally in bijection with the finite set $\Sigma_0 \cap F_X$.

A *family of twists* $\{A_d\}$ *defined by local congruence conditions* is then a subset $\Sigma \subset \Sigma_0$ defined by local congruence conditions. In that case, the characteristic function χ_Σ of Σ factors as

$$\chi_\Sigma = \chi_{\Sigma,\infty} \prod_{v \notin M_\infty} \chi_{\Sigma_v}.$$

For each finite place v of F, let Σ_v be the subset of F_v whose characteristic function is χ_{Σ_v}, and let Σ_∞ be the subset of F_∞ whose characteristic function is $\chi_{\Sigma,\infty}$.

We say that Σ is *large* if Σ_v contains the set $\mathcal{O}_v(2) := \{d \in \mathcal{O}_v : v(d) < 2\}$ for all but finitely many finite places v, and if Σ_∞ is a nonempty union of cosets in $\mathbb{B}(F_\infty)$. By construction, we have $\Sigma_{0,v} = \mathcal{O}_v(2n) \supset \mathcal{O}_v(2)$ for all finite v, so Σ_0 is itself large.

If f is a positive function on $\mathbb{B}(F)$ and $\Sigma \subset \Sigma_0$, we write $\Sigma(X) = \Sigma \cap F_X$ and define

$$\operatorname{avg}_\Sigma f(d) = \lim_{X\to\infty} \frac{1}{\#\Sigma(X)} \sum_{d\in\Sigma(X)} f(d).$$

Our formula for $\operatorname{avg}_\Sigma \operatorname{Sel}_{\phi_d}(A_d)$ is most neatly formulated in terms of the global Selmer ratios $c(\phi_d)$ defined in the introduction, and first defined in [Bhargava et al. 2020].

Theorem 5.2. *Let Σ be a large family of twists A_d. For each $k \in \mathbb{Z}$, let $T_k \subset \Sigma$ be the subset of $d \in \Sigma$ such that $c(\phi_d) = 3^k$. Assume either that $A[\phi]$ is almost everywhere locally a direct summand of $A[\pi]$ or that $\Sigma_v = \mathcal{O}_v(2)$ for all but finitely many v. If T_k is nonempty, then*

$$\operatorname*{avg}_{d\in T_k} \#\operatorname{Sel}_{\phi_d}(A_d) = 1 + 3^k.$$

When they are nonempty, the sets T_k are countable disjoint unions of large sets. Using the uniformity estimate [Bhargava et al. 2015, Theorem 17] and copying the argument from [Bhargava et al. 2020, pp. 319–320], we deduce Theorem 5.2 from the following result. To state it, we define for any $d = (d_v)_{v\in M_\infty} \in \mathbb{B}(F_\infty)$,

$$c_\infty(\phi_d) := \prod_{v\in M_\infty} c_v(\phi_{d_v}).$$

We also let $\overline{\Sigma}_\infty$ denote the image of Σ_∞ in the finite group $\mathbb{B}(F_\infty)$.

Theorem 5.3. *Let Σ be a large family of twists A_d. Assume either that $A[\phi]$ is almost everywhere locally a direct summand of $A[\pi]$ or that $\Sigma_v = \mathcal{O}_v(2)$ for all but finitely many v. Then*

$$\operatorname{avg}_\Sigma \#\operatorname{Sel}_{\phi_d}(A_d) = 1 + \frac{\displaystyle\int_{d\in\overline{\Sigma}_\infty} c_\infty(\phi_d)\mu_\infty(d)}{\displaystyle\int_{d\in\overline{\Sigma}_\infty} \mu_\infty(d)} \prod_{v\notin M_\infty} \frac{\displaystyle\int_{d\in\Sigma_v} c_v(\phi_d)\mu_v(d)}{\displaystyle\int_{d\in\Sigma_v} \mu_v(d)},$$

where μ_v is any Haar measure on $\mathcal{O}_v$ and μ_∞ is the uniform measure on $\mathbb{B}(F_\infty)$.

Proof. Assume at first that $A[\phi](F) \simeq \mathbb{Z}/3\mathbb{Z}$, as in Section 5A. Then the left-hand side is equal to the limit as $X \to \infty$ of the average number of $\mathrm{SL}_2(F)$-orbits of discriminant d in $V(F)^{\mathrm{sel}}$, for d in the finite set $\Sigma \cap F_X$. Given our key result Theorem 5.1, Theorem 5.3 now follows exactly as in the proof of [Bhargava et al. 2020, Theorem 11], and we refer the reader there for the details. Note that two proofs were given in that paper, one in the case $F = \mathbb{Q}$ and one for general number fields F. The proof over general fields F relies on the geometry-of-numbers machinery developed in the preprint [Bhargava et al. $\geq$ 2025], which has still not appeared. One can alternatively make use of the techniques in [Bhargava et al. 2015], again using Theorem 5.1 as a key input, to deduce the formula for general F. Note that in [Bhargava et al. 2015], the authors count cubic extensions of F with prescribed local conditions (ordered by discriminant), exactly by counting *integral* $\mathrm{SL}_2(F)$-orbits of binary cubic forms with prescribed local conditions.

The proof in the general case where $A[\phi](F) \neq \mathbb{Z}/3\mathbb{Z}$ is exactly the same except we identify elements of $\mathrm{Sel}_{\phi_d}(A_d)$ with orbits of binary cubic forms of discriminant dk, where $k \in \mathcal{O}_F$ is chosen so that $F(\sqrt{k}) = F(A[\phi])$, as is done in [Bhargava et al. 2019, §4]. □

Taking $\Sigma = \Sigma_0$ in Theorem 5.3, we deduce Theorem 1.8, and taking $\Sigma = \{d : d \in \mathcal{O}_v(2)\ \forall v\}$, we deduce Theorem 1.9.

5C. ***Explicit Selmer rank bounds.*** The following consequence of Theorem 5.2 will be helpful in giving explicit average rank bounds.

Proposition 5.4. *Let $\phi : A \to B$, Σ, and T_k be as in Theorem 5.2. Then*:

(i) *For each $d \in F^\times$, we have*

$$c(\phi_d) = \frac{\#\,\mathrm{Sel}(\phi_d)}{\#\,\mathrm{Sel}(\hat{\phi}_d)} \cdot \frac{\#\widehat{B}_d[\hat{\phi}_d](F)}{\#A_d[\phi_d](F)}.$$

(ii) *If T_k is nonempty, then it has positive density and*

$$\operatorname*{avg}_{d \in T_k} \dim_{\mathbb{F}_3} \mathrm{Sel}(\phi_d) \oplus \mathrm{Sel}(\hat{\phi}_d) \leq |k| + 3^{-|k|}.$$

(iii) *For a proportion of at least $1 - \frac{1}{2 \cdot 3^{|k|}}$ of $d \in T_k$, we have* $\dim_{\mathbb{F}_3} \mathrm{Sel}(\phi_d) \oplus \mathrm{Sel}(\hat{\phi}_d) = |k|$.

Proof. Since ϕ and $\hat{\phi}$ determine dual local conditions in their respective Selmer groups [Česnavičius 2017, B.1], the Greenberg–Wiles formula [Neukirch et al. 2000, 8.7.9] applies and gives (i). The argument for (ii) is then exactly the same as [Bhargava et al. 2020, Theorem 50].

For (iii), we may assume $k \geq 0$, by switching ϕ and $\hat{\phi}$ if necessary. By (i), we have

$$\dim \mathrm{Sel}(\phi_d) + \dim \mathrm{Sel}(\hat{\phi}_d) = k \quad \text{if and only if} \quad \dim \mathrm{Sel}(\hat{\phi}_d) = 0,$$

at least for the 100% of d such that $\#\widehat{B}_d[\hat{\phi}_d](F) = \#A_d[\phi_d](F) = 1$. By Theorem 5.2, the average size of $\mathrm{Sel}(\hat{\phi}_d)$ for $d \in T_k$ is $1 + 3^{-k}$. Hence, if s_0 is the lim inf of the natural density of $d \in T_k$ with $\#\,\mathrm{Sel}(\hat{\phi}_d) = 1$, then

$$s_0 + 3(1 - s_0) \leq 1 + 3^{-k},$$

and hence $s_0 \geq 1 - \frac{1}{2 \cdot 3^{|k|}}$. □

Proposition 5.5. *Let $\phi : A \to B$ and T_k be as in Theorem 5.2, with Σ the set of squarefree twists. Let S be the set of places of F dividing $3\mathfrak{f}_A\infty$, where $\mathfrak{f}_A$ is the conductor of A. Then $T_k = \varnothing$ if $|k| > \#S$. As a consequence, we have* $\operatorname{avg}_\Sigma \dim_{\mathbb{F}_3} \operatorname{Sel}(\phi_d) \oplus \operatorname{Sel}(\hat{\phi}_d) \leq \#S + 3^{-\#S}$.

Proof. If $v \nmid 3$ is a prime of good reduction, then by Theorem 4.11 and the assumption that d is squarefree, $c_v(\phi_d) = 1$. On the other hand, directly from the definition, we see that $c_v(\phi_d) \geq \frac{1}{3}$ for any v. Hence,

$$c(\phi_d) = \prod_{v \mid 3\mathfrak{f}_A\infty} c_v(\phi_d) \geq \prod_{v \mid 3 f_A\infty} \frac{1}{3} = 3^{-\#S}.$$

For almost all $d \in \Sigma$, we have $A_d[\phi_d](F) = 0 = \widehat{B}_d[\hat{\phi}_d](F)$, in which case Proposition 5.4(i) gives

$$c(\phi_d)c(\hat{\phi}_d) = \frac{\#\operatorname{Sel}(\hat{\phi}_d)}{\#\operatorname{Sel}(\phi_d)} \cdot \frac{\#\operatorname{Sel}(\phi_d)}{\#\operatorname{Sel}(\hat{\phi}_d)} = 1.$$

By the above, we also have $c(\hat{\phi}_d) \geq 3^{-\#S}$, and hence $c(\phi_d) \leq 3^{\#S}$. Since $3^{-\#S} \leq c(\phi_d) \leq 3^{\#S}$, we have $T_k = \varnothing$ if $|k| > \#S$. The second claim now follows from Proposition 5.4(ii). □

6. The average rank is bounded in cyclotomic twist families

In this section, we apply Theorems 1.8 and 1.9 to prove Theorems 1.1 and 1.3, i.e., to bound the average Mordell–Weil rank of $A_d(F)$.

Proof of Theorem 1.1(ii). Since $A[\pi]$ admits a full flag, by Lemma 2.4, so does $A^{(i)}[\pi]$ for each $i = 1, \ldots, 2 \cdot 3^{m-1} - 1$. Hence, for each i, there is a sequence of ζ-linear 3-isogenies

$$A^{(i)} = B_0^{(i)} \xrightarrow{\phi_1^{(i)}} B_1^{(i)} \to \cdots \to B_{k-1}^{(i)} \xrightarrow{\phi_k^{(i)}} B_k^{(i)} = A^{(i+1)}.$$

By Theorem 1.9, for each i, j, the average rank of $\operatorname{Sel}(\phi_{j,d}^{(i)})$, for squarefree $d \in F^\times/F^{\times 2n}$ is bounded.

Recall [Bhargava et al. 2019, Lemma 9.1], that if $\psi_1 : A_1 \to A_2$ and $\psi_2 : A_2 \to A_3$ are isogenies of abelian varieties, then there is an exact sequence

$$\operatorname{Sel}_{\psi_1}(A_1) \to \operatorname{Sel}_{\psi_2 \circ \psi_1}(A_1) \to \operatorname{Sel}_{\psi_2}(A_2). \tag{6-1}$$

Hence, for each d, we have

$$\begin{aligned} \operatorname{rk} A_d(F) &\leq \dim_{\mathbb{F}_3} \frac{A_d(F)}{3A_d(F)} \leq \dim_{\mathbb{F}_3} \operatorname{Sel}_3(A_d) \\ &\leq \sum_{i=0}^{2\cdot 3^{m-1}-1} \sum_{j=1}^{k} \dim_{\mathbb{F}_3} \operatorname{Sel}(\phi_{j,d}^{(i)}) \\ &\leq \sum_{i=0}^{2\cdot 3^{m-1}-1} \sum_{j=1}^{k} \#\operatorname{Sel}(\phi_{j,d}^{(i)}). \end{aligned}$$

Here, the second inequality follows inductively from (6-1), and the third inequality follows from the trivial inequality $k \leq 3^k$ for all $k \in \mathbb{Z}$. Taking the average over d, the result follows from Theorem 1.9. □

In the remainder of this section, we prove Theorem 1.1(i). The proof just given does not apply: by assumption, $B_0^{(i)}[\phi_1^{(i)}]$ is a direct summand of $A^{(i)}[\pi]$, however, there is no reason that $B_{j-1}^{(i)}[\phi_j^{(i)}]$ should be a direct summand of $B_{j-1}^{(i)}[\pi]$. Hence, we cannot directly apply Theorem 1.8 to get the result. Instead, we exploit the fact that we can take the isogenies $\phi_j^{(i)}$ in any order, together with the following dual version of Theorem 1.8:

Corollary 6.1. *Let B, C be abelian varieties over F with ζ-multiplication and a ζ-linear 3-isogeny $\phi : B \to C$ defined over F. Suppose that $\widehat{C}[\hat{\phi}]$ is almost everywhere locally a direct summand of $\widehat{C}[\pi]$. Then the average size of the Selmer group* $\mathrm{Sel}(\phi_d)$ *is bounded.*

Proof. Let $\hat{\phi} : \widehat{C} \to \widehat{B}$ be the dual isogeny, which is itself a ζ-linear 3-isogeny, with respect to the natural ζ-multiplication structure on $\widehat{C}$ and $\widehat{B}$. Let $\mathbb{B} = \mathbb{B}_{2n}$ as in Section 5. For each $d \in \mathbb{B}(F)$, Proposition 5.4(i) gives

$$\frac{\#\mathrm{Sel}(\phi_d)}{\#\mathrm{Sel}(\hat{\phi}_d)} = \frac{\#B_d[\phi_d](F)}{\#\widehat{C}_d[\hat{\phi}_d](F)} c(\phi_d).$$

The classes d for which $B_d[\phi_d](F) = \mathbb{Z}/3\mathbb{Z}$ or $\widehat{C}_d[\hat{\phi}_d](F) = \mathbb{Z}/3\mathbb{Z}$, form a union of at most two square-classes in $\mathbb{B}(F)$, so we can ignore these values of d when trying to bound the average size of $\mathrm{Sel}(\phi_d)$. In any case, up to a harmless factor of 3, the formula reads

$$\frac{\#\mathrm{Sel}(\phi_d)}{\#\mathrm{Sel}(\hat{\phi}_d)} = c(\phi_d).$$

By Theorem 1.8 applied to $\hat{\phi}$, the average size of $\#\mathrm{Sel}(\hat{\phi}_d)$ is bounded. In fact, if we restrict to the set $T_k = T_k(\hat{\phi}) \subset \mathbb{B}(F)$ where $c(\hat{\phi}_d) = 3^k$, then the average size is equal to $1 + 3^k$ by Theorem 5.2. Since $c(\phi_d) = 1/c(\hat{\phi}_d)$, the average size of $\mathrm{Sel}(\phi_d)$ is equal to $3^{-k} + 1$, and in particular converges, at least on this set T_k.

To see that the average size of $\#\mathrm{Sel}(\phi_d)$ converges (to the expected number) on all of $\mathbb{B}(F)$, we can argue as in [Bhargava et al. 2020, §6.4], using the uniformity estimate [Bhargava et al. 2013, Proposition 29] to give a tail bound for those binary cubic forms with large square part in their discriminant. The uniformity estimate applies to elements of $\#\mathrm{Sel}(\hat{\phi}_d)$ since, under the hypotheses of the corollary, they are represented by integral binary cubic forms. However, we need to bound the average size of $c(\phi_d)\#\mathrm{Sel}(\hat{\phi}_d)$ and not $\#\mathrm{Sel}(\hat{\phi}_d)$, so we also need to control the size of $c(\phi_d)$. By Table 3, we have $c(\phi_d) = O(3^m)$, where m is the number of primes v of F with $v(d)$ even and positive. Moreover, each element of $\mathrm{Sel}(\hat{\phi}_d)$ corresponds to a binary cubic form whose discriminant has norm divisible by q_v^2, for all such v. Since

$$3^m \prod_{v(d)=2} q_v^{-2} \gg 3q_{v_1}^{-2} \gg q_{v_1}^{-2},$$

the same argument as in [loc. cit.] goes through. We refer there for the details. □

Lemma 6.2. *Let F be any field and suppose there is a commutative diagram of isogenies of abelian varieties over F,*

$$\begin{array}{ccc} A & \xrightarrow{\phi_1} & B_1 \\ \downarrow{\scriptstyle \phi_2} & & \downarrow{\scriptstyle \psi_2} \\ B_2 & \xrightarrow{\psi_1} & C \end{array}$$

such that ϕ_2 maps $A[\phi_1]$ isomorphically onto $B_2[\psi_1]$. Then ψ_2 induces an injection

$$\frac{B_1(F)}{\phi_1(A(F))} \hookrightarrow \frac{C(F)}{\psi_1(B_2(F))}.$$

In particular, there is an embedding $\mathrm{Sel}_{\phi_1}(A) \hookrightarrow \mathrm{Sel}_{\psi_1}(B_2)$.

Proof. Taking cohomology, we obtain the following commutative diagram, with exact rows:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A[\phi_1](F) & \longrightarrow & A(F) & \xrightarrow{\phi_1} & B_1(F) & \xrightarrow{\delta_1} & H^1(F, A[\phi_1]) \\ & & \| {\scriptstyle \phi_2} & & \downarrow {\scriptstyle \phi_2} & & \downarrow {\scriptstyle \psi_2} & & \| {\scriptstyle \phi_2} \\ 0 & \longrightarrow & B_2[\psi_1](F) & \longrightarrow & B_2(F) & \xrightarrow{\psi_1} & C(F) & \xrightarrow{\partial_1} & H^1(F, B_2[\psi_1]) \end{array}$$

Consider the composite of maps

$$B_1(F) \xrightarrow{\psi_2} C(F) \to \frac{C(F)}{\psi_1(B_2(F))}.$$

We show that the kernel is exactly $\phi_1(A(F))$ and, therefore, that there is an injection

$$\psi_2 : \frac{B_1(F)}{\phi_1(A(F))} \hookrightarrow \frac{C(F)}{\psi_1(B_2(F))},$$

from which the result follows. If $x \in A(F)$, then $\psi_2(\phi_1(x)) = \psi_1(\phi_2(x)) = 0$. Conversely, if $y \in B_1(F)$ and $\psi_2(y) = \psi_1(z)$ for some $z \in B_2(F)$, then

$$\phi_2(\delta_1(y)) = \partial_1(\psi_2(y)) = \partial_1(\psi_1(z)) = 0.$$

Since ϕ_2 is an isomorphism on cohomology, it follows that $\delta_1(y) = 0$ and, hence, that y is in the image of ϕ_1. □

Recall that $A[\pi]$ decomposes as a direct sum of characters, so that $A[\pi](\bar{F}) = \langle P_1, \ldots, P_k \rangle$, and $\mathrm{Gal}(\bar{F}/F)$ stabilizes each of the subgroups $\langle P_i \rangle$. Thus, π factors as a product of ζ-linear 3-isogenies:

$$\begin{array}{ccccccc} A = B_0 & & \xrightarrow{\qquad\qquad\pi\qquad\qquad} & & & & A^{(1)} = B_k \\ & \searrow^{\phi_1} & & & & \nearrow_{\phi_k} & \\ & & B_1 = A/\langle P_1 \rangle & \xrightarrow{\phi_2} & \cdots & \xrightarrow{\phi_{k-1}} & B_{k-1} = A/\langle P_1, \ldots, P_{k-1} \rangle \end{array}$$

Moreover, each of the maps $\phi_i : B_{i-1} \to B_i$ can be twisted to a map $\phi_{i,d} : B_{i-1,d} \to B_{i,d}$.

Corollary 6.3. *For each $i = 1, \dots, k$, the average size of $\mathrm{Sel}_{\phi_{i,d}}(B_{i-1,d})$, for $d \in \mathbb{B}(F)$, is bounded.*

Proof. The assumption that $A[\pi]$ decomposes as a sum of characters means that we can take the P_i's in any order. Hence, for each i, d, there is a commutative diagram

$$\begin{array}{ccc} B_{i-1,d} & \xrightarrow{\phi_{i,d}} & B_{i,d} \\ \downarrow & & \downarrow \\ B'_{k-1,d} & \xrightarrow{\psi_{i,d}} & A_d^{(1)} \end{array}$$

where

$$B'_{k-1,d} := \frac{A_d}{\langle P_1, \dots, P_{i-1}, P_{i+1} \dots, P_k \rangle}$$

and $\ker(\psi_{i,d}) = \ker(\phi_{i,d}) = \langle P_i \rangle$.

By Lemma 2.4, since $A[\pi]$ is completely reducible, so is $A^{(1)}[\pi]$ and hence so is $\hat{A}^{(1)}[\pi]$. Thus $\hat{A}^{(1)}[\hat{\psi}_{i,d}]$ is a direct summand of $\hat{A}^{(1)}[\pi]$. It follows from Corollary 6.1 that the average size of $\mathrm{Sel}_{\psi_{i,d}}(B'_{k-1,d})$ is bounded. By Lemma 6.2, we have embeddings $\mathrm{Sel}_{\phi_{i,d}}(B_{i-1,d}) \hookrightarrow \mathrm{Sel}_{\psi_{i,d}}(B'_{k-1,d})$ for each $d \in \mathbb{B}(F)$, so the average size of $\mathrm{Sel}_{\phi_{i,d}}(B_{i-1,d})$ is bounded as well. □

Proof of Theorem 1.1(i). By Lemma 2.4, since $A[\pi]$ is completely reducible, so is $A^{(i)}[\pi]$ for all i. Hence, we can factor $A^{(i)} \to A^{(i+1)}$ as

$$A^{(i)} = B_0^{(i)} \xrightarrow{\phi_1^{(i)}} B_1^{(i)} \to \cdots \to B_{k-1}^{(i)} \xrightarrow{\phi_k^{(i)}} B_k^{(i)} = A^{(i+1)}.$$

By Corollary 6.3, for each i, j, the average rank of $\mathrm{Sel}(\phi_{j,d}^{(i)})$ is bounded. The result now follows exactly as in the proof of Theorem 1.1(ii). □

Proof of Theorem 1.3. The hypotheses imply that the Rosati involution associated to the polarization restricts to complex conjugation on the subring $\mathbb{Z}[\zeta]$. Thus, after identifying $A \simeq \hat{A}$ and $A_{-3^n} \simeq \hat{A}_{-3^n}$, we can factor multiplication by -3 on A_d as the composition

$$A_d \xrightarrow{\pi_d^{3^{m-1}}} A_d^{(3^{m-1})} = A_{-3^n d} \xrightarrow{\hat{\pi}_d^{3^{m-1}}} \hat{A}_d,$$

where the middle equality is Remark 2.8. As in the proof of Theorem 1.1(ii), we can factor $\pi_d^{3^{m-1}}$ as a product of $\dim A$ 3-isogenies $\phi_{j,d}$. By duality, $\hat{\pi}_d^{3^{m-1}}$ factors as the product of the dual isogenies $\hat{\phi}_{j,d}$. Thus, for each d, we have

$$\begin{aligned} \mathrm{rk}\, A_d(F) &\le \dim_{\mathbb{F}_3} \frac{A_d(F)}{3A_d(F)} \le \dim_{\mathbb{F}_3} \mathrm{Sel}_3(A_d) \\ &\le \sum_{j=1}^{\dim A} \dim_{\mathbb{F}_3} \mathrm{Sel}(\phi_{j,d}) \oplus \mathrm{Sel}(\hat{\phi}_{j,d}) \end{aligned}$$

and hence, by Proposition 5.5,

$$\mathrm{avg}_d\, \mathrm{rk}\, A_d(F) \le \dim A \cdot (\#S + 3^{-\#S}).$$

□

7. The average rank in twist families of trigonal Jacobians

Next we use Theorem 1.1 to prove Corollary 1.2. In this section, F is a number field and $\zeta \in \bar{F}$ is a primitive 3^m-th of unity, for some $m \geq 1$. Let $n = 3^m$, as always.

7A. ***Trigonal Jacobians.*** Let $f(x) \in F[x]$ be a monic separable polynomial such that $f(0) \neq 0$, and let C be the smooth projective curve with affine model

$$C : y^3 = xf(x^{3^{m-1}}).$$

If $m > 1$, then C has a unique rational point ∞ at infinity, and has genus $g = 3^{m-1}\deg(f)$. In fact, we will assume $m > 1$, since the case $m = 1$ will be subsumed by the results of Section 7B.

Let $J = \mathrm{Jac}(C)$ be the Jacobian of C, a g-dimensional principally polarized abelian variety over F. The automorphism $(x, y) \mapsto (\zeta^3 x, \zeta y)$ induces an automorphism of $J_{\bar{F}}$ of order 3^m, which we again call ζ, and which endows J with ζ_n-multiplication. As in Section 2, the endomorphism $1 - \zeta \in \mathrm{End}_{\bar{F}}(J)$ descends to an isogeny $\pi : J \to J^{(1)}$ over F.

Lemma 7.1. *Let G be an extension of $(\mathbb{Z}/2\mathbb{Z})^k$ by a 3-group H. Then every irreducible representation $\rho : G \to \mathrm{GL}_N(\mathbb{F}_3)$ is one-dimensional. Consequently, any G-representation V over $\mathbb{F}_3$ admits a full flag.*

Proof. Since $H \lhd G$ is normal and ρ is semisimple, $\rho|_H$ is also semisimple. Now, any nontrivial representation V over $\mathbb{F}_3$ of a 3-group contains a nonzero fixed vector [Serre 1977, Proposition 26]. Thus, by semisimplicity and induction on $\dim V$, we see that $\rho|_H$ is trivial. Thus, ρ factors through a representation of $(\mathbb{Z}/2\mathbb{Z})^k$. For any $g \in (\mathbb{Z}/2\mathbb{Z})^k$, $\rho(g) \in \mathrm{GL}_N(\mathbb{F}_3)$ has order at most 2, so its minimal polynomial, either $X \pm 1$ or $X^2 - 1$, has distinct $\mathbb{F}_3$-rational roots. Hence, $\rho(g)$ is diagonalizable, and since $(\mathbb{Z}/2\mathbb{Z})^k$ is abelian, the operators $\rho(g)$ are simultaneously diagonalizable. In other words, ρ is a direct sum of characters. Since ρ is irreducible, it follows that ρ is one-dimensional. □

We first prove Corollary 1.2 for Jacobians of the curves $C : y^3 = xf(x^{3^{m-1}})$. In Theorem 7.4, we will address the curves $C : y^{3^m} = f(x)$.

Proof of Corollary 1.2. By assumption $\mathrm{Gal}(f)$ is an extension of $(\mathbb{Z}/2\mathbb{Z})^k$ by a 3-group H. By Theorem 1.1, it is enough to show that the Galois representation $J[\pi]$ has a full flag, and splits as a direct sum of characters if $H = 1$.

Lemma 7.2. $J[\pi^{3^{m-1}}] = J[1 - \zeta_3]$ *is a maximal isotropic $\mathbb{F}_3$-subspace of $J[3]$ of dimension g.*

Proof. Degree considerations show that $\dim J[1 - \zeta_3] = g$, so we need only show that $J[1 - \zeta_3]$ is isotropic with respect to the Weil pairing $J[3] \times J[3] \to \mu_3$. Now, the Rosati involution $\dagger$ sends the ideal $(1 - \zeta_3) = (\sqrt{-3})$ to itself (by Lemma 10.4 below), and $\langle \alpha P, Q\rangle = \langle P, \alpha^\dagger Q\rangle$ for all $\alpha \in \mathrm{End}(J_{\bar{F}})$ and $P, Q \in J[3]$. If $P, Q \in J[(\sqrt{-3})]$, we may write $Q = \sqrt{-3}(R)$, for some $R \in J[3]$. We then compute

$$\langle P, Q\rangle = \langle P, \sqrt{-3}(R)\rangle = \langle -\sqrt{-3}(P), R\rangle = 1,$$

showing that $J[1 - \zeta_3] = J[\sqrt{-3}]$ is isotropic. □

It follows that $\dim_{\mathbb{F}_3} J[\pi] = \deg(f)$. Explicitly, if α is a root of f and $\beta^{3^{m-1}} = \alpha$, then the divisor

$$(\beta, 0) + (\zeta^3\beta, 0) + (\zeta^{3\cdot 2}\beta, 0) + \cdots + (\zeta^{3^m-3}\beta, 0) - 3^{m-1}\infty$$

is fixed by ζ, and $J[\pi]$ is generated by the above divisor classes. Moreover, if K is the splitting field of f over F, then each of these divisors defines an element of $J(K)$.

The action of G_F on $J[\pi]$ induces a homomorphism $\rho : G_F \to \mathrm{GL}_N(\mathbb{F}_3)$, where $N = \deg(f)$, whose kernel is exactly G_K. The image is therefore isomorphic to $\mathrm{Gal}(K/F)$, which is an extension of $(\mathbb{Z}/2\mathbb{Z})^k$ by a 3-group H. Hence, by Lemma 7.1, $J[\pi]$ admits a full flag. If, moreover, $H = 1$, then $J[\pi]$ is completely reducible, as explained in the proof of Lemma 7.1. □

7B. ***Iterated triple covers and Pryms.*** For our second class of Jacobians, let $f(x) \in F[x]$ be a monic separable polynomial of degree $N > 1$. Let C be the smooth projective curve with affine model $y^{3^m} = f(x)$. The degree $n = 3^m$ map $C \to \mathbb{P}^1$, sending $(x, y) \mapsto x$, has Galois group μ_n, at least over $\bar{F}$. If $3 \nmid N$, then the fiber above infinity is a unique F-rational point and C has genus $g = \frac{1}{2}(N-1)(3^m-1)$. If $3 \mid N$, then the fiber above infinity may have more than one point and they may not be F-rational.

Let $J = \mathrm{Jac}(C)$ be the Jacobian. The order 3^m automorphism $(x, y) \mapsto (x, \zeta y)$ of C induces an automorphism ζ on J. When $m = 1$, this endows J with ζ_3-multiplication, and we are in a trigonal situation similar to Section 7A. However, if $m \geq 2$, the automorphism $\zeta \in \mathrm{Aut}_{\bar{F}}(J)$ does not give rise to a ζ_n-multiplication on J, as we have defined it in this paper.

Example 7.3. Consider the curve $C : y^9 = x^2 - 1$ of genus 4. This admits a degree three map to the elliptic curve $E : y^3 = x^2 - 1$, and the Jacobian $J = \mathrm{Jac}(C)$ is isogenous to $A \times E$, for some abelian 3-fold $A \subset J$. The order 9 automorphism ζ induces an order 9 automorphism on A and an order 3 automorphism on E. It thereby endows A with ζ_9-multiplication and E with ζ_3-multiplication, but it does *not* give a ζ_9-multiplication on J. Indeed, the minimal polynomial for $\zeta \in \mathrm{End}_{\bar{F}} J$ is $\Phi_9(x)\Phi_3(x)$ and not $\Phi_9(x) = x^6 + x^3 + 1$.

While J does not admit ζ_n-multiplication, it is nonetheless isogenous to a product of abelian varieties $P_1 \times P_2 \times \cdots \times P_m$ where each P_i has ζ_{3^i}-multiplication (see Lemma 7.5). In any case, we have $\mu_{2n} \subset \mathrm{Aut}_{\bar{F}} J$, and we may speak of the twists J_d, for each $d \in F^\times/F^{\times 2n}$.

Theorem 7.4. *Assume that* $\mathrm{Gal}(f)$ *is an extension of* $(\mathbb{Z}/2\mathbb{Z})^k$ *by a* 3*-group* H*. Then the average rank of the twists* J_d *for squarefree* $d \in F^\times/F^{\times 2n}$ *is bounded. If* $H = 1$*, then the average rank of the twists* J_d*, for all* $d \in F^\times/F^{\times 2n}$ *is bounded.*

Proof. Let $C' : y^{3^{m-1}} = f(x)$, and let J' be its Jacobian. Note that when $m = 1$, we have $C' \simeq \mathbb{P}^1$ and $J' = 0$. The map $q : C \to C'$ sending $(x, y) \mapsto (x, y^3)$ induces a surjection $q_* : J \to J'$, and we let P be the identity component of the kernel, i.e., P is the (generalized) Prym variety for the cover q. Since q is a ramified triple cover, the map $q^* : J' \to J$ is injective. We may identify $q^* = \hat{q}_*$, and it follows that the kernel of q_* is already connected, and hence equal to P.

Lemma 7.5. *P admits ζ_{3^m}-multiplication.*

Proof. It is enough to show that the minimal polynomial for ζ_{3^m}, as an endomorphism of P, is $\Phi_{3^m}(x) = x^{2\cdot 3^{m-1}} + x^{3^{m-1}} + 1$. For this, it is enough to show that the characteristic polynomial of ζ_{3^m}, acting on the homology lattice $H_1(P_{\mathbb{C}}, \mathbb{Z})$ is $\Phi_{3^m}(x)^{N-1}$. By [Arul 2021, Lemma 3.16], the characteristic polynomial of ζ_{3^m} acting on $H_1(C_{\mathbb{C}}, \mathbb{Z}) \simeq H_1(J_{\mathbb{C}}, \mathbb{Z})$ is $(1 + x + x^2 + \cdots + x^{3^m-1})^{N-1}$. The claim now follows, by induction on m. □

Let $\pi : P \to P^{(1)}$ be the isogeny over F descending $1-\zeta$ over $F(\zeta)$, as usual. Note that $P[\pi] \subset P[3]$ since P has ζ_{3^m}-multiplication, whereas $J[1-\zeta]$ is not in general contained in $J[3]$.

Lemma 7.6. *The representation $G_F \to \mathrm{Aut}_{\mathbb{F}_3} P[\pi]$ factors through* $\mathrm{Gal}(f)$.

Proof. If $\alpha_1, \ldots, \alpha_N$ are the roots of $f(x)$, then the divisor classes $(\alpha_i, 0) - (\alpha_j, 0)$ generate the group $J[1-\zeta]$ [Schaefer 2018, §3] and so the G_F-action on $P[\pi] \subset J[1-\zeta]$ factors through $\mathrm{Gal}(f)$. □

Now we finish the proof of Theorem 7.4. By Lemmas 7.1 and 7.6, the Galois module $P[\pi]$ has a full flag, and is completely reducible if $H = 1$. Thus Theorem 1.1 says that the average rank of P_d, for squarefree $d \in F^\times/F^{\times 2\cdot 3^m}$, is bounded (and without the squarefree condition if $H = 1$). Up to isogeny, we have $J_d \simeq P_d \times J'_d$, where J'_d is the twist of J' (which has $\mu_{2\cdot 3^{m-1}}$ twists) by the image of d under $F^\times/F^{\times 2\cdot 3^m} \to F^\times/F^{\times 2\cdot 3^{m-1}}$. By induction, the average rank of J'_d, for $d \in F^\times/F^{\times 2\cdot 3^m}$ is bounded. (We view the family J'_d over $\mathbb{B}_{2\cdot 3^m}$, instead of the more natural $\mathbb{B}_{2\cdot 3^{m-1}}$, but the same proof works for this slightly "redundant" family as well.) Thus the average rank of J_d is bounded. □

Remark 7.7. Combining the two families considered in this section, we obtain similar results for the curves $C_{k,j} : y^{3^k} = xf(x^{3^j})$. The Jacobian $\mathrm{Jac}(C_{k,j})$ is isogenous to $\prod_{r=1}^{k} P_{r,j}$, where each $P_{r,j} = \mathrm{Prym}(C_{r,j} \to C_{r-1,j})$ is a generalized Prym variety with $\zeta_{3^{r+j}}$-multiplication.

8. CM abelian varieties

Next, we prove Theorem 1.4. Let $\zeta = \zeta_{3^m}$ be a primitive 3^m-th root of unity. We recall our definition of complex multiplication:

Definition 8.1. An abelian variety A over a number field F *has complex multiplication by* $\mathbb{Z}[\zeta]$ if A has dimension 3^{m-1} and a G_F-equivariant ring embedding $\mathbb{Z}[\zeta] \hookrightarrow \mathrm{End}_{\overline{F}} A$.

Proof of Theorem 1.4. Set $g = 3^{m-1} = \dim A$. The assumption $\mathbb{Q}(\zeta) \subset F$ ensures that $A \simeq A^{(1)}$ by Lemma 2.3. Hence, we can view the 3-isogeny $\pi : A \to A^{(1)}$ as an endomorphism of A, and we have $\pi^{2g} = 3u$ for some automorphism u of A. By the multiplicativity of the global Selmer ratio [Shnidman 2021, Corollary 3.5], we have

$$c(\pi_d)^{2g} = c(\pi_d^{2g}) = c([3]u) = c([3]). \tag{8-1}$$

We claim that $c([3]) = 1$. If v is an infinite prime, then since $F \supset \mathbb{Q}(\zeta)$, we have $F_v \simeq \mathbb{C}$ and $c_v([3]) = \#A[3](\mathbb{C})^{-1} = 3^{-2g}$. If $v \nmid 3$ is a finite prime then $c_v([3]) = c_v(A_d)/c_v(A_d) = 1$. Finally, $\prod_{v|3} c_v([3]) = 3^{g[F:\mathbb{Q}]}$, by [Shnidman 2021, Proposition 3.1]. Let $[F : \mathbb{Q}] = 2N$. Then F has N complex places, so $c([3]) = 3^{-2gN} \cdot 3^{g\cdot 2N} = 1$, as claimed. By (8-1), we also have $c(\pi_d) = 1$ for all $d \in F^\times/F^{\times 6g}$.

It follows from Theorem 5.2 that the average size of $\mathrm{Sel}_{\pi_d}(A_d)$ is $1+1=2$. Since $2r \le 3^r - 1$ for all integers r, we have for all d:

$$\mathrm{rk}_{\mathbb{Z}[\zeta]}\, A_d(F) \le \dim_{\mathbb{F}_3} \mathrm{Sel}_{\pi_d}(A_d) \le \frac{1}{2}(3^{\mathrm{rk}\,\mathrm{Sel}_{\pi_d}(A_d)} - 1) = \frac{1}{2}(\#\,\mathrm{Sel}_{\pi_d}(A_d) - 1).$$

Thus, the average $\mathbb{Z}[\zeta]$-rank of $A_d(F)$ is at most $\frac{1}{2}$. The second part of the theorem follows from Proposition 5.4(ii). □

Over general number fields F, it is no longer true that $c(\pi_d) = 1$ for all d, even for abelian varieties with complex multiplication. Moreover, one must consider more than one 3-isogeny to bound the average rank of $A_d(F)$, in general. However, one can still prove upper bounds which are significantly stronger than Theorem 1.3. For example, for CM abelian varieties A of dimension $g = 3^{m-1}$ over $\mathbb{Q}$ with ζ_{3^m}-multiplication, we can show that the average rank of $A_d(F)$ is at most $\frac{13}{9}g$. Over the totally real field $\mathbb{Q}(\zeta_{3^m} + \bar{\zeta}_{3^m})$, we can also prove that an explicit positive proportion of twists have $\mathrm{rk}\, A_d(F) = 0$. We omit these proofs, since they are somewhat technical and could probably be optimized further. Finally, we remark that the only other result in the literature in this direction that we are aware of is [Diaconu and Tian 2005], which proves that an infinite but density zero set of twists of the degree p Fermat Jacobian over $\mathbb{Q}(\zeta_p + \bar{\zeta}_p)$ have rank 0.

9. Rational points on hyperbolic varieties

Proof of Theorem 1.5. It is not enough to simply invoke Theorem 1.1 and [Kühne 2021, Theorem 4], since $\mathrm{Jac}(C_d)$, is not the d-th sextic twist of $J = \mathrm{Jac}(C)$ in our sense. Indeed, the twists C_d come from the involution $\tau(x, y) = (-x, y)$, which does not induce -1 on $\mathrm{Jac}(C)$. Instead, we consider the Prym variety $P = \ker(\mathrm{Jac}(C) \to \mathrm{Jac}(C/\tau))$, and its dual $\widehat{P}$. Note that ζ preserves P, and τ restricts to -1 on P. Thus, Theorem 1.1 applies to the twist family P_d, for $d \in F^\times/F^{\times 6}$, and it follows that $\mathrm{avg}_d\, \mathrm{rk}\, \widehat{P}_d(F) = \mathrm{avg}_d\, \mathrm{rk}\, P_d(F)$ is bounded.

The inclusion $P_d \hookrightarrow \mathrm{Jac}(C_d)$ induces a surjection $q : \mathrm{Jac}(C_d) \to \widehat{P}_d$. Suppose that $C_d(F)$ is nonempty and that $z_0 \in C_d(F)$. Then composing with the Abel–Jacobi map $C_d \hookrightarrow \mathrm{Jac}(C_d)$, using z_0 as base point, we obtain a map $j : C_d \to \widehat{P}_d$. We prove that j is injective on points, except possibly at the fixed points of τ (the points where $x = 0$ and the point(s) at infinity). If $j(w) = j(z)$, then $w - z$ is the pullback of a divisor on C_d/τ. Thus $\tau(w - z) \equiv w - z$, and so $\tau(w) + z \equiv w + \tau(z)$. But if D is a divisor of degree 2 on a nonhyperelliptic curve C, then by Riemann–Roch, we have $h^0(D) = 2 + 1 - g + h^0(K - D) = 3 - g + g - 2 = 1$. Thus, we must have $\tau(w) + z = w + \tau(z)$. If $\tau(w) = \tau(z)$, then $w = z$, as desired. The only other possibility is that $\tau(w) = w$ and $\tau(z) = z$. This proves the claimed injectivity.

Thus, to prove Theorem 1.5, we may replace C_d with the image of j, which is a closed irreducible curve in $\widehat{P}_d$. By [Gao et al. 2021, Theorem 1.1], there is a constant c such that $\#C_d(F) \le c^{1+\mathrm{rk}\,\widehat{P}_d(F)}$, for all d. But since $\mathrm{avg}_d\, \mathrm{rk}\, \widehat{P}_d(F)$ is bounded, this implies that for all $\varepsilon > 0$, there exists N_ε such that $C_d(F) \le N_\varepsilon$ for at least $1 - \varepsilon$ of twists d. □

In order to set up the proof (and statement) of Theorem 1.6, we need to give a precise description of theta divisors for the curves C with affine model $y^3 = f(x)$. Recall that for any smooth projective curve C/F of genus $g \geq 2$, the symmetric power $C^{(g-1)}$ parametrizes effective divisors D on C of degree $g-1$. Given $\kappa \in \mathrm{Div}(C)$ of degree $g-1$, there is a morphism $C^{(g-1)} \to \mathrm{Jac}(C)$ sending $D \mapsto D - \kappa$. Its image is the theta divisor, denoted $\Theta = \Theta_\kappa$. The divisor itself depends on κ, though its class in the Néron–Severi group of $\mathrm{Jac}(C)$ does not. If 2κ is canonical, then Θ is preserved by the involution -1 on $\mathrm{Jac}(C)$, by Riemann–Roch. Such a κ exists over $\bar{F}$, but need not exist over F, in general. If in addition there exists $\mu_n \subset \mathrm{Aut}_{\bar{F}}(C)$ which fixes κ, then for each $d \in F^\times/F^{\times 2n}$, we may consider the twist Θ_d of Θ, which is a divisor in $\mathrm{Jac}(C)_d$.

In our case, let $f : C \to \mathbb{P}^1$ be the degree three map $(x, y) \mapsto x$. The ramification divisor has the form $2D$, and satisfies $K_C = f^*K_{\mathbb{P}^1} + 2D$. Hence $\kappa = D - f^*0$ is a half-canonical divisor (over F) which is fixed by the μ_3-action. We therefore obtain sextic twists $\Theta_d \subset \mathrm{Jac}(C)_d$, for each $d \in F^\times/F^{\times 6}$, as in the statement of Theorem 1.6.

Proof of Theorem 1.6. This now follows from Theorem 1.1 and [Gao et al. 2021, Theorem 1.1]. □

10. Abelian surfaces with ζ_3-multiplication

For our final application, we study a family of abelian surfaces with ζ_3-multiplication, arising as Prym varieties. We prove results on the Mordell–Weil groups in sextic twist families of such surfaces, and give applications to explicit uniform bounds on rational points in sextic twist families of bielliptic trigonal curves of genus 3 (Theorem 1.7). For some recent results on rank statistics in larger families of Prym surfaces, see [Laga 2023].

Let F be a number field and let $f(x) = x^2 + ax + b \in F[x]$ be a quadratic polynomial with nonzero discriminant and $b \neq 0$. Then $y^3 = f(x^2) = x^4 + ax^2 + b$ is an affine model of a smooth projective plane quartic curve C. Note the double cover $\pi : (x, y) \mapsto (x^2, y)$ to the elliptic curve $E : y^3 = f(x)$. We refer to these genus-3 curves as *bielliptic Picard curves*; see [Laga and Shnidman 2023]. As in Section 9, we consider the Prym variety $P = \mathrm{Prym}_{C/E}$, i.e., the kernel of the map $J = \mathrm{Jac}(C) \to E$ induced by Albanese functoriality. The Prym P need not be principally polarized over $\mathbb{Q}$, but it admits a polarization $\lambda : P \to \widehat{P}$ whose kernel is order 4 [Mumford 1974]. The ζ_3-multiplication on J induces ζ_3-multiplication on P, and hence we may speak of the sextic twists P_d. In fact, P_d is itself the Prym variety of $C_d : y^3 = x^4 + adx^2 + bd^2$, which covers the elliptic curve $E_d : y^3 = x^2 + adx + bd^2$.

Lemma 10.1. *Let $\pi : P \to P_{-27}$ denote the descent of $1 - \zeta$ to F. Then $P[\pi](\bar{F}) \simeq (\mathbb{Z}/3\mathbb{Z})^2$ is spanned by $(s, 0) - (-s, 0)$ and $(t, 0) - (-t, 0)$, where $\pm s, \pm t$ are the four roots of $f(x^2)$.*

In order to apply our result to P, we assume that $f(x)$ has linear factors over F, so that $P[\pi]$ decomposes as a direct sum of two 1-dimensional Galois modules, corresponding to the quadratic characters $G_F \to \mathbb{F}_3^\times$ cut out by the fields $F(s)$ and $F(t)$. Then Theorem 1.1 says that the average rank of $P_d(F)$ is bounded, and it is interesting to ask whether there is some positive proportion of d with $\mathrm{rk}\, P_d(\mathbb{Q}) \leq 1$, so that we may apply the Chabauty method. We do not quite prove that such a positive

proportion of d exists for general Pryms of this type, but we can prove it in seemingly any given example with the help of explicit computations. We demonstrate the idea by proving the following result stated in the introduction:

Theorem 1.7. *Consider the sextic twist family $C_d : y^3 = (x^2 - d)(x^2 - 4d)$ of genus-3 curves. For at least $\frac{1}{3}$ of squarefree $d \in \mathbb{Z}$ such that $d \equiv 2$ or $11 \pmod{36}$, we have $\#C_d(\mathbb{Q}) \leq 5$.*

We will use the following variant of Chabauty's method:

Theorem 10.2 (Stoll). *Let C be a smooth projective curve of genus $g \geq 2$ over a number field F, and let H be a G_F-stable subgroup of $\mathrm{Aut}_{\bar{F}} C$. Embed $C \hookrightarrow \mathrm{Jac}(C)$ using any positive degree H-invariant divisor as basepoint. Suppose there is a quotient B of $\mathrm{Jac}(C)$ such that the composition $\iota : C \hookrightarrow \mathrm{Jac}(C) \to B$ is an embedding, and suppose there exists $H \hookrightarrow \mathrm{Aut}_{\bar{F}} B$, compatible with the H-action on C, via ι. Then for all but finitely many H-twists C_ξ with $\mathrm{rk}\, B_\xi(F) < \dim B$, we have*

$$\#C_\xi(F) \leq f_C(\mathrm{rk}\, B_\xi(F) + g - \dim B) + \#C_\xi^{\mathrm{triv}}(F) + \#C_\xi^{\mathrm{triv,non\text{-}tors}}(\bar{F}) \setminus C_\xi^{\mathrm{triv}}(F).$$

Here, f_C is the explicit function defined in [Stoll 2006, §3] *and C_ξ^{triv} is the subscheme of points fixed by some nontrivial automorphism in H, and $C_\xi^{\mathrm{triv,non\text{-}tors}}$ is the subscheme of trivial points which map to nontorsion points of B.*

Proof. This is a straightforward generalization of [Stoll 2006, Theorem 5.1], which is the special case where $B = \mathrm{Jac}(C)$. (We have stated an ineffective version of the result, which is sufficient for our purposes. This is what allows us to use the function f_C as opposed to Stoll's $\tilde{f}_C$.) The proof is the same, except that instead of the nondegenerate pairing

$$\Omega(C/F) \times \mathrm{Jac}(C)(F) \otimes \mathbb{Q} \to F$$

used in [Stoll 2006, §6], we use the nondegenerate pairing $\Omega(C/F)^B \times B(F) \otimes \mathbb{Q} \to F$, where $\Omega(C/F)^B$ is the image of $\iota^* : \Omega(B/F) \to \Omega(C/F)$. □

We deduce Theorem 1.7 from Theorem 10.2 and the following theorem, whose proof will occupy the remainder of this section:

Theorem 10.3. *Let $C : y^3 = (x^2 - 1)(x^2 - 4)$, and let P be the corresponding Prym variety. Let Σ be the set of squarefree $d \in \mathbb{Z}$ such that $d \equiv 2$ or $11 \pmod{36}$.*[2] *Then the average rank of P_d, for $d \in \Sigma$, is at most $\frac{7}{3} \approx 2.33$. Moreover, for at least $\frac{1}{3}$ of $d \in \Sigma$, we have $\mathrm{rk}\ P_d \leq 1$.*

Proof of Theorem 1.7. The curve C embeds in $B = \widehat{P} = J/\pi^* E$, the dual of the Prym P; see [Barth 1987, 1.12]. We apply Theorem 10.2 to the cyclic group H of order six generated by $(x, y) \mapsto (-x, \zeta_3 y)$. We embed C in its Jacobian using the point $\infty = [0:1:0]$. Consulting [Stoll 2006, Lemma 3.1], we have $f_C(r) \leq 4$ when C is a plane quartic and $r \leq 2$. We have $C_d^{\mathrm{triv}}(\mathbb{Q}) = \{\infty\}$ for all $d \in \Sigma$, so the second term in Theorem 10.2 is 1. The third term is 0 since all eight of the trivial points on C map to torsion points of B. Indeed, the points with $y = 0$ map to 3-torsion points on $\mathrm{Jac}(C_d)$, and if $P = (0, y_0) \in C_d$,

[2] Σ is the set of d such that $d, -3d \notin \mathbb{Q}_2^{\times 2} \cup \mathbb{Q}_3^{\times 2}$.

then $2P - 2\infty \in \pi^* E_d$; hence P is sent to a 2-torsion point on $B = J_d/\pi^* E_d$. Altogether we get a bound of $C_\xi(F) \leq 5$ in Stoll's theorem, which combined with Theorem 10.3 proves Theorem 1.7. □

10A. ***Bielliptic Picard curves.*** Before proving Theorem 10.3, we prove some preliminary lemmas.

Lemma 10.4. *Let (J, λ) be the Jacobian of a curve C with a nontrivial automorphism ζ inducing ζ-multiplication on J. Then the Rosati involution $\alpha \mapsto \lambda^{-1}\hat{\alpha}\lambda$ on* $\mathrm{End}(J)$ *restricts to complex conjugation on $\mathbb{Z}[\zeta] \subset \mathrm{End}(J)$.*

Proof. Let D_0 be a degree $g-1$ divisor fixed by ζ. Consider the theta divisor

$$\Theta = \{D - D_0 : \deg(D) = g - 1,\ D \text{ effective}\} \subset J$$

and set $\mathcal{L} = \mathcal{O}_J(\Theta)$. We have $\lambda = \varphi_{\mathcal{L}} : J \to \hat{J}$. Since Θ is fixed by ζ, we have $\zeta^*\mathcal{L} \simeq \mathcal{L}$ and hence $\varphi_{\mathcal{L}} = \varphi_{\zeta^*\mathcal{L}} = \hat{\zeta}\varphi_{\mathcal{L}}\zeta$. Rearranging, we see that the Rosati involution sends ζ to $\zeta^{-1} = \bar{\zeta}$. □

Remark 10.5. The proof shows, more generally, that if α is an automorphism of a curve C, and α^* is the induced automorphism of $\mathrm{Jac}(C)$, then the Rosati involution sends α^* to its inverse.

Now let $C : y^3 = x^4 + ax^2 + b$ be a bielliptic Picard curve defined over $\mathbb{Q}$. Let P be the Prym surface for the covering $C \to E$ where $E : y^3 = x^2 + ax + b$. Since P has ζ_3-multiplication, the endomorphism $[-3] : P \to P$ factors as $[-3] = \pi_{-27} \circ \pi$, where $\pi : P \to P_{-27}$ is the canonical $(3,3)$-isogeny coming from Lemma 2.2 (see also Remark 2.8). Let $\pi_d : P_d \to P_{-27d}$ be the sextic twist family of $(3,3)$-isogenies, and let $\hat{\pi}_d : \widehat{P}_{-27d} \to \widehat{P}_d$ denote the dual isogeny.

Lemma 10.6. $\mathrm{Sel}(\pi_{-27d}) \simeq \mathrm{Sel}(\hat{\pi}_d)$.

Proof. Let $C_d : y^3 = x^4 + adx^2 + bd^2$, let $E_d : y^3 = x^2 + adx + bd^2$, and let $J_d = \mathrm{Jac}(C_d)$.[3] The abelian variety P_d is, by definition, $\ker(J_d \to E_d)$, where the map $J_d \to E_d$ is induced by the double cover $C_d \to E_d$. Let λ_J denote the principal polarization of J_d, and let ζ_J be the automorphism of J_d induced by the map $(x, y) \mapsto (x, \zeta_3 y)$ on C_d. By Lemma 10.4 we have a commutative diagram

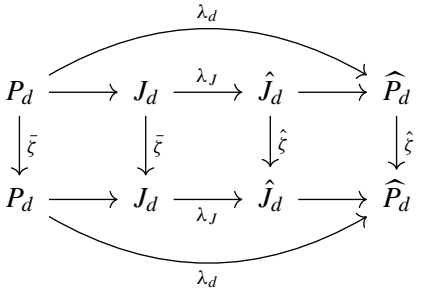

[3]Note that this is not the same as the d-th sextic twist coming from the ζ-multiplication on J. The latter is isomorphic to the d-th quadratic twist of the Jacobian of $dy^3 = f(x^2)$, and is in general not a Jacobian.

It follows that $\zeta^{-1} = \lambda_d^{-1}\hat{\zeta}\lambda_d$ in $\mathrm{End}(P_d)$, and hence

$$(1-\hat{\zeta}^{-1})\circ\lambda_d = \lambda_d \circ (1-\zeta)$$

over $\bar{F}$. Over F we must therefore have $\hat{\pi}_{-27d}\circ\lambda_d = \lambda_{-27d}\circ\pi_d$, and since λ_d is prime-to-3, we deduce $\mathrm{Sel}(\pi_d) \simeq \mathrm{Sel}(\hat{\pi}_{-27d})$. □

Since $[-3] = \pi_{-27}\circ\pi$, it follows that

$$\mathrm{rk}(P_d) \le \dim_{\mathbb{F}_3}\mathrm{Sel}_3(P_d) \le \dim_{\mathbb{F}_3}(\mathrm{Sel}(\pi_d)\oplus\mathrm{Sel}(\pi_{-27d})) = \dim_{\mathbb{F}_3}\mathrm{Sel}(\pi_d) + \dim_{\mathbb{F}_3}\mathrm{Sel}(\hat{\pi}_d).$$

The following result relates the parity of $\dim_{\mathbb{F}_3}\mathrm{Sel}_3(P_d)$ to the global Selmer ration $c(\pi_d)$.

Proposition 10.7. *Let $d \in \mathbb{Q}^\times$ be such that $P_{-27d}[\pi_{-27d}](\mathbb{Q}) = 0$, and write $c(\pi_d) = 3^m$. Then we have the congruence* $\dim_{\mathbb{F}_3}\mathrm{Sel}_3(P_d) \equiv m \pmod 2$.

Proof. By the Greenberg–Wiles formula, we have $\#\mathrm{Sel}(\pi_d)/\#\mathrm{Sel}(\hat{\pi}_d) = c(\pi_d) = 3^m$. Since $[-3] = \pi_d\circ\pi_{-27d}$, we have an exact sequence

$$0 \to \mathrm{Sel}(\pi_d) \to \mathrm{Sel}_3(P_d) \to \mathrm{Sel}(\pi_{-27d}) \to \frac{Ш(P_{-27d})[\pi_{-27d}]}{\pi_d(Ш(P_d)[3])} \to 0.$$

Exactness on the left is because $P_{-27d}[\pi_{-27d}](\mathbb{Q}) = 0$. By Lemma 10.6, there is an isomorphism $\mathrm{Sel}(\pi_{-27d}) \simeq \mathrm{Sel}(\hat{\pi}_d)$, so we see that

$$m \equiv \dim_{\mathbb{F}_3}\mathrm{Sel}_3(P_d) + \dim_{\mathbb{F}_3}\frac{Ш(P_{-27d})[\pi_{-27d}]}{\pi_d(Ш(P_d)[3])} \pmod 2. \tag{10-1}$$

Let

$$\langle\cdot,\cdot\rangle : Ш(P_{-27d}) \times Ш(\widehat{P}_{-27d}) \to \mathbb{Q}/\mathbb{Z}$$

be the Cassels–Tate pairing. Using the polarization $\lambda_{-27d} : P_{-27d} \to \widehat{P}_{-27d}$, define

$$\langle\cdot,\cdot\rangle_\lambda : Ш(P_{-27d}) \times Ш(P_{-27d}) \to \mathbb{Q}/\mathbb{Z}$$

by $\langle x, y\rangle_\lambda = \langle x, \lambda_{-27d}(y)\rangle$. As in [Shnidman 2021, Theorems 4.3, 4.4], if $x \in Ш(P_{-27d})[\pi_{-27d}]$, then y is in the image of $\pi_d : Ш(P_d) \to Ш(P_{-27d})$ if and only if $\langle x, \lambda_{-27d}(y)\rangle_\lambda = 0$ for all $y \in Ш(P_{-27d})[\pi_{-27d}]$. Thus, the Cassels–Tate pairing $\langle\cdot,\cdot\rangle_\lambda$ restricts to a nondegenerate paring on the finite group

$$\frac{Ш(P_{-27d})[\pi_{-27d}]}{\pi(Ш(P_d)[3])}.$$

Moreover, since both P_d and P_{-27d} are prime-to-3 polarized, this pairing is antisymmetric, and therefore alternating. The nondegeneracy implies that it has even $\mathbb{F}_3$-rank. Combining with (10-1), we deduce the desired congruence modulo 2. □

The following general lemma will be used to compute local Selmer ratios below.

Lemma 10.8. *Let $\alpha : A \to B$ be an isogeny of abelian varieties over a nonarchimedean characteristic* 0 *local field F. Then $c_\ell(\alpha)c_\ell(\hat{\alpha}) = \#(\mathcal{O}_F/\deg(\alpha)\mathcal{O}_F)$.*

Proof. By [Česnavičius 2017, B.1], the groups $B(F)/\alpha A(F)$ and $\hat{A}(F)/\hat{\alpha}\widehat{B}(F)$ are orthogonal complements under Tate–Shatz local duality

$$H^1(F, A[\alpha]) \times H^1(F, \widehat{B}[\hat{\alpha}]) \to \mathbb{Q}/\mathbb{Z}.$$

Thus

$$c_\ell(\alpha)c_\ell(\hat{\alpha}) = \frac{\#B(F)/\alpha A(F)}{\#A(F)[\alpha]} \cdot \frac{\#\hat{A}(F)/\hat{\alpha}\widehat{B}(F)}{\#\widehat{B}(F)[\hat{\alpha}]} = \frac{\#H^1(F, A[\alpha])}{\#A(F)[\alpha] \cdot \#\widehat{B}(F)[\hat{\alpha}]} = \#(\mathcal{O}_F/\deg(\alpha)\mathcal{O}_F),$$

where the final equality follows from the Euler–Poincaré characteristic formula. □

10B. ***The example.*** Now specialize to the context of Theorem 10.3 and the specific curves $C_d : y^3 = (x^2 - d)(x^2 - 4d)$.

The isogeny $\pi : P \to P_{-27}$ factors as

$$P \xrightarrow{\phi} B \xrightarrow{\psi} P_{-27},$$

where $B = P/\langle(1, 0) - (-1, 0)\rangle$. Since $(1, 0) - (-1, 0) \in P[\pi]$, we obtain twists $\phi_d : P_d \to B_d$ and $\psi_d : B_d \to P_{-27d}$ by Lemma 2.6.

Let us compute the local Selmer ratios for ϕ_d and ψ_d, for all $d \in \Sigma$ (where Σ is as in Theorem 10.3). For any $d \in \mathbb{Q}^\times$, we have $c_\infty(\phi_d) = c_\infty(\psi_d)$, since the kernels of ϕ and ψ are both $\mathbb{Z}/3\mathbb{Z}$. Note that P has good reduction at all $p > 3$, since C does. Thus, for $d \in \Sigma$ and for all $p \nmid 6\infty$, by Theorem 4.11, we have $c_p(\phi_d) = 1 = c_p(\psi_d)$. If $p = 2$, then since $d \equiv 2, 3 \pmod 4$, neither d nor $-3d$ is a square in $\mathbb{Q}_2$, so by Table 2, $c_2(\phi_d) = 1 = c_2(\psi_d)$ as well. To compute the ratios $c_3(\phi_d)$ and $c_3(\psi_d)$, we use Lemma 10.8. By multiplicativity, we have

$$c_3(\phi_d)c_3(\psi_d)c_3(\phi_{-27d})c_3(\psi_{-27d}) = c_3([3]) = 9.$$

Since $d, -27d \notin \mathbb{Q}_3^{\times 2}$, all four of these local Selmer ratios are integers, and the same is true for the dual isogenies. Thus, by Lemma 10.8, each ratio must be either 1 or 3. Hence, of the four Selmer ratios $c_3(\phi_d), c_3(\psi_d), c_3(\phi_{-27d}), c_3(\psi_{-27d})$, exactly two are 3 and the other two are 1.

Lemma 10.9. *If $d \in \Sigma$, then $c_3(\phi_d) \neq c_3(\psi_d)$.*

Proof. By Proposition 10.7, the parity of $\dim_{\mathbb{F}_3} \mathrm{Sel}_3(P_d)$ is odd if and only if $c(\pi_d) = c(\phi_d)c(\psi_d)$ is an odd power of 3. Hence, by our local computations at all other primes $p \neq 3$ given above, the parity of $\dim_{\mathbb{F}_3} \mathrm{Sel}_3(P_d)$ is odd if and only if $c_3(\phi_d) \neq c_3(\psi_d)$. Since $\mathrm{Jac}(C_d)$ is prime-to-3-isogenous to $P_d \times E_d$, we have $\mathrm{Sel}_3(J_d) = \mathrm{Sel}_3(P_d) \oplus \mathrm{Sel}_3(E_d)$. Also, letting $K = \mathbb{Q}(\zeta_3)$ and $X \in \{J_d, P_d, E_d\}$, we have

$$\dim \mathrm{Sel}_3(X) \equiv \dim \mathrm{Sel}_\pi(X) + \dim \mathrm{Sel}_{\pi_{-27}}(X_{-27}) = \dim \mathrm{Sel}_\pi(X_K) \pmod 2,$$

where π is the map induced by $1 - \zeta$ on divisors, and X_K is the base change to K. It follows that $c(\phi_d) \neq c(\psi_d)$ if and only if $\dim \mathrm{Sel}_\pi(J_{d,K}) - \dim \mathrm{Sel}_\pi(E_{d,K})$ is odd. The latter two π-Selmer groups can be computed in Magma [Bosma et al. 1997] for any choice of d, using the command `PhiSelmerGroup`.

In fact, it is enough to take $d = 2$, since the isomorphism class of A_d over $\mathbb{Q}_3$ depends only on the image of d in $\mathbb{Q}_3^\times/\mathbb{Q}_3^{\times 6}$, and all elements of Σ map to the sixth-power class of $d = 2$.[4] For $d = 2$, we find that $\dim \mathrm{Sel}_\pi(J_{d,K}) - \dim \mathrm{Sel}_\pi(E_{d,K}) = 1$. □

To recap, for $d \in \Sigma$, we have $c(\phi_d) = c_3(\phi_d)c_\infty(\phi_d)$ and $c(\psi_d) = c_3(\psi_d)c_\infty(\psi_d)$, and we know $c_\infty(\phi_d) = c_\infty(\psi_d)$ (which is equal to 1 or 1/3, depending on the sign of d) and $\{c_3(\phi_d), c_3(\psi_d)\} \subset \{1, 3\}$. Combining this with Lemma 10.9, we conclude that exactly one of $c(\phi_d)$ and $c(\psi_d)$ is 1 and the other is $3^{\pm 1}$.

Proof of Theorem 10.3. For all d, we have

$$\mathrm{rk}\, P_d(\mathbb{Q}) \le \mathrm{rk}(\mathrm{Sel}(\pi_d) \oplus \mathrm{Sel}(\hat{\pi}_d)) \le \mathrm{rk}(\mathrm{Sel}(\phi_d) \oplus \mathrm{Sel}(\hat{\phi}_d)) + \mathrm{rk}(\mathrm{Sel}(\psi_d) \oplus \mathrm{rk}\, \mathrm{Sel}(\hat{\psi}_d)).$$

For $d \in \Sigma$, we have seen that exactly one of $c(\phi_d)$ and $c(\psi_d)$ is 1 and the other is $3^{\pm 1}$. Thus, the average rank of $P_d(\mathbb{Q})$ for $d \in \Sigma$ is at most $\left(1 + \frac{4}{3}\right) = \frac{7}{3}$, by Proposition 5.4. This proves the first claim of Theorem 10.3.

Next we show that $\dim_{\mathbb{F}_3} \mathrm{Sel}_3(P_d) = 1$ for at least $\frac{1}{3}$ of $d \in \Sigma$. Without loss of generality we may assume that $c(\phi_d) = 1$ and $c(\psi_d) = 3^{\pm 1}$. By Proposition 5.4(iii), for at least $\frac{1}{2}$ of $d \in \Sigma$, we have $\mathrm{Sel}(\phi_d) = 0 = \mathrm{Sel}(\hat{\phi}_d)$, and for at least $\frac{5}{6}$ of $d \in \Sigma$, we have $\dim_{\mathbb{F}_3} \mathrm{Sel}(\psi_d) \oplus \mathrm{Sel}(\hat{\psi}_d) = 1$. Thus, for at least $\frac{5}{6} - \frac{1}{2} = \frac{1}{3}$ of $d \in \Sigma$, we have

$$\dim_{\mathbb{F}_3} \mathrm{Sel}_3(P_d) \le \dim \mathrm{Sel}(\phi_d) + \dim \mathrm{Sel}(\hat{\phi}_d) + \dim \mathrm{Sel}(\psi_d) + \dim \mathrm{Sel}(\hat{\psi}_d) \le 1.$$

This implies that $\mathrm{rk}\, P_d(\mathbb{Q}) \le 1$ for at least $\frac{1}{3}$ of $d \in \Sigma$. □

10C. ***More general curves.*** It is plausible that for *every* Prym P associated to some curve $C_{a,b} : y^3 = (x^2 - a)(x^2 - b)$, our method shows that $\mathrm{rk}\, P_d \le 1$ for a positive proportion of d. This holds if one can check a certain 3-adic condition on the numbers $c_3(\phi_d)$ and $c_3(\psi_d)$, exactly as in the proof of Lemma 10.9. This condition is satisfied in all examples we checked, but we do not have a proof in general. Since the local Selmer ratios $c_3(\phi_{a,b,d})$ and $c_3(\psi_{a,b,d})$ are locally constant as functions on $\mathbb{Q}_3^3 = \{(a, b, d)\}$, we can at least say that this condition holds for a large class of bielliptic Picard curves $C_{a,b}$, with a and b satisfying certain congruence conditions modulo a power of 3.

In [Shnidman and Weiss 2023], we prove that a positive proportion of P_d have rank at most 1, in the case where a/b is a square, using an extra argument which avoids the local 3-adic computation. In general, we have the following result, whose proof is an easier version of the argument given above, so we omit it.

Theorem 10.10. *Fix $a \in \mathbb{Q} \setminus \{0, \pm 1\}$. For $d \in \mathbb{Q}^\times/\mathbb{Q}^{\times 6}$, let $P_{a,d}$ be the Prym surface for the genus-3 curve $y^3 = (x^2 - d)(x^2 - ad)$. Then $\mathrm{rk}\, P_{a,d}(\mathbb{Q}) \le 2$ for a positive proportion of d.*

[4] To check this, use the fact that $\mathbb{Z}_3^{\times 6} = 1 + 9\mathbb{Z}_3$.

Acknowledgments

We are grateful to Nils Bruin, Jef Laga, Max Lieblich, Dino Lorenzini, Beth Malmskog, Zach Scherr, and Michael Stoll for helpful conversations and correspondences. We thank the referees for their helpful and detailed comments and corrections, and for suggesting to us cleaner proofs of Lemmas 3.8, 4.8, and 10.8. Shnidman was supported by the Israel Science Foundation (grant 2301/20). Weiss was supported by an Emily Erskine Endowment Fund postdoctoral fellowship at the Hebrew University of Jerusalem, by the Israel Science Foundation (grant 1963/20), and by the US-Israel Binational Science Foundation (grant 2018250).

References

[Arul 2021] V. Arul, "Division by $1-\zeta$ on superelliptic curves and Jacobians", *Int. Math. Res. Not.* **2021**:4 (2021), 3143–3185. MR Zbl

[Barth 1987] W. Barth, "Abelian surfaces with (1, 2)-polarization", pp. 41–84 in *Algebraic geometry* (Sendai, Japan, 1985), edited by T. Oda, Adv. Stud. Pure Math. **10**, North-Holland, Amsterdam, 1987. MR Zbl

[Bhargava 2004] M. Bhargava, "Higher composition laws, I: A new view on Gauss composition, and quadratic generalizations", *Ann. of Math.* (2) **159**:1 (2004), 217–250. MR Zbl

[Bhargava and Gross 2014] M. Bhargava and B. H. Gross, "Arithmetic invariant theory", pp. 33–54 in *Symmetry*: *representation theory and its applications*, edited by R. Howe et al., Progr. Math. **257**, Birkhäuser, New York, 2014. MR Zbl

[Bhargava and Ho 2016] M. Bhargava and W. Ho, "Coregular spaces and genus one curves", *Camb. J. Math.* **4**:1 (2016), 1–119. MR Zbl

[Bhargava and Shankar 2015a] M. Bhargava and A. Shankar, "Binary quartic forms having bounded invariants, and the boundedness of the average rank of elliptic curves", *Ann. of Math.* (2) **181**:1 (2015), 191–242. MR Zbl

[Bhargava and Shankar 2015b] M. Bhargava and A. Shankar, "Ternary cubic forms having bounded invariants, and the existence of a positive proportion of elliptic curves having rank 0", *Ann. of Math.* (2) **181**:2 (2015), 587–621. MR Zbl

[Bhargava et al. 2013] M. Bhargava, A. Shankar, and J. Tsimerman, "On the Davenport–Heilbronn theorems and second order terms", *Invent. Math.* **193**:2 (2013), 439–499. MR Zbl

[Bhargava et al. 2015] M. Bhargava, A. Shankar, and X. Wang, "Geometry-of-numbers methods over global fields, I: Prehomogeneous vector spaces", preprint, 2015. arXiv 1512.03035

[Bhargava et al. 2019] M. Bhargava, Z. Klagsbrun, R. J. Lemke Oliver, and A. Shnidman, "3-isogeny Selmer groups and ranks of abelian varieties in quadratic twist families over a number field", *Duke Math. J.* **168**:15 (2019), 2951–2989. MR Zbl

[Bhargava et al. 2020] M. Bhargava, N. Elkies, and A. Shnidman, "The average size of the 3-isogeny Selmer groups of elliptic curves $y^2 = x^3 + k$", *J. Lond. Math. Soc.* (2) **101**:1 (2020), 299–327. MR Zbl

[Bhargava et al. $\geq$ 2025] M. Bhargava, A. Shankar, and X. Wang, "Geometry-of-numbers methods over global fields, II: Coregular representations", in preparation.

[Bosma et al. 1997] W. Bosma, J. Cannon, and C. Playoust, "The Magma algebra system, I: The user language", *J. Symbolic Comput.* **24**:3-4 (1997), 235–265. MR Zbl

[Brumer 1992] A. Brumer, "The average rank of elliptic curves, I", *Invent. Math.* **109**:3 (1992), 445–472. MR Zbl

[Caro and Pasten 2023] J. Caro and H. Pasten, "A Chabauty–Coleman bound for surfaces", *Invent. Math.* **234**:3 (2023), 1197–1250. MR Zbl

[Česnavičius 2016] K. Česnavičius, "Selmer groups as flat cohomology groups", *J. Ramanujan Math. Soc.* **31**:1 (2016), 31–61. MR Zbl

[Česnavičius 2017] K. Česnavičius, "p-Selmer growth in extensions of degree p", *J. Lond. Math. Soc.* (2) **95**:3 (2017), 833–852. MR Zbl

[Delone and Faddeev 1940] B. N. Delone and D. K. Faddeev, "Theory of irrationalities of third degree", *Acad. Sci. URSS. Trav. Inst. Math. Stekloff*, **11** (1940), 3–340. In Russian. MR Zbl

[Diaconu and Tian 2005] A. Diaconu and Y. Tian, "Twisted Fermat curves over totally real fields", *Ann. of Math.* (2) **162**:3 (2005), 1353–1376. MR Zbl

[Dimitrov et al. 2021] V. Dimitrov, Z. Gao, and P. Habegger, "Uniformity in Mordell–Lang for curves", *Ann. of Math.* (2) **194**:1 (2021), 237–298. MR Zbl

[Gan et al. 2002] W. T. Gan, B. Gross, and G. Savin, "Fourier coefficients of modular forms on G_2", *Duke Math. J.* **115**:1 (2002), 105–169. MR Zbl

[Gao et al. 2021] Z. Gao, T. Ge, and L. Kühne, "The uniform Mordell–Lang conjecture", preprint, 2021. arXiv 2105.15085

[Heath-Brown 1994] D. R. Heath-Brown, "The size of Selmer groups for the congruent number problem, II", *Invent. Math.* **118**:2 (1994), 331–370. MR Zbl

[Katz and Sarnak 1999] N. M. Katz and P. Sarnak, *Random matrices, Frobenius eigenvalues, and monodromy*, Amer. Math. Soc. Colloq. Publ. **45**, Amer. Math. Soc., Providence, RI, 1999. MR Zbl

[Klagsbrun et al. 2013] Z. Klagsbrun, B. Mazur, and K. Rubin, "Disparity in Selmer ranks of quadratic twists of elliptic curves", *Ann. of Math.* (2) **178**:1 (2013), 287–320. MR Zbl

[Kühne 2021] L. Kühne, "Equidistribution in families of abelian varieties and uniformity", preprint, 2021. arXiv 2101.10272

[Laga 2023] J. Laga, "Arithmetic statistics of Prym surfaces", *Math. Ann.* **386**:1-2 (2023), 247–327. MR Zbl

[Laga and Shnidman 2023] J. Laga and A. Shnidman, "The geometry and arithmetic of bielliptic Picard curves", preprint, 2023. arXiv 2308.15297

[Levi 1914] F. Levi, "Kubische Zahlkörper und binäre kubische Formenklassen", *Sitz.ber. Sächs. Akad. Wiss. Leipz. Math.-Nat.wiss. Kl.* **66** (1914), 26–37. Zbl

[Mazur and Rubin 2007] B. Mazur and K. Rubin, "Finding large Selmer rank via an arithmetic theory of local constants", *Ann. of Math.* (2) **166**:2 (2007), 579–612. MR Zbl

[Milne 1986] J. S. Milne, *Arithmetic duality theorems*, Persp. Math. **1**, Academic Press, Boston, MA, 1986. MR Zbl

[Mumford 1974] D. Mumford, "Prym varieties, I", pp. 325–350 in *Contributions to analysis*, edited by L. V. Ahlfors et al., Academic Press, New York, 1974. MR Zbl

[Neukirch et al. 2000] J. Neukirch, A. Schmidt, and K. Wingberg, *Cohomology of number fields*, Grundl. Math. Wissen. **323**, Springer, 2000. MR Zbl

[Schaefer 2018] E. F. Schaefer, "Explicit descent for Jacobians of prime power cyclic covers of the projective line", *Trans. Amer. Math. Soc.* **370**:5 (2018), 3487–3505. MR Zbl

[Serre 1977] J.-P. Serre, *Linear representations of finite groups*, Grad. Texts in Math. **42**, Springer, 1977. MR Zbl

[Shnidman 2021] A. Shnidman, "Quadratic twists of abelian varieties with real multiplication", *Int. Math. Res. Not.* **2021**:5 (2021), 3267–3298. MR Zbl

[Shnidman and Weiss 2023] A. Shnidman and A. Weiss, "Rank growth of elliptic curves over n-th root extensions", *Trans. Amer. Math. Soc. Ser. B* **10** (2023), 482–506. MR Zbl

[Smith 2017] A. Smith, "2^∞-Selmer groups, 2^∞-class groups, and Goldfeld's conjecture", preprint, 2017. arXiv 1702.02325

[Stoll 2006] M. Stoll, "Independence of rational points on twists of a given curve", *Compos. Math.* **142**:5 (2006), 1201–1214. MR Zbl

[Thorne 2013] J. A. Thorne, "Vinberg's representations and arithmetic invariant theory", *Algebra Number Theory* **7**:9 (2013), 2331–2368. MR Zbl

Communicated by Melanie Matchett Wood
Received 2022-05-27 Revised 2023-12-12 Accepted 2024-01-22

ariel.shnidman@mail.huji.ac.il *Einstein Institute of Mathematics, The Hebrew University of Jerusalem, Jerusalem, Israel*

weiss.742@osu.edu *Department of Mathematics, The Ohio State University, Columbus, OH, United States*

Department of Mathematics, Ben-Gurion University of the Negev, Be'er Sheva, Israel

Algebra & Number Theory

msp.org/ant

See inside back cover or msp.org/ant for submission instructions.

The subscription price for 2025 is US $565/year for the electronic version, and $820/year (+$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP.

Algebra & Number Theory (ISSN 1944-7833 electronic, 1937-0652 printed) at Mathematical Sciences Publishers, 798 Evans Hall #3840, c/o University of California, Berkeley, CA 94720-3840 is published continuously online.

ANT peer review and production are managed by EditFLOW® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

Algebra & Number Theory

Volume 19 No. 1 2025

A modification of the linear sieve, and the count of twin primes 1
JARED DUKER LICHTMAN

Ranks of abelian varieties in cyclotomic twist families 39
ARI SHNIDMAN and ARIEL WEISS

Picard rank jumps for K3 surfaces with bad reduction 77
SALIM TAYOU

Curves with few bad primes over cyclotomic $\mathbb{Z}_\ell$-extensions 113
SAMIR SIKSEK and ROBIN VISSER

Vanishing results for the coherent cohomology of automorphic vector bundles over the Siegel variety in positive characteristic 143
THIBAULT ALEXANDRE

Super-Hölder vectors and the field of norms 195
LAURENT BERGER and SANDRA ROZENSZTAJN