Hilbert's tenth problem over the rational numbers

OpenAI

October 6, 2026

Abstract

We give a negative answer to Hilbert's tenth problem over the rational numbers: no algorithm decides whether a polynomial with integer coefficients has a rational zero. The number of variables is part of the input.

Contents

1 Introduction **3**
- 1.1 Historical setting . . . 3
- 1.2 Proof strategy . . . 5

2 Finite rational tests for integer solvability **6**
- 2.1 Elliptic indices and the local comparison . . . 6
- 2.2 Contact primes and ordinary integer witnesses . . . 9
- 2.3 A recursive theory with an ordinary integer model . . . 11
- 2.4 Finite rational tests and compactness . . . 12
- 2.5 Odd valuations select a tested branch . . . 14
- 2.6 Index congruences at every contact prime . . . 15
- 2.7 The height contradiction . . . 16

3 An existential condition excluding odd poles **18**
- 3.1 The formula and the pole restriction . . . 18
- 3.2 Preparing integer witnesses . . . 19
- 3.3 A criterion for one twist . . . 20
- 3.4 One parameter for both twists . . . 24

4 A height bound from odd contact with five points **26**
- 4.1 The five-point quaternionic quotient . . . 27
- 4.2 The abelian surface family . . . 30
- 4.3 Rational fibers and odd-contact ramification . . . 31
- 4.4 Comparison of base height and Faltings height . . . 32
- 4.5 A field of definition for homomorphisms . . . 35
- 4.6 The endomorphism dichotomy and descent up to isogeny . . . 35
- 4.7 Two-dimensional representations and modularity . . . 38
- 4.8 A uniform bound for the modular level . . . 40
- 4.9 The modular factor and the final height bound . . . 42

5 The rank-one elliptic group and its local parameters **44**

6 Prime patterns for the two conjugate local tests **46**
- 6.1 Reduction orders and their possible common factors . . . 48
- 6.2 A finite multiplier set and the initial residue conditions . . . 49
- 6.3 A lower bound for the prime patterns . . . 51
- 6.4 An elementary upper sieve with uniform errors . . . 52
- 6.5 Counting bad prime values . . . 54
- 6.6 A uniform sieve on each affine fiber . . . 55
- 6.7 Completing the prime-pattern construction . . . 55

7 Turing degree and a quartic normal form **56**

References **57**

1 Introduction

Hilbert's tenth problem over the rational numbers asks for a decision procedure for rational zeros of polynomials with integer coefficients. The number of variables is part of the input.

Theorem 1.1. *There is no algorithm which, given a polynomial $f \in \mathbb{Z}[X_1, \ldots, X_n]$, with n part of the input, decides whether f has a zero in $\mathbb{Q}^n$.*

1.1 Historical setting

Hilbert's tenth problem, as stated in his 1900 list of mathematical problems, asks for a finite procedure deciding whether a polynomial equation with integer coefficients has a solution in the ordinary integers [31, Problem 10]. Its negative resolution grew out of a much stronger description of polynomial equations. A relation on the positive integers is *recursively enumerable* if an algorithm can list exactly its satisfying tuples. It is *Diophantine* if membership can be expressed by the existence of additional positive integers satisfying a polynomial equation with integer coefficients. Thus a Diophantine description expresses a condition entirely through polynomial solvability.

Davis, Putnam, and Robinson proved in 1961 that every recursively enumerable relation has such a representation if exponentiation is also allowed. They reduced the ordinary polynomial representation theorem to a growth criterion for a Diophantine relation [14, Theorem and Corollary 3]. Matiyasevich supplied the required growth in 1970 through a Diophantine description of a Fibonacci relation, completing the theorem that every recursively enumerable relation on the positive integers is Diophantine [37]. The equivalent all-integer decision problem is therefore undecidable. This representation theorem explains the strength of the result: integer polynomial equations can encode arbitrary recursively enumerable conditions.

Letting the unknowns range over $\mathbb{Q}$ changes the decision problem. Enumerating rational tuples will eventually find any rational zero, but gives no answer for an equation without one. A possible transfer of the integer theorem would be an existential definition of membership in $\mathbb{Z}$ using only the ring operations of $\mathbb{Q}$. One could then replace each integer unknown by a rational unknown subject to that membership condition, obtaining an equivalent rational existential question. More generally, a *Diophantine model* of integer arithmetic represents the integers by a Diophantine subset of some $\mathbb{Q}^k$, with Diophantine relations for addition and multiplication. Such a model also translates integer equations into rational existential questions [8, Definition 3.1 and Observation 3.3]. Here the defining polynomial equations and all their witnesses are interpreted over $\mathbb{Q}$. These constructions are sufficient routes to undecidability; the decision problem does not require its solution to produce either one.

Julia Robinson showed in 1949 that the integers are first-order definable in the rationals using addition and multiplication. Her formula uses both universal and existential quantifiers, and her work also proves undecidability of the full first-order theory of rational arithmetic [60, Theorems 3.1 and 4.2]. Later work reduced the quantifier complexity. Poonen defined $\mathbb{Z}$ in $\mathbb{Q}$ by a positive formula with two universal quantifiers followed by seven existential quantifiers [54, Theorem 4.1]. Koenigsmann then gave a universal definition of $\mathbb{Z}$; equivalently, its complement in $\mathbb{Q}$ is Diophantine. His result also implies undecidability of the $\forall\exists$ theory [33, Theorem 1 and Corollaries 2–3]. These results give substantial control of definitions inside $\mathbb{Q}$, but do not supply the existential membership condition for the direct transfer above.

The distinction between a Diophantine model and undecidability also has a geometric setting. Mazur conjectured that, for a variety over $\mathbb{Q}$, the closure of its rational points in its real points has only finitely many connected components [38, Conjecture 3]. He observed that this conjecture would rule out a Diophantine definition of $\mathbb{Z}$ in $\mathbb{Q}$. Cornelissen and

Zahidi showed that it would also rule out a Diophantine model of integer arithmetic in a finite Cartesian power of $\mathbb{Q}$ [8, Theorem 3.6]. These obstructions to definitions and models do not by themselves rule out a negative answer to the rational decision problem.

One approach studies the subrings $\mathbb{Z}[S^{-1}] \subseteq \mathbb{Q}$, whose denominators use only primes in S. Poonen constructed disjoint recursive sets of primes T_1, T_2, each of natural density zero, such that positive-integer arithmetic has a Diophantine model in $\mathbb{Z}[S^{-1}]$ whenever $S \supseteq T_1$ and $S \cap T_2 = \varnothing$. Hilbert's tenth problem is undecidable for every such ring. Taking S to be the complement of T_2 gives a recursive set of density one and a proper subring of $\mathbb{Q}$ [53, Theorem 1.3]. The construction uses selected multiples of a point on a rank-one elliptic curve whose real coordinates approximate integers.

Eisenträger, Miller, Park, and Shlapentokh later constructed computably presentable subrings $\mathbb{Z}[S^{-1}]$, including examples in which S has density one, whose Hilbert's tenth problem is Turing equivalent to that over $\mathbb{Q}$ [19, Theorem 3.17]. For these examples, a Turing reduction from integer solvability would already give a Turing reduction from integer solvability to the rational field problem. The density of the inverted primes alone therefore does not settle the relation to the field problem.

Elliptic curves also give a transfer between rings of integers of number fields. Poonen proved that, for $F \subseteq K$, an elliptic curve E/F with

$$\operatorname{rank} E(F) = \operatorname{rank} E(K) = 1$$

makes $\mathcal{O}_F$ Diophantine in $\mathcal{O}_K$ [52, Theorem 1]. If $\mathbb{Z}$ is already Diophantine in $\mathcal{O}_F$, undecidability transfers to $\mathcal{O}_K$. Mazur and Rubin pursued this route through quadratic twists and Selmer groups, conditional on the conjectured evenness of $\dim_{\mathbb{F}_2} \mathrm{Sha}(E/L)[2]$. Assuming this for every elliptic curve over every number field, they obtained undecidability for every infinite finitely generated $\mathbb{Z}$-algebra [39, Theorem 1.13, Corollary 1.14, and Theorem 8.1].

This number-field program has recently been completed unconditionally. Koymans and Pagano combined two-descent with additive combinatorics to prove that $\mathbb{Z}$ is Diophantine in the ring of integers of every number field, and that Hilbert's tenth problem is undecidable for every infinite finitely generated $\mathbb{Z}$-algebra [34]. Alpöge, Bhargava, Ho, and Shnidman gave another route: for every quadratic extension K/F of number fields, they construct an abelian variety over F with positive rank unchanged over K, and deduce that $\mathbb{Z}$ is Diophantine in every number-field ring of integers [2, Theorem 1.1 and Corollary 1.2]. The field $\mathbb{Q}$ is not a finitely generated $\mathbb{Z}$-algebra, so these results concern a different class of rings from the field in Theorem 1.1.

The elliptic approach also illuminates nearby rational decision problems. In a rank-one group, a multiple nP of a nontorsion point can represent the integer n, and local coordinates or heights can reveal arithmetic information about that index. Cornelissen and Zahidi used elliptic divisibility sequences and existential valuation predicates in this direction. Under a conjectural odd primitive-divisor property for one such sequence, they obtained a Diophantine model of $(\mathbb{Z}, +, |)$ and proved undecidability of a positive Π_2 fragment with one universal quantifier [9, Conjecture 3.16 and Theorems 3.17 and 5.3]. The assertion with one universal quantifier uses the correction by Cornelissen and Shlapentokh [7, Remark 2.3 in the author version].

Garcia-Fritz, Pasten, and Vidaux instead adjoined predicates comparing the logarithmic heights of rational tuples to the rational ring language. In this expanded language they gave a positive existential interpretation of $\mathbb{N}$, and hence proved undecidability; comparisons between tuples of length at most three suffice [22, Theorems 1.1–1.2]. Their construction uses elliptic canonical heights to encode consecutive squares. The present paper concerns ordinary rational polynomial solvability, and its height estimates enter the proof of that ordinary statement.

1.2 Proof strategy

Constant polynomials can be decided directly. For each nonconstant integral polynomial f, we construct an effectively generated sequence of finite tests. Each test can be answered by finitely many rational-solvability queries, and the sequence satisfies

$$f \text{ has an integer zero} \quad \Longleftrightarrow \quad \text{every test succeeds.} \tag{1.1}$$

An integer zero supplies witnesses for all the tests. If a rational-solvability algorithm existed, we could run two searches in parallel: enumeration of integer zeros, and enumeration of tests until one fails. The arithmetic work is to prove that a failed test exists whenever f has no integer zero.

Suppose all the finite tests succeed. Compactness then gives a ring R containing a zero $\mathbf{a} = (a_1, \ldots, a_n)$ of f, embedded in an elementary extension of the rational field; we use a tilde for embedded values and points. The ambient extension also carries the integer and valuation data, height functions, elliptic multiple maps, and functions for finite products used in the proof. The tests attach an elliptic point T_a to each $a \in R$, compatibly with addition, and require T_a to be nonidentity when $a \neq 0$. The embedded points lie in the transferred copy of a fixed infinite cyclic subgroup $\mathbb{Z}P$ of finite index in a rank-one elliptic group, with $T_1 = P$. This defines unique indices

$$\widetilde{T}_a = \eta(a)P, \qquad \eta(a) \in {}^*\mathbb{Z}.$$

Here ${}^*\mathbb{Z}$ is the ambient integer structure, whose elements may be nonordinary. The point axioms make $\eta : R \to {}^*\mathbb{Z}$ injective and additive, with $\eta(1) = 1$, but η need not respect multiplication. Thus the modeled equation does not yet give an equation satisfied by the indices.

Two local comparisons address this obstruction. The primary test compares an embedded ring value with a quotient of truncated elliptic logarithms. The formal logarithm can recover the corresponding quotient of point indices once that quotient is integral at the prime. A second comparison, on the conjugate elliptic curve, supplies this control using only additivity of η. Where the comparisons apply, they show that every $\widetilde{a}$ is integral at the prime and agrees with $\eta(a)$ to the prescribed precision. Each pair uses auxiliary choices that work for every ring element, which is essential for the later uniform bound.

The proof arranges these comparisons at primes detected by a height estimate. Fix the five rational points of $\mathbf{P}^1$ supplied by that estimate and primitive integral linear forms vanishing at them; call these the contact forms. For a rational point $s = (u : v)$ outside this set, written in coprime integer coordinates, let $M(s)$ be the product, without repetition, of the nonexceptional primes at which one of these forms has odd valuation. We call the primes counted by $M(s)$ its contact primes. The estimate bounds the absolute logarithmic Weil height by

$$h(s) \leq HM(s)^c$$

for fixed constants H, c. This definition and estimate transfer to the ambient extension. At each prescribed precision, the theory represents every ring element as a sum of three fractions of ring elements, called slopes, and equips factors of their contact forms with the local conditions. A prime-pattern theorem supplies the required primes for these representations in the ordinary integer model. A separate existential condition controls the parity of possible denominator valuations outside a fixed finite set. In a model satisfying the tests, consider the three slopes supplied for any ring element at that precision. After the further fixed exclusions, the parity condition lets the proof apply the local comparisons at every contact prime of those slopes. Enlarging the exceptional set changes the height constant but preserves the exponent c.

Set

$$\delta = f(\eta(a_1), \dots, \eta(a_n)), \qquad B = \max(1, |\eta(a_1)|, \dots, |\eta(a_n)|).$$

If $\delta = 0$, elementarity gives an ordinary integer zero of f. Otherwise δ is a nonzero ambient integer whose size is bounded by a fixed polynomial in B. Choose one ordinary precision K large enough in terms of the degree of f and the exponent c. For any $A \in R$, choose a tested representation at precision K. Every contact prime of its slopes divides this same δ to order at least K, because all modeled root coordinates agree locally with their indices. The height estimate then bounds each slope height by a constant times B. The logarithmic height inequality for their sum gives

$$h(\widetilde{A}) \ll_f B \qquad \text{for every } A \in R,$$

with one ordinary constant.

For each nonzero root coordinate a_j, consider its attached point T_{a_j}. Applying the uniform bound to the ring elements representing $x(T_{a_j})$ gives $h(x(\eta(a_j)P)) \ll_f B$. The transferred canonical-height comparison for P makes this height grow quadratically in $|\eta(a_j)|$. If $B > 1$, choose j with $|\eta(a_j)| = B$; the two bounds give $B^2 \ll_f B$. If $B = 1$, the required bound is immediate. Thus all root indices lie in an ordinary finite interval. Additivity and injectivity identify the modeled root coordinates with those ordinary integers, contradicting $\delta \neq 0$. This proves the finite-test equivalence and hence the undecidability theorem.

The detailed reduction in Section 2 makes these steps precise. Section 3 proves the pole-parity condition; its passage from two-descent data to rational points uses the pointwise 2-converse [48, Theorem 1.1] and the Cassels–Tate pairing [6]. Section 4 proves the five-point height estimate. Its modularity step uses the odd regular Fontaine–Mazur theorem at 2 [49, Theorem 1.1], whose statement has no hypothesis on the residual image. Section 5 proves the rank-one and local elliptic results, and Section 6 supplies the prime patterns and simultaneous local witnesses for the ordinary integer model. Section 7 gives Turing degree $0'$ and the quartic normal form.

At rational primes, valuations are normalized by $v_q(q) = 1$, with $v_q(0) = +\infty$. The support of a nonzero rational number is the set of primes dividing its numerator or denominator in lowest terms. The notation $X \ll Y$ means $X \leq CY$ for a positive constant depending only on the fixed data in the statement; additional dependencies are indicated when they occur.

2 Finite rational tests for integer solvability

We construct the finite rational tests announced in the introduction and prove that success of every test forces an ordinary integer zero. The construction has an ordinary integer model. Conversely, compactness will embed a model of all the tests in an elementary extension of $\mathbb{Q}$, where arithmetic restrictions on its ring values will force an integer solution. The later sections prove the arithmetic inputs used here.

The elliptic point data first give an additive assignment of integer indices. A local comparison will recover those indices from ring values and determine the two local test conditions.

2.1 Elliptic indices and the local comparison

Fix the number field, involution, and elliptic curve

$$F = \mathbb{Q}(\sqrt{2}), \qquad d = 75 - 53\sqrt{2}, \qquad E: \; y^2 = x^3 - d^2 x, \tag{2.1}$$

where σ is the nontrivial automorphism of $F/\mathbb{Q}$. The following result, proved in Section 5, supplies the integer indices.

Proposition 2.1. *The group $E(F)$ has rank one. Consequently there exist a positive integer m and a nontorsion point $P \in E(F)$ such that*

$$mE(F) = \mathbb{Z}P.$$

We fix such m and P once and for all.

We now describe the point data that the recursive theory will impose. Let R be an integral domain of characteristic zero for which $\sqrt{2} \notin \mathrm{Frac}(R)$, and write

$$R_F = R[\sqrt{2}], \qquad F_R = \mathrm{Frac}(R)(\sqrt{2}).$$

For each $a \in R$, consider a point $T_a \in E(F_R)$ subject to

$$\begin{gathered} T_0 = O, \qquad T_1 = P, \qquad T_{a+a'} = T_a + T_{a'} \quad (a, a' \in R), \\ T_a \in mE(F_R) \quad (a \in R), \\ a \neq 0 \Longrightarrow T_a \text{ is finite and } x(T_a)y(T_a) \neq 0. \end{gathered} \tag{2.2}$$

The theory will supply six unary ring-valued functions $x_0, x_1, x_2, y_0, y_1, y_2$. For $a \neq 0$, they are required to satisfy $x_2(a)y_2(a) \neq 0$ and

$$x(T_a) = \frac{x_0(a) + x_1(a)\sqrt{2}}{x_2(a)}, \qquad y(T_a) = \frac{y_0(a) + y_1(a)\sqrt{2}}{y_2(a)}. \tag{2.3}$$

Their values at $a = 0$ are unrestricted, since $T_0 = O$. In the integer model we will take $T_a = aP$. Representing the coordinates by ring elements will also let a uniform height bound on R control their heights.

For the arithmetic analysis, consider an elementary extension of the ordinary many-sorted structure containing $\mathbb{Q}, \mathbb{Z}, \mathbb{R}, F$, the prime and valuation relations, the height functions, the elliptic multiple maps, and the finite-product functions. We call elements of the original structure *ordinary*; starred sorts satisfy the same first-order statements with ordinary parameters. For the moment, suppose that R embeds in the rational sort ${}^*\mathbb{Q}$ of such an extension. Its fraction-field interpretation embeds F_R in ${}^*F$. A tilde denotes the image of a ring element, field element, or point. Later compactness will supply an extension and an embedded ring with these properties.

The identity $mE(F) = \mathbb{Z}P$ transfers to this extension: every point of $mE({}^*F)$ has a unique expression nP, with $n \in {}^*\mathbb{Z}$. Define

$$\widetilde{T}_a = \eta(a)P. \tag{2.4}$$

The group-law axioms give

$$\eta(a + a') = \eta(a) + \eta(a'), \qquad \eta(1) = 1. \tag{2.5}$$

If $\eta(a) = 0$, then $T_a = O$, and the finite-point axiom forces $a = 0$. Thus $\eta : R \to {}^*\mathbb{Z}$ is injective and additive. Conjugating (2.4) gives $\widetilde{T}_a^{\sigma} = \eta(a)P^{\sigma}$. We make no assumption about $\eta(aa')$.

For nonzero $h \in R$, the index ratio of T_{ha} and T_h is $\eta(ha)/\eta(h)$, which need not equal $\eta(a)$. We will recover this ratio locally from point parameters, and use a second comparison to identify it with $\eta(a)$ to the required precision.

For a finite point with $y \neq 0$, put $z = -x/y$, and set $z(O) = 0$. Every nonzero multiple of P avoids the two-torsion points, so $z(nP)$ is defined for every integer n. Write

$$\ell(Z) = Z + \sum_{t \geq 2} c_t Z^t \in F[[Z]], \qquad \ell_K(Z) = Z + \sum_{2 \leq t \leq K} c_t Z^t$$

for the formal logarithm in this parameter and its truncation through degree K. The truncations are effectively computable from the fixed equation. Section 5 proves the following uniform statements about integer multiples of the fixed F-rational points P, P^σ.

Proposition 2.2. *There is a fixed finite set of rational primes S_E with the following properties. Let $p \notin S_E$, let v be any place of F above p, normalized by $v(p) = 1$, and let $n, n' \in \mathbb{Z}$, with $n' \neq 0$. Assume $nP, n'P$ reduce to O at v. Then*

$$v\left(\frac{z(nP)}{z(n'P)}\right) = v_p(n/n'). \tag{2.6}$$

For every integer $K \geq 1$, if in addition $p > K + 1$ and $v_p(n/n') \geq 0$, then $\ell_K(z(n'P)) \neq 0$ and

$$v\left(\frac{\ell_K(z(nP))}{\ell_K(z(n'P))} - \frac{n}{n'}\right) \geq K. \tag{2.7}$$

Both assertions also hold at the same place v after replacing E, P, ℓ_K by $E^\sigma, P^\sigma, \ell_K^\sigma$ and imposing the corresponding reduction-to-O hypotheses on the conjugate points. Here and below the valuation of zero is $+\infty$, so $n = 0$ is included.

The two estimates provide different kinds of local information. Equation (2.6) determines the valuation of an index ratio from the point parameters without assuming that the ratio is integral. Once an index ratio is known to be integral, (2.7) approximates it by a ratio of finite polynomial evaluations. The next lemma combines both forms of control.

Reduction to O has its usual valuation meaning, transferred to the elementary extension. Only the stated assertions about integer multiples of the fixed F-rational points P, P^σ and evaluations of the finite polynomials ℓ_K, ℓ_K^σ are transferred; no infinite formal series is evaluated in the extension.

Lemma 2.3 (Local comparison)**.** *Let $K \geq 1$ be an ordinary integer. Let q be a prime of the ambient integer sort ${}^*\mathbb{Z}$, possibly nonordinary, with $q \notin S_E$ and $q > K + 1$, and let v_* be a place of ${}^*F$ above q, normalized by $v_*(q) = 1$. Suppose there are nonzero $h, j_0 \in R$, chosen independently of a, such that for every $a \in R$:*

(i) The points $\widetilde{T}_h, \widetilde{T}_{ha}$ reduce to O at v_*, the value $\ell_K(z(\widetilde{T}_h))$ is nonzero, and

$$v_*\left(\frac{\ell_K(z(\widetilde{T}_{ha}))}{\ell_K(z(\widetilde{T}_h))} - \widetilde{a}\right) \geq K.$$

(ii) The points $\widetilde{T}^\sigma_{j_0}$ and $\widetilde{T}^\sigma_{(h-4)a}$ reduce to O on the conjugate curve at the same place v_*, and

$$v_*\left(\frac{z(\widetilde{T}^\sigma_{(h-4)a})}{z(\widetilde{T}^\sigma_{j_0})}\right) \geq K.$$

Then every $\widetilde{a}$ is integral at q, and

$$v_q(\widetilde{a} - \eta(a)) \geq K \qquad (a \in R).$$

Here $z(O) = 0$. The denominator in (ii) is nonzero because $j_0 \neq 0$ and the nonzero points in (2.2) have nonzero x- and y-coordinates.

Proof. Transfer the formal-parameter valuation equality of Proposition 2.2 to the conjugate points at v_*. The comparison in (ii) gives

$$v_q\left(\frac{\eta((h-4)a)}{\eta(j_0)}\right) \geq K.$$

The denominator is a nonzero ambient integer, so its q-valuation is nonnegative. Consequently

$$v_q(\eta((h-4)a)) \geq K. \tag{2.8}$$

Put $V_a = \eta(ha)$. Additivity applied to $(h-4)a = ha - 4a$, including the case $a = 1$, gives

$$V_a \equiv 4\eta(a) \pmod{q^K}, \qquad V_1 \equiv 4 \pmod{q^K}.$$

Since $q \neq 2$, V_1 is a q-unit. The ambient integers $V_a, \eta(a)$ are q-integral, and hence

$$V_a/V_1 \text{ is } q\text{-integral}, \qquad V_a/V_1 \equiv \eta(a) \pmod{q^K}. \tag{2.9}$$

The primary points have indices V_1 and V_a. They reduce to O, the denominator index is nonzero, and their ratio is now known to be integral. The truncated-logarithm assertion of Proposition 2.2 therefore gives

$$v_*\left(\frac{\ell_K(z(\widetilde{T}_{ha}))}{\ell_K(z(\widetilde{T}_h))} - \frac{V_a}{V_1}\right) \geq K.$$

The comparison in (i) relates this same logarithm ratio to $\widetilde{a}$. Combining it with the last display and (2.9) proves the required congruence, since v_* restricts to v_q on ${}^*\mathbb{Q}$. It also proves q-integrality of $\widetilde{a}$. Zero numerator indices are covered by the convention $v_q(0) = +\infty$. □

The same h must serve every a, because the calculation uses both a and 1. We will impose ring-language versions of these comparisons and obtain their valuation hypotheses at primes detected by a height estimate.

2.2 Contact primes and ordinary integer witnesses

Throughout the paper, h denotes the absolute logarithmic Weil height. The next input identifies primes whose product controls the height of a rational point. It is proved in Section 4.

Theorem 2.4 (Five-point height estimate)**.** *There exist five distinct points $\mathcal{C} \subset \mathbf{P}^1(\mathbb{Q})$, including ∞, a finite set S of rational primes, and constants $H, c > 0$ with the following property. For each $b \in \mathcal{C}$, choose a primitive integral linear form $L_b(U, V)$ vanishing at b, with $L_\infty(U, V) = V$. For $s = (u : v) \in \mathbf{P}^1(\mathbb{Q}) \setminus \mathcal{C}$, let $M_S(s)$ be the product, taken without repetition, of primes $q \notin S$ for which at least one of the integers*

$$v_q(L_b(u, v)) - \min\{v_q(u), v_q(v)\}, \qquad b \in \mathcal{C}, \tag{2.10}$$

is odd. Here $v_q(q) = 1$ and $v_q(0) = +\infty$. Then

$$h(s) \leq H M_S(s)^c. \tag{2.11}$$

Multiplying (u, v) by a nonzero rational scalar adds the same valuation to both terms in (2.10), so each contact order is independent of the chosen homogeneous coordinates. It is nonnegative after scaling the coordinates to be integral and primitive at q. With fixed primitive integral coordinates, the nonzero integers $L_b(u, v)$ have only finitely many prime divisors, so the product defining $M_S(s)$ is finite.

For a positive integer K and a nonzero integer e satisfying $M_S(s)^K \mid e$, the estimate gives

$$h(s) \le H|e|^{c/K}.$$

In the reduction, the same polynomial error will play the role of e for representations of every ring element. This is why the exponent c must be fixed, although its precise value is not needed. The following immediate consequence permits the further fixed exclusions needed by the local tests.

Corollary 2.5. *If $S' \supset S$ is any fixed finite set of primes, then*

$$h(s) \le H\left(\prod_{q \in S' \setminus S} q\right)^c M_{S'}(s)^c \qquad (s \in \mathbf{P}^1(\mathbb{Q}) \setminus \mathcal{C}).$$

In particular enlarging the exceptional set changes the constant but leaves the exponent unchanged.

Proof. Every prime counted by $M_S(s)$ is either counted by $M_{S'}(s)$ or belongs to $S' \setminus S$. Thus $M_S(s) \le (\prod_{q \in S' \setminus S} q) M_{S'}(s)$, and Theorem 2.4 applies. □

To apply this estimate uniformly, the theory will write each ring element as a sum of three fractions and impose local conditions on factors of their contact forms. The following ordinary integer result supplies witnesses for those axioms. It is a direct consequence of the prime-pattern lemma and the Chinese remainder calculation proved in Section 6.

Corollary 2.6 (Ordinary contact representations)**.** *Let $\mathcal{K} \subset \mathbb{Z}_{>0}$ be the finite multiplier set supplied by Lemma 6.1 in Section 6. It has the following property. For every integer $K \ge 1$, every $A \in \mathbb{Z}$, and every real lower bound $Y_0 > 0$, there exist integers u_1, u_2, v, with $v \ne 0$, such that, on setting $u_3 = Av - u_1 - u_2$, the following hold.*

(1) Representation and primes. *Each of the thirteen numbers*

$$v, \qquad L_b(u_i, v) \quad (1 \le i \le 3,\ b \in \mathcal{C} \setminus \{\infty\}) \tag{2.12}$$

equals $\varepsilon k r$, where $\varepsilon \in \{1, -1\}$, $k \in \mathcal{K}$, and $r > Y_0$ is a positive rational prime. Every such r splits in F, is a prime of good reduction for both E and E^σ, lies outside S_E, and satisfies $r > K + 1$.

(2) Witnesses at each place. *For each occurrence of a prime r in this list and each of the two places v_r of F above it, there are nonzero integers h, j_0, chosen independently of $a \in \mathbb{Z}$, such that every $a \in \mathbb{Z}$ satisfies conditions (i) and (ii) of Lemma 2.3 in the ordinary setting: take $R = \mathbb{Z}$, its identity embedding in the ordinary structure, $T_a = aP$, $q = r$, and $v_* = v_r$, normalized by $v_r(r) = 1$. In this specialization the tildes disappear. The choices of h, j_0 may differ at different occurrences and at the two places.*

The set $\mathcal{K}$ is independent of K, A, Y_0.

The two places in part (2) are the two residue branches above a split prime. At either one, conditions (i) and (ii) use the two conjugate curves at that same place. If $\mathfrak{p}$ is the prime ideal of $\mathbb{Z}[\sqrt{2}]$ corresponding to v_r, its valuation ring is $(\mathbb{Z}[\sqrt{2}])_{\mathfrak{p}}$, with uniformizer r. Thus the two valuation inequalities in part (2) are exactly membership in $r^K(\mathbb{Z}[\sqrt{2}])_{\mathfrak{p}}$; for a finite indicated point, reduction to O is equivalent to $1/x \in r(\mathbb{Z}[\sqrt{2}])_{\mathfrak{p}}$. We now encode these ordinary witnesses by formulas that also make sense for a general ring.

2.3 A recursive theory with an ordinary integer model

Let $\mathcal{T}$ be a theory in a countable recursive expansion of the language of rings. Its ring R is required to be an integral domain of characteristic zero with $\sqrt{2} \notin \mathrm{Frac}(R)$. Add the six unary ring-valued functions above. Interpret T_0 as O and, for $a \neq 0$, require the two fractions in (2.3) to define the coordinates of $T_a \in E(F_R)$, with both denominators nonzero. Impose the point conditions (2.2).

Use the five points $\mathcal{C}$, primitive contact forms L_b, and finite multiplier set $\mathcal{K}$ of the preceding inputs. For each ordinary integer $K \geq 1$, add the following representation axiom. For every $A \in R$ there exist $u_1, u_2, v \in R$, with $v \neq 0$, such that, on setting

$$u_3 = Av - u_1 - u_2,$$

each member of

$$v, \qquad L_b(u_i, v) \quad (1 \leq i \leq 3,\ b \in \mathcal{C} \setminus \{\infty\}) \tag{2.13}$$

equals $\varepsilon k r$, where $\varepsilon \in \{1, -1\}$, $k \in \mathcal{K}$, and $r \in R$ has the following properties. The choices of r and its auxiliary data are made separately for the different members of the list.

First require $r \neq 0$, rR proper, and R/rR a field. Require $B_r \in R$ satisfying

$$B_r^2 \equiv 2 \pmod{rR}, \qquad 2B_r \text{ invertible in } R/rR.$$

These conditions define two maximal ideals of R_F:

$$\mathfrak{p}_\pm = \ker(R_F \longrightarrow R/rR,\ \sqrt{2} \longmapsto \pm B_r). \tag{2.14}$$

They are the two residue branches to be tested. In the integer model they correspond to the two places of F above a supplied split prime r. In a general model they are defined by the quotient field, independently of an ambient valuation.

For each $\mathfrak{p} \in \{\mathfrak{p}_+, \mathfrak{p}_-\}$, require nonzero $h, j_0 \in R$, chosen at that branch and independently of a, such that the following hold for every $a \in R$. A point is called *formal for this localization* if it is O, or if $1/x \in r(R_F)_\mathfrak{p}$.

(i) The points T_h, T_{ha} are formal for this localization, $\ell_K(z(T_h)) \neq 0$, and

$$\frac{\ell_K(z(T_{ha}))}{\ell_K(z(T_h))} - a \in r^K (R_F)_\mathfrak{p}. \tag{2.15}$$

(ii) The points $T^\sigma_{j_0}, T^\sigma_{(h-4)a}$ are formal for the conjugate curve at the same localization, and

$$\frac{z(T^\sigma_{(h-4)a})}{z(T^\sigma_{j_0})} \in r^K (R_F)_\mathfrak{p}. \tag{2.16}$$

Here $z(O) = 0$, and the denominator in (2.16) is nonzero by $j_0 \neq 0$ and (2.2). The pairs h, j_0 may differ between the two branches. The scheme contains every ordinary K; the termination proof will later select a single precision after f is fixed.

All these conditions have recursive first-order expressions in R. The absence of $\sqrt{2}$ from $\mathrm{Frac}(R)$ is expressed by

$$\forall a, b \in R \quad b \neq 0 \Longrightarrow a^2 \neq 2b^2.$$

Elements of R_F are pairs of ring elements, and elements of F_R are fractions $(A + B\sqrt{2})/D$, with $A, B, D \in R$ and $D \neq 0$. Equations with fixed coefficients in F can therefore be separated into two coordinates and cleared of denominators. The elliptic group law, including the identity and exceptional addition cases, is given by finitely many field equations and

inequations. Membership in $mE(F_R)$ is expressed by existential coordinates for a point whose m-multiple is the prescribed point. These conditions are ring formulas under the same interpretation.

Saying that R/rR is a field means that $1 \notin rR$ and every element outside rR has an inverse modulo rR. For $D' \in F_R$, the assertion

$$D' \in r^i (R_F)_{\mathfrak{p}_\pm} \tag{2.17}$$

means that $D' = r^i X/Y$ for some $X, Y \in R_F$ such that the evaluation of Y under $\sqrt{2} \mapsto \pm B_r$ is nonzero modulo rR, hence a unit in that field. These representations and inverse conditions are expressed by ring quantifiers. The exponent i is fixed in each axiom; neither variable exponentiation nor a valuation function is part of $\mathcal{T}$.

Lemma 2.7. *The theory $\mathcal{T}$ is recursive and has a model whose ring is $\mathbb{Z}$. In this model $T_a = aP$ for every integer a.*

Proof. The ring and elliptic axioms hold for $\mathbb{Z}$ with $T_a = aP$. For nonzero a, the coordinates lie in F and admit integer numerator and denominator representatives; choose arbitrary integer values for the six coordinate functions at $a = 0$. Nonzero multiples of the nontorsion point P avoid O and the two-torsion points, so their x- and y-coordinates are nonzero.

Fix $K \geq 1$ and $A \in \mathbb{Z}$. Corollary 2.6 supplies (2.13) with the required fixed multipliers and positive split primes. For each such r, the quotient $\mathbb{Z}/r\mathbb{Z}$ is a field. Splitness supplies a root B_r of 2 modulo r, and $2B_r$ is invertible because r is odd. The ideals (2.14) correspond to the two places above r. At each one, the corollary supplies nonzero h, j_0, fixed for all integers a, with the required formal points and both valuation comparisons. Its localization interpretation gives exactly (2.15) and (2.16). Thus every representation axiom holds.

The coefficients of ℓ_K are computable from the fixed equation, and the interpretations and axiom schemes above are effective. The finite sets, fixed integers, rational coefficients, and coordinates of P are finitely many exact constants that can be hardcoded in the theory. Their existence suffices for the existence of the decision procedure constructed below; no effective search for those fixed choices is required. No list of all true arithmetic sentences is used. □

2.4 Finite rational tests and compactness

The finite tests will also impose a positive existential condition on every ring value. This input holds on all integers and controls the parity of possible poles; Section 3 constructs it.

Proposition 2.8. *There are a positive existential formula $\Phi(b)$ in the language of rings, with fixed rational coefficients, and a finite set S_1 of rational primes such that:*

1. *$\mathbb{Q} \models \Phi(b)$ for every $b \in \mathbb{Z}$;*
2. *if $b \in \mathbb{Q}^\times$ and $\mathbb{Q} \models \Phi(b)$, then, for every $q \notin S_1$, the inequality $v_q(b) < 0$ implies $v_q(b) \in 2\mathbb{Z}$.*

The second property does not rule out even pole orders. Its later use rests on a simple product observation. If $q \notin S_1$ and nonzero $a, c \in \mathbb{Q}$ satisfy $\Phi(a)$ and $\Phi(ac)$, with both $v_q(a)$ and $v_q(ac)$ negative, then

$$v_q(c) = v_q(ac) - v_q(a) \in 2\mathbb{Z}.$$

Section 2.5 applies this observation when Φ holds on every element of the ring supplied by the finite tests.

An integer zero can therefore supply rational witnesses for all the conditions imposed below. In the converse direction, pole parity will connect a contact prime in the ambient integer structure to one of the tested localizations of the embedded ring.

Assume that an algorithm decides rational solvability of integral polynomials. Given $f \in \mathbb{Z}[X_1, \ldots, X_n]$, the constant case can be settled immediately, so assume f is nonconstant. Add constants $\mathbf{a} = (a_1, \ldots, a_n)$ and the equation $f(\mathbf{a}) = 0$ to $\mathcal{T}$, and Skolemize the resulting theory. This is effective for a recursive first-order theory: after prenex conversion, each existential quantifier is replaced by a function of the preceding universal variables. All resulting function symbols are ring-valued. The Skolemized axioms are universal sentences with quantifier-free matrices.

A ground term is a term with no variables, built from the named constants using the ring operations and added functions in this language. A compatible assignment of values to all ground terms will interpret the added functions on the ring of term values. Enumerate these terms, associate a rational variable to each, and enumerate the following constraints:

(a) the constants $0, 1$ and the ring operations have their usual values;

(b) equal input tuples have equal outputs for each function symbol;

(c) every ground instance of every universal axiom holds;

(d) the value of every ground term satisfies the positive existential formula Φ from Proposition 2.8.

The witnesses for Φ in (d) are auxiliary rational variables. They are not values of additional functions of the ring and do not create new ground terms subject to (d). The same distinction applies to the inverse variables introduced only when translating a finite test to polynomial equations; they are not added to the term-generated ring.

Every finite initial part of this list is decidable by the assumed rational-solvability algorithm. Indeed, after introducing the witnesses for Φ, the finite collection is an existential Boolean combination of polynomial equalities and inequations. An inequation $g \neq 0$ becomes $gy = 1$ with a new variable; a finite Boolean formula can be put into disjunctive normal form; and a conjunction $g_1 = \cdots = g_s = 0$ is equivalent over $\mathbb{Q}$ to $\sum_i g_i^2 = 0$. Clearing fixed rational denominators gives finitely many integral polynomial queries.

If any test fails, report that f has no integer zero. This answer is sound: an integer zero extends the model of Lemma 2.7 and admits Skolem functions; all its ground-term values are integers and hence satisfy Φ.

Lemma 2.9. *If every finite rational test succeeds, there exist an ambient elementary extension of the many-sorted structure specified in §2.1, a model R of $\mathcal{T}$, a tuple $\mathbf{a} \in R^n$ with $f(\mathbf{a}) = 0$, and an embedding $R \hookrightarrow {}^*\mathbb{Q}$ into its rational sort such that every element of the image satisfies Φ.*

Proof. Take the ordinary many-sorted structure specified in §2.1, with all the indicated functions and relations. Expand this ambient structure by constants naming all ordinary elements in their respective sorts. Its elementary diagram consists of all first-order sentences in its language that are true in the expanded structure. Use this diagram and add constants assigning a value to every ground term of the Skolemized ring language, the ring and functionality constraints, all ground universal instances, and Φ for each term value. Every finite subset has a model: its finitely many constraints are among a successful rational test, and they can be interpreted in the ordinary structure. Finitely many sentences of its elementary diagram are already true there. The compactness theorem therefore supplies a simultaneous assignment in an elementary extension; see [36].

Let R be the set of values of the ground terms. Constraint (a) makes it a subring of ${}^*\mathbb{Q}$. Interpret each added function by applying its symbol to representative ground terms. Constraint (b) makes the result independent of representatives. Every element of R is represented by a ground term, so (c) implies every universal Skolemized axiom in R. Its reduct is the required model of $\mathcal{T}$, and (d) gives the pole-parity condition on every element. The interpretation of fraction fields extends the embedding to $F_R \hookrightarrow {}^*F$.

The ambient diagram is used only in this existence argument. It is not required to be recursive and is never queried by the finite rational tests. □

We now analyze any model and embedding supplied by Lemma 2.9. All starred objects refer to its one ambient extension. The point conditions give the additive injection η of (2.4)–(2.5).

2.5 Odd valuations select a tested branch

The height radical supplies primes of the ambient integer sort, whereas the tests are imposed at ideals above an element $r \in R$. We now connect these two kinds of data. Fix an ordinary precision $K \geq 1$. Enlarge the finite exceptional set S in the height estimate to a finite set S_K containing: the exceptional set for Φ, the supports of all multipliers in $\mathcal{K}$, the primes needed to keep the contact forms distinct with their required unit determinants, the bad or ramified primes of the elliptic data and F, and the finitely many primes excluded by Proposition 2.2 at precision K.

Write $M_K(s) = M_{S_K}(s)$ for the radical with these larger exceptions. Corollary 2.5 keeps the exponent unchanged: if $M(s) = M_S(s)$ is the original radical, then

$$M(s) \leq \left(\prod_{q \in S_K \setminus S} q \right) M_K(s), \qquad h(s) \leq H_K M_K(s)^c$$

with an ordinary constant H_K. Enlarge c to a positive integer once and for all if necessary.

Choose the representation axiom at precision K for any $A \in R$, and put $\widetilde{s}_i = (\widetilde{u}_i : \widetilde{v})$. All its contact forms are nonzero, so $\widetilde{s}_i \notin \mathcal{C}$. Consider any ambient prime q counted in some $M_K(\widetilde{s}_i)$. Here and below a radical, prime, or valuation on starred sorts is the transfer of its ordinary definition.

Let $t = \min\{v_q(\widetilde{u}_i), v_q(\widetilde{v})\}$. At $q \notin S_K$, the minimum of the valuations of the five contact forms equals t: all their coefficients are integral and a pair has unit determinant. Since some contact depth is odd, these five valuations cannot all have the same parity. In particular one of them is odd. For its expression $\varepsilon k r$ in (2.13), k is a q-unit, and hence

$$v_q(\widetilde{r}) \text{ is odd.} \tag{2.18}$$

The odd valuation links this ring element r to the ambient prime q, but it does not make the tested localization a subring of the ambient valuation ring: R may have denominators at q. We will quotient the value group by the convex subgroup generated by the negative valuations of nonzero ring elements. Pole parity ensures that the odd value of r survives positively in this quotient. The center of the resulting valuation will then select one of the two tested branches.

Lemma 2.10. *For the prime q and element r in (2.18), there is a branch $\mathfrak{p} \in \{\mathfrak{p}_+, \mathfrak{p}_-\}$ of that representation axiom and an ambient place v_* of ${}^*F$ over q such that, for every ordinary integer $l \geq 1$ and every $D' \in F_R$,*

$$D' \in r^l (R_F)_{\mathfrak{p}} \quad \Longrightarrow \quad v_*(\widetilde{D}') \geq l. \tag{2.19}$$

Every point formal for this tested localization actually reduces to O for v_*.

Proof. Let J be the convex subgroup of ${}^*\mathbb{Z}$ generated by the negative values of v_q on $\widetilde{R}\setminus\{0\}$. Thus $\gamma \in J$ if $|\gamma|$ is bounded by a finite sum of absolute values of such negative values. This definition is external; we will use it only to construct a valuation on the embedded ring, not as part of an algorithm or a transferred formula.

Every valuation of a nonzero ring element that lies in J is even. If $J = 0$, this is immediate. Otherwise, given $a \neq 0$ with $v_q(\widetilde{a}) \in J$, finite products of ring elements having negative value give $b \in R\setminus\{0\}$ with $v_q(\widetilde{b}) < -|v_q(\widetilde{a})|$. Both $\widetilde{b}$ and $\widetilde{a}\widetilde{b}$ have negative valuation. The soundness of Φ, transferred from Proposition 2.8, says that these two valuations are even. Their difference $v_q(\widetilde{a})$ is therefore even as well.

By (2.18), $v_q(\widetilde{r})$ is outside J; it is positive, because every negative ring value lies in J. Extend v_q to an ambient place v_* on ${}^*F$. Outside the fixed ramified primes, its value group is still ${}^*\mathbb{Z}$ with the same normalization, by transfer of the corresponding ordinary place facts. Coarsen it to the ordered quotient ${}^*\mathbb{Z}/J$, obtaining w_*. This coarsened valuation is nonnegative on $\widetilde{R}$: all its negative original values become zero. It is also nonnegative on $\widetilde{R}_F$, since $\sqrt{2}$ is integral.

The pullback of the center on the embedded R_F is

$$\mathfrak{p} = \{x \in R_F : w_*(\widetilde{x}) > 0\}.$$

It is a proper prime ideal, and it contains r because $w_*(\widetilde{r}) > 0$. On the other hand,

$$R_F/rR_F \simeq (R/rR)[X]/(X^2-2) \simeq (R/rR)\times(R/rR),$$

where the last isomorphism uses $B_r^2 = 2$ and the invertibility of $2B_r$. A prime of this product is one of its two maximal ideals. Consequently $\mathfrak{p}$ is exactly one of the tested branch ideals $\mathfrak{p}_\pm$.

If $D' = r^l X/Y$ as in (2.17), then $Y \notin \mathfrak{p}$, so its image has w_*-value zero. Thus

$$w_*(\widetilde{D}') \geq l w_*(\widetilde{r}).$$

For $D' = 0$, (2.19) is automatic. Otherwise, if $J = 0$, the displayed inequality gives $v_*(\widetilde{D}') \geq l v_q(\widetilde{r}) \geq l$. If $J \neq 0$, its convexity implies that it contains every ordinary integer. A positive value in the ordered quotient is represented by an ambient integer larger than every element of J; the same inequality again gives $v_*(\widetilde{D}') \geq l$.

Finally, a tested finite formal point has $1/x \in r(R_F)_{\mathfrak{p}}$, so its actual x-valuation is negative. At the good integral model this is exactly the condition for reduction to O. The identity point has that reduction as well. □

2.6 Index congruences at every contact prime

Proposition 2.11. *Fix an ordinary $K \geq 1$ and use the exceptional set S_K above. Choose a representation supplied at precision K for an arbitrary $A \in R$, and put $\widetilde{s}_i = (\widetilde{u}_i : \widetilde{v})$. For every prime q of ${}^*\mathbb{Z}$ counted by any $M_K(\widetilde{s}_i)$, and every $a \in R$, the element $\widetilde{a}$ is integral at q and*

$$v_q(\widetilde{a} - \eta(a)) \geq K. \tag{2.20}$$

Proof. The odd-contact argument selects a factor r with odd q-valuation. Lemma 2.10 supplies one of its tested branches and a place v_* above q. At this branch the theory supplies nonzero h, j_0, fixed for all $a \in R$. Every tested formal point is an actual formal point at v_*, for its own curve. The same lemma sends the memberships (2.15) and (2.16) to the two valuation comparisons of Lemma 2.3, each at precision K and at this same place. The primary logarithm denominator is nonzero by its axiom and the embedding; the conjugate parameter denominator is nonzero by $j_0 \neq 0$ and the point conditions. Finally $q \notin S_K$ ensures $q \notin S_E$ and $q > K+1$. All hypotheses of Lemma 2.3 therefore hold, and it gives the assertion. □

2.7 The height contradiction

Proposition 2.12. *Let $f \in \mathbb{Z}[X_1, \dots, X_n]$ be nonconstant. If every finite system of rational constraints constructed in §2.4 for f is solvable, then f has a zero in $\mathbb{Z}^n$.*

Proof. Choose the model supplied by Lemma 2.9, with its root tuple $\mathbf{a}$ and the additive map η of (2.4)–(2.5). Set

$$\delta = f(\eta(a_1), \dots, \eta(a_n)) \in {}^*\mathbb{Z}.$$

If $\delta = 0$, the ambient integer sort satisfies $\exists \mathbf{x} \in \mathbb{Z}^n \ f(\mathbf{x}) = 0$; by elementarity the ordinary integers satisfy that sentence. It remains to consider $\delta \neq 0$.

Write

$$B = \max(1, |\eta(a_1)|, \dots, |\eta(a_n)|), \qquad D = \max(1, \deg f). \tag{2.21}$$

For an ordinary constant $C_f > 0$, transfer of the elementary polynomial bound gives

$$|\delta| \leq C_f B^D. \tag{2.22}$$

Choose an ordinary integer $K \geq cD$, where c is the fixed positive integer exponent of the height estimate. All exceptions and constants below refer to this fixed K.

Figure 1 summarizes how the chosen representations for all $A \in R$ use the same δ.

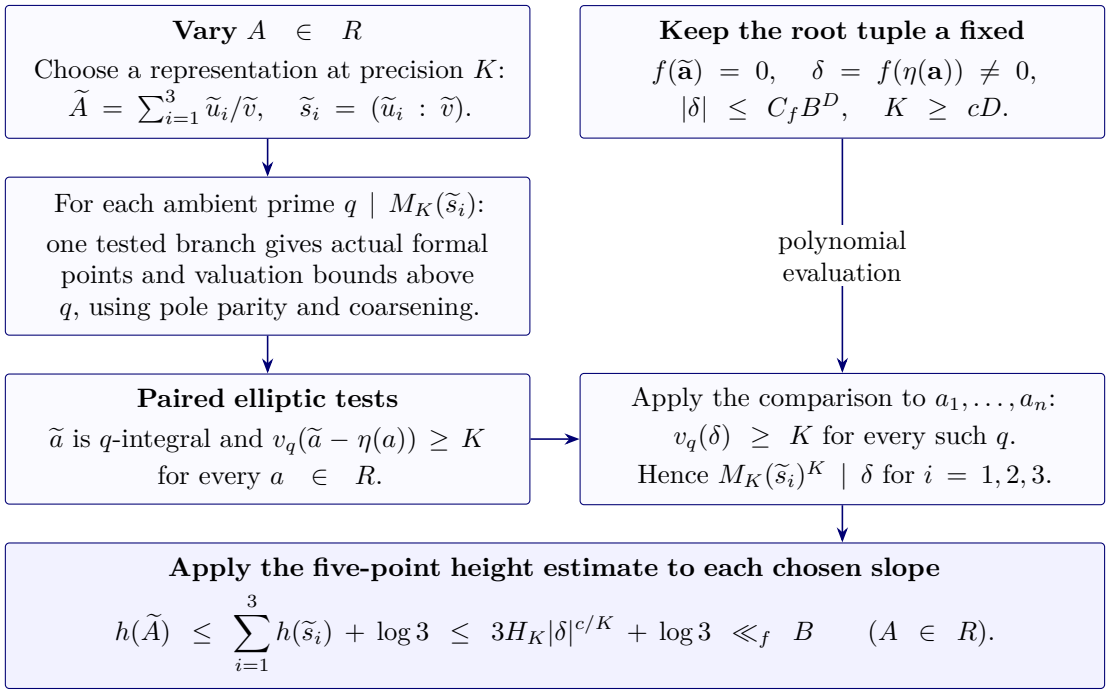


Figure 1: The use of one nonzero δ for every $A \in R$. All slopes shown come from the representation chosen at the fixed ordinary precision K, and M_K uses the fixed exceptional set for K. The comparison holds for every $a \in R$, so it applies to the fixed root tuple. The implied constant in the final bound is ordinary and independent of A and $\mathbf{a}$.

For any $A \in R$, take its representation at precision K and any prime q counted in any $M_K(\widetilde{s}_i)$. Proposition 2.11 applies to each a_j. Both the embedded coordinates and their integer indices are q-integral. Since $f(\widetilde{\mathbf{a}}) = 0$ in the ambient field and the coefficients of f are integers, their coordinatewise congruences give

$$v_q(\delta) \geq K. \tag{2.23}$$

For each fixed slope, (2.23) now holds for every ambient prime counted by its radical. This is the internal universal premise in the transfer of the ordinary divisibility fact: if every prime in a squarefree product divides an integer to order at least K, then the Kth power of that product divides the integer. Using the transferred finite-product function gives

$$M_K(\widetilde{s}_i)^K \mid \delta. \tag{2.24}$$

The coarsenings may have been chosen externally one prime at a time; only the universal assertion just established enters this transferred implication. Although A and its representation were arbitrary, the integer δ is the same for all of them.

As $\delta \neq 0$, (2.24) and (2.22) give

$$M_K(\widetilde{s}_i)^c \leq |\delta|^{c/K} \leq C_f^{c/K} B^{cD/K} \ll_f B.$$

The height estimate therefore implies

$$h(\widetilde{s}_i) \ll_f B. \tag{2.25}$$

Since $v \neq 0$ and $u_1 + u_2 + u_3 = Av$, the ordinary logarithmic height inequality for a sum of three numbers, transferred to the ambient structure, gives

$$h(\widetilde{A}) = h\left(\frac{\widetilde{u}_1}{\widetilde{v}} + \frac{\widetilde{u}_2}{\widetilde{v}} + \frac{\widetilde{u}_3}{\widetilde{v}}\right) \leq \sum_{i=1}^{3} h(\widetilde{s}_i) + \log 3 \ll_f B \qquad (A \in R). \tag{2.26}$$

The implicit constant is ordinary and independent of A and of the root tuple. The dependence on f comes only from its degree and coefficients and the resulting fixed choice of K.

For each nonzero a_j, formula (2.3) expresses $x(T_{a_j})$ using the three ring elements $x_0(a_j), x_1(a_j), x_2(a_j)$, with $x_2(a_j) \neq 0$. Applying (2.26) to these three elements and using the logarithmic height inequalities over the fixed number field F gives

$$h(x(\eta(a_j)P)) \ll_f B. \tag{2.27}$$

Canonical height on the fixed elliptic curve gives, for ordinary integer $n \neq 0$, and hence for ambient integer indices by transfer,

$$h(x(nP)) = 2\widehat{h}(P)n^2 + O(1), \qquad \widehat{h}(P) > 0; \tag{2.28}$$

see [67]. The error term is uniformly bounded for the fixed curve and point. If $B > 1$, choose a maximizing nonzero index in (2.21). Equations (2.27)–(2.28) give

$$2\widehat{h}(P)B^2 \leq CB + C'$$

for ordinary constants, so B is bounded by an ordinary constant. This is also true if $B = 1$.

Every element of ${}^*\mathbb{Z}$ in an ordinary finite interval is an ordinary integer, by transfer of the finite description of that interval. Thus $\eta(a_j) = n_j \in \mathbb{Z}$ for each j. By additivity and $\eta(1) = 1$, $\eta(n_j) = n_j$; injectivity then gives $a_j = n_j$ in R. The relation $f(\mathbf{a}) = 0$ is consequently an ordinary integer solution. In the case currently under consideration this also contradicts $\delta \neq 0$. □

Proof of Theorem 1.1. Constant polynomials can be decided directly. For nonconstant f, if every finite rational test succeeds, Proposition 2.12 gives an integer zero. Therefore, when f has no integer zero, some finite test must fail. Run the finite-test search and enumeration of integer tuples in parallel, allocating successive finite steps to each. The enumeration terminates when a zero exists; the test search terminates when none exists. Their answers are sound. Rational decidability would therefore give integer decidability, contradicting the Davis–Putnam–Robinson–Matiyasevich theorem [14, 37]. □

3 An existential condition excluding odd poles

We prove Proposition 2.8, used in the finite rational tests of Section 2. It requires a positive existential condition that holds on all integers and restricts the denominators of any rational number satisfying it. Square classes below mean classes in the multiplicative group modulo squares.

Earlier existential formulas make related valuation information accessible on restricted sets of primes. For a fixed global field K of characteristic different from 2 and a fixed quadratic extension L/K, Demeyer and Van Geel give an existential formula which, for nonzero $x, y \in K$, is equivalent to

$$v_{\mathfrak{p}}(x) \text{ odd} \quad \Longrightarrow \quad v_{\mathfrak{p}}(x/y^2) > 0 \qquad \text{for every nonarchimedean prime } \mathfrak{p} \text{ inert in } L/K.$$

They also permit a fixed finite union of such inert-prime sets [18, Theorem 14 and Corollary 15 in the author version]. Cornelissen and Zahidi use related odd-valuation predicates in their elliptic-divisibility approach [9, Sections 3.3 and 3.11]. The existential definitions of Demeyer and Van Geel control the chosen inert primes. Proposition 2.8 instead supplies the stated one-sided restriction outside a fixed finite exceptional set.

We first define a formula from finite multiplier lists and prove the pole restriction. We then choose the lists once and for all and prove integer completeness by realizing the same square class on two elliptic curves.

3.1 The formula and the pole restriction

For now, let $\mathcal{M} \subset \mathbb{Q}^\times$ be a finite nonempty list, and for each $m \in \mathcal{M}$ let $\mathcal{D}_m \subset \mathbb{Q}^\times$ be a finite nonempty list. The construction below associates a formula Φ to these lists. Let S_1 be a finite set containing 2, 3, and the supports of every member of all the lists. We will make one fixed choice of the lists in the next subsection.

For $l \in \mathbb{Q}^\times$, write

$$E_l : \quad y^2 = x(x-l)(x+3l). \tag{3.1}$$

The three roots are distinct, so this is an elliptic curve with its point at infinity O as origin. Define $\Phi(b)$ to hold if $b = 0$, or if $b \neq 0$ and, for some $m \in \mathcal{M}$, $m_d \in \mathcal{D}_m$, and $t, e \in \mathbb{Q}$, the following conditions hold:

$$d = m_d b, \qquad b' = mb, \qquad U = b\frac{1-dt^2}{1+dt^2}, \qquad H' = b'(b'-U^2), \tag{3.2}$$

$$1 + dt^2 \neq 0, \qquad eH' \neq 0,$$

and each of E_e and $E_{eH'}$ has a rational point whose x-coordinate belongs to $b'\mathbb{Q}^{\times 2}$.

For each fixed choice of the lists, this is a positive existential formula. The choices of multipliers are finite disjunctions. A nonzero condition $a \neq 0$ is the existential equation $az = 1$; the condition on x is expressed by $x = b'z^2$ with $z \neq 0$. Introduce U as a variable and multiply its defining equality by $1 + dt^2$. All remaining conditions are polynomial equalities with fixed rational coefficients, which can be cleared to integer coefficients.

Lemma 3.1. *For any finite lists as above, if $b \in \mathbb{Q}^\times$ satisfies $\Phi(b)$ and $q \notin S_1$, then $v_q(b) < 0$ implies $v_q(b) \in 2\mathbb{Z}$.*

Proof. Suppose that $q \notin S_1$, that witnesses for $\Phi(b)$ have been chosen, and that $k = v_q(b) < 0$ is odd. The multipliers are q-adic units, so $v_q(d) = v_q(b') = k$. If $t = 0$, then $U = b$. If $t \neq 0$, the valuation $v_q(dt^2) = k + 2v_q(t)$ is odd and hence nonzero. When this valuation is positive,

both $1 - dt^2$ and $1 + dt^2$ are units. When it is negative, both have valuation $v_q(dt^2)$. In either case $v_q(U) = k$. Since $2k < k$,

$$v_q(b' - U^2) = 2k, \qquad v_q(H') = 3k.$$

Consequently $v_q(e)$ and $v_q(eH')$ have opposite parities. Let l be the one of these two parameters whose valuation is even, say $v_q(l) = 2a$. Replacing

$$l = q^{2a}L, \qquad x = q^{2a}X, \qquad y = q^{3a}Y$$

gives $Y^2 = X(X - L)(X + 3L)$ with $L \in \mathbb{Z}_q^\times$ and does not change the square class of x.

Every nonzero X on this unit-parameter curve has even valuation. Indeed, if $v_q(X) < 0$, all three factors have valuation $v_q(X)$, so the right side has valuation $3v_q(X)$. If $v_q(X) > 0$, the other two factors are units because $q \neq 3$, so its valuation is $v_q(X)$. In both cases this valuation must be even. The case $v_q(X) = 0$ already has the desired parity. This contradicts $x \in b'\mathbb{Q}_q^{\times 2}$, since $v_q(b') = k$ is odd. □

3.2 Preparing integer witnesses

We now choose the multiplier lists so that the formula also holds on every integer. These lists will be fixed independently of the input b. Choose a finite list $\mathcal{M} \subset \mathbb{Q}^\times$ with the following property: for every $b \in \mathbb{Q}^\times$, some $m \in \mathcal{M}$ makes

$$b' = mb > 0, \qquad b' \in \mathbb{Q}_2^{\times 2}, \qquad b' \in \text{the unit nonsquare class of } \mathbb{Q}_3^\times. \tag{3.3}$$

Here the last condition means that b' has even 3-adic valuation and nonsquare unit part. Such a list exists because $\mathbb{R}^\times/\mathbb{R}^{\times 2}$, $\mathbb{Q}_2^\times/\mathbb{Q}_2^{\times 2}$, and $\mathbb{Q}_3^\times/\mathbb{Q}_3^{\times 2}$ are finite, their classes are open, and weak approximation supplies a rational representative for each specified finite tuple of classes. For each $m \in \mathcal{M}$, set

$$S(m) = \{2, 3\} \cup \mathrm{Supp}(m).$$

The same argument gives a finite list $\mathcal{D}_m \subset \mathbb{Q}^\times$ such that, for any $b \neq 0$, some $m_d \in \mathcal{D}_m$ makes

$$d = m_d b > 0, \qquad d \in \mathbb{Q}_q^{\times 2} \quad (q \in S(m)). \tag{3.4}$$

Use these fixed lists in Φ, and fix S_1 containing 2, 3, and all their supports as above. The normalizations (3.3)–(3.4) are used only to select entries m, m_d for an integer input; they are not additional conjuncts of Φ. Lemma 3.1 already proves the pole restriction for this choice.

Fix $b \in \mathbb{Z} \setminus \{0\}$ and choose m, m_d so that (3.3) and (3.4) hold. We choose t to serve two purposes. The value U should be small at the places where d was made a local square, while at primes outside $S(m)$ where $v_q(b')$ is odd it should retain the valuation of b. These properties will make H' a local square at every prime where $v_q(b')$ is odd, and b' a local square at every prime where $v_q(H')$ is odd.

At the real place and at every $q \in S(m)$, choose a local solution of $dt^2 = 1$. As t tends to this solution, the rational function U tends to zero: its denominator tends to the nonzero number 2. Thus, in sufficiently small neighborhoods, $H' > 0$ at the real place, and

$$H' = b'^2\left(1 - \frac{U^2}{b'}\right) \in \mathbb{Q}_q^{\times 2} \quad (q \in S(m)),$$

by openness of the local square subgroup. At every $q \in \mathrm{Supp}(m_d) \setminus S(m)$, require t to be so small that $v_q(dt^2) > 0$. Weak approximation satisfies all these conditions simultaneously with a rational t. In particular $1 + dt^2 \neq 0$ and $H' > 0$.

Let β, θ be the unique positive squarefree integers representing the square classes of b', H', and let

$$B = \mathrm{Supp}(\beta), \qquad J = \mathrm{Supp}(\theta).$$

Lemma 3.2. *The sets B, J are disjoint and avoid $2, 3$. Moreover,*

$$\theta \in \mathbb{Q}_q^{\times 2} \quad (q \in B), \qquad \beta \in \mathbb{Q}_q^{\times 2} \quad (q \in J). \tag{3.5}$$

Both β and θ are squares in $\mathbb{Q}_2^\times$; at 3, θ is a square and β is a unit nonsquare. In particular $B \neq \varnothing$.

Proof. The assertions at $2, 3$ follow from the choices already made and are unchanged on replacing numbers by square-class representatives. At a prime $q \in B \cap S(m)$, H' is a local square by construction. Now suppose $q \in B \setminus S(m)$. Since m is a unit at q and b is an integer, $v_q(b) = v_q(b')$ is a positive odd integer. If $q \notin \mathrm{Supp}(m_d)$ and $t \neq 0$, then $v_q(dt^2) = v_q(b) + 2v_q(t)$ is odd and nonzero, and the quotient $(1 - dt^2)/(1 + dt^2)$ is a unit, just as in Lemma 3.1. For $t = 0$ this quotient is 1. If instead $q \in \mathrm{Supp}(m_d)$, our additional smallness condition on t also makes it a unit. Hence $v_q(U) = v_q(b)$, and

$$v_q(U^2/b') = v_q(b) > 0.$$

As q is odd, $1 - U^2/b'$ is a local square. The displayed formula for H' proves the first assertion of (3.5), and also proves $B \cap J = \varnothing$.

For the converse assertion, let $q \in J$. Such a prime is outside $S(m)$, since H' is square there, and $v_q(b')$ is even because $q \notin B$. If $U = 0$, then $H' = b'^2$ and there is no such prime, so suppose $U \neq 0$. Put

$$\alpha = v_q(b'), \qquad \gamma = 2v_q(U).$$

Both are even. If $\alpha \neq \gamma$, then $v_q(b' - U^2) = \min(\alpha, \gamma)$ is even, contrary to the oddness of $v_q(H')$. Therefore $\alpha = \gamma$. Since $v_q(H') = \alpha + v_q(b' - U^2)$ is odd, cancellation is strict: $v_q(b' - U^2) > \alpha$. Thus $b'/U^2 \in 1 + q\mathbb{Z}_q$, a square in $\mathbb{Q}_q^\times$, proving that β is square there. This argument also covers possible denominators in t or U. Finally β cannot be 1, since its class at 3 is nonsquare. □

Write $H' = \theta c^2$ with $c \in \mathbb{Q}^\times$. For every $n \in \mathbb{Q}^\times$ the isomorphism

$$E_{n\theta} \longrightarrow E_{nH'}, \qquad (x, y) \longmapsto (c^2 x, c^3 y) \tag{3.6}$$

preserves x-coordinate square classes. Thus, taking $e = n$ in the formula, it remains to find one n for which both E_n and $E_{n\theta}$ have a rational point with nonzero x-coordinate of class β, the class of b'. We first give a criterion for one parameter l, and then construct a single n for which it applies to both twists.

3.3 A criterion for one twist

We recall the exact part of full two-descent that we use; see [67, Chapter X, Proposition 1.4, p. 315]. For a split separable cubic $y^2 = \prod_{i=1}^3 (x - e_i)$ over a field k of characteristic zero, choosing two roots identifies $H^1(k, E[2])$ with $(k^\times/k^{\times 2})^2$. In the coordinate corresponding to e_i, the Kummer map takes a point to the class of $x - e_i$; at $(e_i, 0)$ the value is replaced by $\prod_{j \neq i}(e_i - e_j)$, and at O it is 1. The Kummer map is an injection of $E(k)/2E(k)$ into this group. The group $\mathrm{Sel}_2(E/\mathbb{Q})$ is the subgroup of rational pairs lying in the Kummer image at every completion. The standard Selmer exact sequence [67, Chapter X, Theorem 4.2, p. 333] is

$$0 \longrightarrow E(\mathbb{Q})/2E(\mathbb{Q}) \longrightarrow \mathrm{Sel}_2(E/\mathbb{Q}) \longrightarrow \mathrm{Sha}(E/\mathbb{Q})[2] \longrightarrow 0. \tag{3.7}$$

We use the coordinates $(x+3l, x-l)$ for E_l. For a finite rational point T with $x(T) \notin \{0, l, -3l\}$, the curve equation gives

$$[x(T)] = [x(T)+3l]\,[x(T)-l].$$

Consequently, if the class $(\beta, 1)$ is represented by a rational point and is distinct from the classes of all four rational two-torsion points, that point has the required nonzero x-coordinate class β. The next calculation identifies conditions that make $(\beta, 1)$ the one additional Selmer direction beyond rational two-torsion.

For the Selmer graph-matrix viewpoint underlying this calculation, see Monsky's appendix to Heath-Brown [30, Appendix, pp. 365–370]. The matrix and simultaneous rank calculations below adapt that viewpoint to the present curve family.

At the primes dividing l, the local square conditions will be recorded by binary residue symbols. For an odd prime q and a q-adic unit a, write $[a/q] \in \mathbb{F}_2$ for the nonsquare bit: it is 0 for a square residue and 1 for a nonsquare residue. For a finite set Q of distinct primes different from $2, 3$, indexed by i, and a vector $w \in \mathbb{F}_2^Q$, write D_w for the diagonal matrix with diagonal w. Put

$$s_i = [-1/q_i], \qquad \lambda_i = [3/q_i], \qquad (Ku)_i = \sum_{j \neq i} [q_j/q_i](u_j + u_i), \qquad N = K + D_s. \tag{3.8}$$

All matrices and vectors in this calculation are over $\mathbb{F}_2$. Write $\mathbf{1}$ for the all-ones vector and $\mathbf{1}_B$ for the indicator of B in Q. Quadratic reciprocity gives

$$K + K^{\mathsf{t}} = ss^{\mathsf{t}} + D_s. \tag{3.9}$$

In particular, if $s^{\mathsf{t}}\mathbf{1} = 1$, then

$$K\mathbf{1} = K^{\mathsf{t}}\mathbf{1} = 0, \qquad N = K^{\mathsf{t}} + ss^{\mathsf{t}}. \tag{3.10}$$

The family $y^2 = x(x-n)(x+3n)$ also occurs in the study of $\pi/3$- and $2\pi/3$-congruent numbers. Mokrani adapted Monsky matrices to these families [43]. For this family, Wei and Guo give a two-Selmer matrix formulation and write the same quadratic-reciprocity identity for the associated Rédei matrix as (3.9) [72, Section 4 and Theorem 6.2]. The simultaneous kernel conditions for the two twists required here are constructed below.

Lemma 3.3. *Let $l > 0$ be squarefree, prime to 6, with $l \equiv 3 \pmod 4$ and $l \equiv 2 \pmod 3$. Suppose that its support Q contains the nonempty set B of Lemma 3.2 and a prime outside B, and suppose that*

$$\ker K = \langle \mathbf{1}, \mathbf{1}_B \rangle, \qquad \lambda \notin \operatorname{im} K. \tag{3.11}$$

Then

$$\dim_{\mathbb{F}_2} \mathrm{Sel}_2(E_l/\mathbb{Q}) = 3,$$

and $(\beta, 1)$ is a Selmer class outside the subgroup of rational two-torsion classes, in Kummer coordinates $(x+3l, x-l)$.

Proof. In the stated coordinates, the points $(l, 0)$ and $(-3l, 0)$ have classes

$$\tau_1 = (l, 1), \qquad \tau_2 = (3, -l), \tag{3.12}$$

respectively. For example, the replacement coordinate at $(l, 0)$ is $(l+3l)l = 4l^2$, and the replacement at $(-3l, 0)$ is $(-3l)(-4l) = 12l^2$. These two classes are independent globally: l is nonsquare, and the second coordinate of τ_2 is negative.

We first bound the Selmer group using necessary local conditions. Afterward we verify that $(\beta, 1)$ lies in the actual local Kummer image everywhere. For any odd prime q, the group $E_l(\mathbb{Q}_q)/2E_l(\mathbb{Q}_q)$ has order four. To see this, choose an open formal subgroup H on which multiplication by 2 is an isomorphism. The quotient $G = E_l(\mathbb{Q}_q)/H$ is finite, and $E_l(\mathbb{Q}_q)/2E_l(\mathbb{Q}_q) \simeq G/2G$. Moreover $E_l(\mathbb{Q}_q)[2] \simeq G[2]$: injectivity follows from $H[2] = 0$, and a lift of any element of $G[2]$ can be corrected by an element of H, using the surjectivity of multiplication by 2 on H. Finally $|G/2G| = |G[2]| = 4$.

At an odd prime $q \nmid 3l$, the roots are integral and pairwise distinct modulo q. If $v_q(x) < 0$, all three factors have that valuation, which must be even. If x is integral, at most one factor has positive valuation, and that valuation must be even. The replacement classes at the roots are units as well. Hence the local Kummer image is contained in the subgroup of pairs having even coordinate valuations. This unramified subgroup has order four, so the containment is equality.

At a support prime $q_i \mid l$, the valuation pairs of τ_1 and τ_2 are $(1, 0)$ and $(0, 1)$ modulo 2; therefore they generate the entire local image, again by its order four. At 3, l is a unit nonsquare and $-l$ is a square. Thus τ_1 and τ_2 have independent first-coordinate classes and square second coordinates. The local image at 3 is consequently

$$(\mathbb{Q}_3^\times/\mathbb{Q}_3^{\times 2}) \times \{1\}. \tag{3.13}$$

At 2 we need the necessary condition that both Kummer coordinates have even valuation. For a point away from the roots with $v_2(x) < 0$, all three factors have valuation $v_2(x)$, so that valuation is even. If x is integral and even, $x + 3l$ and $x - l$ are units. Suppose x is odd. If $x - l = 2u$ with u odd, then

$$(y/2)^2 = (l + 2u)u(u + 2l) \equiv l \pmod 8.$$

The right side is an odd unit congruent to 3 or 7 modulo 8, which is impossible for a square. In every remaining nonroot case, $v_2(x - l) \geq 2$ and $v_2(x + 3l) \geq 2$. Their difference is $4l$, so their minimum valuation is 2. Since x is a unit, their valuation sum is even, and hence both valuations are even. The root and identity classes also have even valuations, as follows from (3.12). This establishes the claimed necessary condition at 2; no sufficiency is being asserted.

At the real place the first coordinate is always positive. Indeed, for a nonroot real point the allowable x intervals are $(-3l, 0)$ and (l, ∞), and the replacement first coordinates at the roots are positive as well. The sign of the second coordinate therefore gives a homomorphism from the Selmer group to $\{\pm 1\}$, and it is surjective because of τ_2. Its kernel, the classes with both coordinates positive, has index two. Every class in this kernel has a unique representative of the form

$$\left(3^\delta \prod_i q_i^{u_i},\ \prod_i q_i^{v_i}\right), \qquad \delta \in \mathbb{F}_2, \quad u, v \in \mathbb{F}_2^Q. \tag{3.14}$$

Here the good-prime restrictions and the even valuations at 2 exclude every other prime, and (3.13) excludes 3 from the second coordinate.

At q_i, compare (3.14) with the local torsion class

$$\tau_1^{u_i}\tau_2^{v_i} = (l^{u_i}3^{v_i}, (-l)^{v_i}).$$

The coordinate valuations already agree modulo 2. Requiring the resulting unit quotients to be squares gives exactly

$$Ku + D_\lambda v = \delta\lambda, \qquad Nv = 0. \tag{3.15}$$

For example the first unit quotient has nonsquare bit $\sum_{j\neq i}[q_j/q_i](u_j+u_i)+(\delta+v_i)\lambda_i$; the second has bit $(Kv)_i+s_iv_i$. The requirement that the second coordinate be square at 3 adds

$$(\lambda+s)^{\mathsf{t}}v=0, \tag{3.16}$$

using $[q_i/3]=\lambda_i+s_i$.

The residue assumptions imply $s^{\mathsf{t}}\mathbf{1}=1$ and $\lambda^{\mathsf{t}}\mathbf{1}=0$. Equations (3.10) give

$$\ker N=\ker K^{\mathsf{t}}\cap s^{\perp}. \tag{3.17}$$

Indeed, left multiplying $Nv=0$ by $\mathbf{1}^{\mathsf{t}}$ gives $s^{\mathsf{t}}v=0$, and the identity $N=K^{\mathsf{t}}+ss^{\mathsf{t}}$ then gives the claim in both directions. By (3.11), $\ker K^{\mathsf{t}}$ has dimension two. It contains $\mathbf{1}$, on which s^{t} is nonzero, so (3.17) is a line. The functional λ^{t} is nonzero on $\ker K^{\mathsf{t}}$ because $\lambda\notin\operatorname{im}K=(\ker K^{\mathsf{t}})^{\perp}$. Since it vanishes on $\mathbf{1}$, it is nonzero on the complementary line (3.17). Equations (3.16) and (3.17) therefore force $v=0$. The first equation in (3.15) now gives $Ku=\delta\lambda$, so $\delta=0$ and $u\in\ker K=\langle\mathbf{1},\mathbf{1}_B\rangle$. There are at most four sign-normalized classes, and hence at most eight Selmer classes.

For equality, we verify that

$$\gamma=(\beta,1)$$

is an actual local Kummer class everywhere. At 2 and at the real place it is the identity class, since β is a positive local square. At 3 it is τ_1, since β and l are both unit nonsquares. At $q_i\mid l$, take $u=\mathbf{1}_B$ and $v=0$: the comparison with $\tau_1^{u_i}$ has square coordinate quotients precisely because $K\mathbf{1}_B=0$. At every other finite odd prime it is unramified, so lies in the local image already described. This proves $\gamma\in\mathrm{Sel}_2(E_l/\mathbb{Q})$. It is neither the identity nor τ_1: B is nonempty, and l has a prime factor outside B. It cannot equal τ_2 or $\tau_1\tau_2$, whose second coordinates are negative. Thus τ_1,τ_2,γ are three independent Selmer classes, and the upper bound is attained. □

It remains to realize this particular Selmer class by a rational point. The consequence we use from Theorem 1.1 of [48] is the following: an elliptic curve $C/\mathbb{Q}$ with $C(\mathbb{Q})[2]\neq 0$ and $\operatorname{corank}_{\mathbb{Z}_2}\mathrm{Sel}_{2^\infty}(C/\mathbb{Q})\in\{0,1\}$ has finite $\mathrm{Sha}(C/\mathbb{Q})$. The next proof verifies the full Selmer-corank hypothesis before applying it.

Lemma 3.4. *Under the hypotheses of Lemma 3.3, the curve E_l has a rational point with nonzero x-coordinate in $\beta\mathbb{Q}^{\times 2}$.*

Proof. Put $r=\operatorname{rank}E_l(\mathbb{Q})$. Since E_l has full rational two-torsion, the Kummer sequence (3.7) and Lemma 3.3 imply

$$r+\dim_{\mathbb{F}_2}\mathrm{Sha}(E_l/\mathbb{Q})[2]=1. \tag{3.18}$$

This also controls the full 2-primary Selmer group before any finiteness conclusion. The usual exact sequence is

$$0\longrightarrow E_l(\mathbb{Q})\otimes\mathbb{Q}_2/\mathbb{Z}_2\longrightarrow\mathrm{Sel}_{2^\infty}(E_l/\mathbb{Q})\longrightarrow\mathrm{Sha}(E_l/\mathbb{Q})[2^\infty]\longrightarrow 0.$$

These are cofinitely generated $\mathbb{Z}_2$-modules. If the last group has corank t, its structure is $(\mathbb{Q}_2/\mathbb{Z}_2)^t\oplus F$ with F finite, so $t\leq\dim_{\mathbb{F}_2}\mathrm{Sha}(E_l/\mathbb{Q})[2]$. It follows from (3.18) that

$$0\leq\operatorname{corank}_{\mathbb{Z}_2}\mathrm{Sel}_{2^\infty}(E_l/\mathbb{Q})=r+t\leq 1.$$

The curve has nonzero rational two-torsion, so [48, Theorem 1.1] now applies and proves that $\mathrm{Sha}(E_l/\mathbb{Q})$ is finite.

The Cassels–Tate pairing on this finite group is perfect and alternating. Alternation here uses the principal polarization represented by the rational divisor (O); see [6]; see also [55, Section 1 and Corollary 9]. For completeness, an alternating perfect pairing on a finite abelian 2-group forces an even number of cyclic factors. Choose an element x of maximal order 2^a. Perfection supplies y whose pairing with x has exact order 2^a. The subgroup generated by x, y is a nondegenerate copy of $(\mathbb{Z}/2^a\mathbb{Z})^2$, and the whole group is its direct sum with its orthogonal complement. Induction gives a decomposition into such paired cyclic factors. Hence

$$\dim_{\mathbb{F}_2} \mathrm{Sha}(E_l/\mathbb{Q})[2] \quad \text{is even.}$$

This conclusion uses the perfect pairing on the full finite 2-primary group; its restriction to the subgroup killed by 2 need not be perfect. Equation (3.18) therefore gives $\mathrm{Sha}(E_l/\mathbb{Q})[2] = 0$ and $r = 1$.

The Kummer sequence now shows that $\gamma = (\beta, 1)$ is the class of a rational point T. By Lemma 3.3 its class differs from all four two-torsion point classes. Thus T is finite and none of $x(T)$, $x(T) - l$, $x(T) + 3l$ is zero. The curve equation gives, modulo squares,

$$[x(T)] = [x(T) + 3l]\,[x(T) - l] = \beta,$$

as required. □

Remark 3.5. The unrestricted 2-converse [50, Theorem 1.1] can replace the pointwise theorem in the preceding proof. It gives finiteness of the whole $\mathrm{Sha}(E_l/\mathbb{Q})$ from the same full Selmer-corank bound, without any rational-two-torsion hypothesis.

3.4 One parameter for both twists

The criterion is now a condition on the prime support of one twist. We construct a single positive squarefree n for which it holds both on the support of n and on that of $n\theta$. The mutual local-square relations from Lemma 3.2 are what allow the second support to be added without disturbing the first construction.

Lemma 3.6. *For β, θ, B, J as in Lemma 3.2, there is a positive squarefree integer n, prime to 6θ, whose support P contains B and at least one auxiliary prime outside B, such that the following statements hold for both $l = n$ and $l = n\theta$. The number l satisfies $l \equiv 3 \pmod 4$ and $l \equiv 2 \pmod 3$. On its prime support $Q = P$ or $Q = P \cup J$, the matrices (3.8) satisfy (3.11).*

Proof. Since β, θ are odd squares at 2, their products of (-1)-symbols are 1. At 3, reciprocity gives $[q/3] = [3/q] + [-1/q] = \lambda_q + s_q$. Consequently

$$\sum_{i\in B} s_i = \sum_{j\in J} s_j = 0, \qquad \sum_{i\in B} \lambda_i = 1, \qquad \sum_{j\in J} \lambda_j = 0. \tag{3.19}$$

The fixed B–J symbols need not vanish individually. Their row sums, however, satisfy

$$\sum_{j\in J} [q_j/q_i] = 0 \quad (i \in B), \qquad \sum_{i\in B} [q_i/q_j] = 0 \quad (j \in J), \tag{3.20}$$

by (3.5).

Introduce one auxiliary prime q_0 of type $(s_0, \lambda_0) = (1, 1)$, and a set I of $|B| - 1$ auxiliary primes of type $(0, 0)$. These types can be imposed by the residue classes 7 and 1 modulo 12, respectively. We first prescribe all required Legendre symbols abstractly, and realize them by actual primes at the end. Until then, each unchosen symbol is a bit whose reverse

is constrained by quadratic reciprocity; the matrix reciprocity identity therefore holds for every completion of these choices. On each of the sets

$$Q_0 = B \cup I \cup \{q_0\}, \qquad Q_1 = Q_0 \cup J,$$

the total s-sum is 1 and the total λ-sum is 0. The former gives $l \equiv 3 \pmod 4$, while the sum of $s + \lambda$ gives $l \equiv 2 \pmod 3$ for the corresponding prime product.

We first show that, subject to $K\mathbf{1}_B = 0$, both matrix targets follow from one nonsingularity condition. Fix $b_* \in B$ and let a have entries 1 at q_0, b_* and 0 elsewhere. Thus

$$a^{\mathsf{t}}\mathbf{1} = 0, \qquad a^{\mathsf{t}}\mathbf{1}_B = 1.$$

We shall impose $K\mathbf{1}_B = 0$ on both supports. Assuming this condition for the moment, set $K^{\#} = K + \lambda a^{\mathsf{t}}$. The total symbol sums and (3.10) show that $K^{\#}$ kills $\mathbf{1}$ on both sides, and

$$K^{\#}\mathbf{1}_B = \lambda. \tag{3.21}$$

Put

$$T = \mathrm{Id} - \mathbf{1}_B a^{\mathsf{t}}.$$

This is a projection onto $a^{\perp}$, with kernel $\langle \mathbf{1}_B \rangle$, and it fixes $\mathbf{1}$. By (3.21),

$$K = K^{\#}T. \tag{3.22}$$

Suppose that $\ker K^{\#} = \langle \mathbf{1} \rangle$. If $Kx = 0$, then $Tx = c\mathbf{1}$, so $x = c\mathbf{1} + (a^{\mathsf{t}}x)\mathbf{1}_B$; the converse follows from $K\mathbf{1} = K\mathbf{1}_B = 0$. This gives the required kernel of K. If $Kx = \lambda$, put $y = Tx$. Then $a^{\mathsf{t}}y = 0$ and $K^{\#}(y - \mathbf{1}_B) = 0$, so $y - \mathbf{1}_B \in \langle \mathbf{1} \rangle$. Applying a^{t} would give $-1 = 0$, a contradiction. Thus $\lambda \notin \operatorname{im} K$.

For either support, let M be the matrix obtained from $K^{\#}$ by deleting the q_0 row and column. It is enough to make M nonsingular. Indeed, if $K^{\#}x = 0$, subtract $x_{q_0}\mathbf{1}$ from x. The resulting vector still lies in the kernel, has zero q_0 coordinate, and its remaining coordinates are killed by M. Nonsingularity makes that vector zero, proving $\ker K^{\#} = \langle \mathbf{1} \rangle$. We therefore have two tasks for the symbol choices: impose $K\mathbf{1}_B = 0$ and make these deleted matrices nonsingular on both supports.

For the smaller support, let C be a B-by-I matrix whose columns are a basis of the even-sum hyperplane

$$\{x \in \mathbb{F}_2^B : \mathbf{1}_B^{\mathsf{t}} x = 0\}.$$

These columns will form the off-diagonal block of the deleted matrix. When $|B| = 1$, C is the empty matrix. Prescribe $[q_i/q_b] = [q_b/q_i] = C_{bi}$ for $b \in B, i \in I$. Every other off-diagonal symbol incident with I is set to zero. Prescribe

$$[q_0/q_b] = \sum_{i \in I} C_{bi} \qquad (b \in B),$$

and obtain reverse symbols by reciprocity. In particular $[q_b/q_0] = [q_0/q_b] + s_b$. The q_0–J symbols remain free.

These prescriptions ensure $K\mathbf{1}_B = 0$ on both supports. At a B row, the terms from I and q_0 cancel, and the sum from J, when present, is zero by (3.20). At an I row the sum is a column sum of C, hence zero. At the q_0 row it is the sum of all entries of C plus $\sum_B s_b$, again zero. At a J row it is zero by the second equality in (3.20). It therefore remains to make the deleted matrices nonsingular.

On Q_0, the deleted matrix has the form

$$M_0 = \begin{pmatrix} A & C \\ C^{\mathsf{t}} & 0 \end{pmatrix}, \qquad \mathbf{1}_B^{\mathsf{t}} A \mathbf{1}_B = 1. \tag{3.23}$$

The equality follows from (3.21) and $\sum_B \lambda_b = 1$. The lower right block is zero since the only potentially nonzero symbols incident with I are in C, whose column sums vanish. The added matrix λa^{t} changes neither off-diagonal block, because a and λ both vanish on I. No symmetry of A is required.

To prove M_0 nonsingular, let $M_0(x, y)^{\mathsf{t}} = 0$. The second block equation says $C^{\mathsf{t}} x = 0$. Since the columns of C span the even-sum hyperplane, its annihilator is $\langle \mathbf{1}_B \rangle$; hence $x = c\mathbf{1}_B$. Left multiplying the first block equation $cA\mathbf{1}_B + Cy = 0$ by $\mathbf{1}_B^{\mathsf{t}}$ gives $c = 0$ by (3.23). The injectivity of C then gives $y = 0$. This proof applies also when $|B|$ is even; when $|B| = 1$ it says simply that $M_0 = (1)$.

On adding J, the B, I block of the deleted matrix stays equal to M_0. Indeed, the only possible changes to its diagonal are the B–J row sums, which vanish by (3.20); all I–J symbols were set to zero. Now varying the bit $[q_0/q_j]$ for one $j \in J$ also varies its reverse by the same bit, as their reciprocity discrepancy is fixed. After deleting the q_0 row and column, the only surviving change is a toggle of the j-th diagonal entry. Thus these choices independently toggle the J-diagonal entries of the enlarged matrix. Its determinant is a multilinear polynomial in these bits whose coefficient of their full product is $\det M_0 = 1$. A nonzero multilinear polynomial over $\mathbb{F}_2$ cannot vanish at every point of the Boolean cube: this follows by induction on the number of variables, writing it as $f_0 + X f_1$. Hence some choice makes the enlarged matrix nonsingular. If J is empty, the smaller-support calculation already suffices. The chosen bits do not affect the smaller matrix.

The deleted matrices are now nonsingular on both supports, so the preceding reduction proves (3.11) on both. Finally realize the prescribed symbols. Choose the auxiliary primes successively. At each step, specify the type modulo 12 and the Legendre symbols modulo every already fixed prime in $B \cup J$ or among the earlier auxiliary primes. Each symbol can be imposed by a nonzero residue modulo that prime; reverse symbols are consistent by the prescribed reciprocity rule. The Chinese remainder theorem gives a reduced residue class, and Dirichlet's theorem supplies infinitely many primes in it. Excluding the finitely many previously used primes causes no difficulty. The resulting product $n = \prod_{q \in B \cup I \cup \{q_0\}} q$ has all the asserted properties. □

Completion of the proof of Proposition 2.8. The value $b = 0$ satisfies Φ by definition. For a nonzero integer b, make the local choices preceding Lemma 3.2 and choose n by Lemma 3.6. Both $l = n$ and $l = n\theta$ satisfy the hypotheses of Lemma 3.4, so each curve has a point with nonzero x-coordinate of class β, which is also the class of b'. The isomorphism (3.6) sends the point on $E_{n\theta}$ to one on $E_{nH'}$ with the same x-coordinate square class. Taking $e = n$ gives the two points required in the definition of Φ, on E_e and $E_{eH'}$. This proves integer completeness, while Lemma 3.1 proves the asserted pole restriction. □

4 A height bound from odd contact with five points

We prove Theorem 2.4, whose contact radical supplies the primes used in Section 2. For $s \in \mathbf{P}^1(\mathbb{Q}) \setminus \mathcal{C}$, the target is

$$h(s) \le H M_S(s)^c,$$

where $M_S(s)$ is the squarefree product of the nonexceptional primes of odd contact with the five fixed points, as defined in that theorem.

Throughout this section, h is the absolute logarithmic Weil height, and h_{F} is the original stable Faltings height. Unless another dependence is stated, constants depend only on the geometric objects and auxiliary levels fixed in the proof.

The geometric part of the proof constructs a finite cover $\pi : Y \to \mathbf{P}^1$ with ramification index two exactly above $\mathcal{C}$, carrying a family of principally polarized abelian surfaces. Outside fixed exceptional primes, the local equation becomes $t = z^2$ after a finite unramified extension. Thus only odd contact can ramify the splitting field of a rational fiber, as Lemma 4.1 makes precise. The polarized isomorphism class of the surfaces varies on each component, which also allows their Faltings height to control the base height (Lemma 4.2).

The surfaces carry quaternionic multiplication, and fibers above a common base point are isogenous. For a rational base point, this gives isogenies between a surface and its Galois conjugates. Fibers with extra endomorphisms have uniformly bounded height by complex multiplication. For the other fibers, a controlled splitting of the isogeny obstruction produces a two-dimensional Galois representation. Modularity and a conductor estimate give a weight-two newform whose level is bounded by a fixed power of $M_S(s)$. The surface is then an isogeny factor of the corresponding modular Jacobian over a field of degree $O(M_S(s))$. Height bounds for that Jacobian and the quantitative isogeny theorem complete the estimate.

4.1 The five-point quaternionic quotient

We first construct the cover whose ramification will detect the odd contact orders in Theorem 2.4. Let $B/\mathbb{Q}$ be the indefinite quaternion algebra of discriminant $210 = 2 \cdot 3 \cdot 5 \cdot 7$, let $\mathcal{O}$ be a maximal order, and put $G = B^{\times}$. Fix $B \otimes_{\mathbb{Q}} \mathbf{R} \simeq M_2(\mathbf{R})$. The quaternionic Shimura datum is $(G, \mathfrak{H}^{\pm})$, where the two half planes parametrize the conjugates of the elliptic-curve Hodge homomorphism. Its reflex field is $\mathbb{Q}$: over $\mathbf{C}$ its cocharacter has the usual GL_2 conjugacy class, and this class is invariant under the Galois action for the inner form G.

We use canonical models and functoriality for Shimura varieties [16, 17]. In the form needed here, they give the model over the reflex field with complex points

$$G(\mathbb{Q})\backslash(\mathfrak{H}^{\pm} \times G(\mathbf{A}_f)/K)$$

for a compact open subgroup K, and define level maps and finite-adelic right actions over that field. We use the full Shimura varieties over $\mathbb{Q}$, including all their geometric components.

At maximal level, put

$$X = \mathrm{Sh}_{\widehat{\mathcal{O}}^{\times}}(G, \mathfrak{H}^{\pm}).$$

This curve is proper because B is a division algebra. Strong approximation for B^1, the reduced-norm theorem, and $\mathrm{Nrd}(\mathcal{O}_p^{\times}) = \mathbb{Z}_p^{\times}$ give

$$G(\mathbf{A}_f) = G(\mathbb{Q})^{+}\widehat{\mathcal{O}}^{\times}.$$

Indeed, match the finite-adelic norm modulo $\widehat{\mathbb{Z}}^{\times}$ by a positive rational reduced norm and then apply strong approximation to the norm-one part. There are also rational elements of negative norm, since no real place is ramified. Consequently $X_{\mathbf{C}}$ is the connected compact curve

$$(\mathcal{O}^1/\{\pm 1\})\backslash\mathfrak{H}.$$

The uniformizing group has no nontrivial elliptic stabilizers. A noncentral norm-one unit fixing a point has integral trace of absolute value less than two, hence trace 0 or ± 1. It would generate $\mathbb{Q}(i)$ or $\mathbb{Q}(\sqrt{-3})$. These fields cannot embed in B, since they split at the

ramified primes 5 and 7, respectively. The Eichler area formula, in the normalization of [68, Theorem 39.1.2], therefore gives

$$\frac{\operatorname{area}(X(\mathbf{C}))}{2\pi} = \frac{1}{6}\prod_{p\mid 210}(p-1) = 8.$$

There are no cusps or elliptic corrections, so $2g(X)-2=8$ and $g(X)=5$.

The finite-adelic normalizer of $\widehat{\mathcal{O}}$ induces the Atkin–Lehner group

$$\mathcal{W} \simeq (\mathbb{Z}/2\mathbb{Z})^4$$

over $\mathbb{Q}$. At a split prime, the normalizer is scalars times order units. At a division prime, its quotient by scalars and order units has order two, detected by parity of reduced-norm valuation. Finite adelic scalars act trivially, since $\mathbf{A}_f^\times = \mathbb{Q}^\times\widehat{\mathbb{Z}}^\times$. The resulting action is faithful: a generic domain point has rational stabilizer equal to the center, which accounts precisely for the scalars already removed. Let $X^* = X/\mathcal{W}$.

The geometry of this discriminant-210 quotient has been tabulated before. Long, Maclachlan, and Reid list the signature $(0;2^5;0)$, recording a genus-zero quotient with five elliptic cycles of order two and no cusps [35, Table 3]. Nualart Riera gives genus 5 for the original curve and the five involution labels $30, 42, 70, 105, 210$, and identifies the full quotient over $\mathbb{Q}$ with $\mathbf{P}^1$ [47, Propositions 4.1–4.2]. The calculation here recovers this fixed configuration and establishes the rational branch values used in the reduction.

A finite stabilizer of a point on a smooth characteristic-zero curve acts faithfully on its tangent line and is cyclic. Since every nonidentity element of $\mathcal{W}$ has order two, each nontrivial stabilizer has order two. If r is the number of geometric branch values of $X \to X^*$, Riemann–Hurwitz gives

$$2g(X^*) - 2 + \frac{r}{2} = \frac{2g(X)-2}{16} = \frac{1}{2}. \tag{4.1}$$

Thus $r = 5 - 4g(X^*) \le 5$.

We exhibit five different stabilizer labels. Take

$$m = 30, 42, 70, 105, 210.$$

The field $\mathbb{Q}(\sqrt{-m})$ is nonsplit at every prime dividing 210. For the first three values, the only prime of 210 not dividing m is respectively $7, 5, 3$, and the residues of $-m$ there are the nonsquares $5, 3, 2$. For $m = 105$, the quadratic discriminant is -420, so 2 is ramified too. For $m = 210$, all four primes are ramified. The quaternion embedding criterion from Albert–Brauer–Hasse–Noether, as in [68], gives an embedding of each field in B.

An embedded $\sqrt{-m}$ lies in some maximal order: at split local places an integral element stabilizes a lattice, and at division places it lies in the unique maximal order. Intersecting these local orders gives a global maximal order containing the embedded element. Strong approximation implies that all maximal orders in this indefinite rational quaternion algebra are conjugate by $B^\times$ [68, Theorems 28.2.10 and 28.2.11(b)], so a conjugate of the embedded element lies in our fixed $\mathcal{O}$. It normalizes $\mathcal{O}$ because it is a unit at every prime outside 210, and at a division prime every element normalizes the unique maximal order. Its norm is m and it fixes a point of $\mathfrak{H}$, so it gives a fixed point of the involution w_m with the corresponding norm-parity label.

The five labels are different. A fiber of the quotient is one $\mathcal{W}$-orbit, and the stabilizers along this orbit agree because $\mathcal{W}$ is abelian. Distinct labels thus yield distinct branch values, proving $r \ge 5$. Equation (4.1) now forces $r = 5$ and $g(X^*) = 0$. There is exactly one branch

value for each of the five labels. Every w_m is defined over $\mathbb{Q}$, so Galois preserves its uniquely labelled branch value. Each branch value is therefore individually rational. In particular X^* is a genus-zero curve with a rational point, hence is isomorphic over $\mathbb{Q}$ to $\mathbf{P}^1$. Fix an isomorphism sending one branch value to ∞, and denote the five branch values by $\mathcal{C}$.

4.2 The abelian surface family

We next pass to a fine level so that the cover carries abelian surfaces. The construction will provide a finite level map $Y \to X$, unramified over X, and a principally polarized abelian surface scheme $\mathcal{A} \to Y$. Write π for the composite

$$Y \longrightarrow X \longrightarrow X^* \simeq \mathbf{P}^1, \qquad \pi : Y \longrightarrow \mathbf{P}^1.$$

Every fiber $\mathcal{A}_y$ with $y \in Y(\overline{\mathbb{Q}})$ will have an embedding $B \hookrightarrow \mathrm{End}^0_{\overline{\mathbb{Q}}}(\mathcal{A}_y)$ given by left multiplication. Fibers with the same image under π will be geometrically isogenous, and the polarized moduli map will be nonconstant on every geometric component. We now construct the family and prove these properties.

Canonical-model functoriality associates algebraic maps to morphisms of Shimura data at compatible levels [16, 17]. A Siegel variety at principal level at least three carries the universal principally polarized abelian scheme. The required morphism to a Siegel datum comes from the following symplectic representation.

On the four-dimensional rational space $V = B$, put

$$\rho(g)(x) = x\bar{g}, \qquad \psi(x, y) = \mathrm{Trd}(\Delta x \bar{y}),$$

where the bar is the standard quaternion involution and Δ is a pure quaternion with $\Delta^2 < 0$. Such a rational Δ exists by real approximation in the trace-zero subspace. The identities $\bar{\Delta} = -\Delta$ and cyclicity of reduced trace show that ψ is alternating. It is nondegenerate because the reduced-trace pairing is nondegenerate and Δ is invertible. Moreover

$$\psi(x\bar{g}, y\bar{g}) = \mathrm{Nrd}(g)\psi(x, y).$$

Quaternion conjugation reverses products, as does composition of right multiplications; together these identities give $\rho(gh) = \rho(g)\rho(h)$. Thus ρ is a faithful symplectic-similitude representation.

The Hodge homomorphism makes $V_{\mathbf{R}}$ a complex vector space by right multiplication. At a base point represented by $J = \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$, the associated real quadratic form is

$$\psi(x, x\bar{J}) = \mathrm{tr}(J\Delta x x^{\mathrm{t}}).$$

Since Δ has trace zero, $J\Delta$ is symmetric, and its determinant is positive; it is therefore definite. Choose its sign on one half plane. Conjugation by $G(\mathbf{R})$ gives all the other complex structures, and the displayed similitude identity gives the corresponding definite sign on the other half plane as well. This is the two-component Siegel datum. A real scalar acts on V with the homological weight of an elliptic curve, which is the required weight.

Choose a rational symplectic basis for V. The standard lattice in that basis is self-dual for ψ. Intersect a sufficiently small compact open subgroup of $\widehat{\mathcal{O}}^\times$ with the inverse image of a principal Siegel level, and call the resulting subgroup K_Y. Put

$$Y = \mathrm{Sh}_{K_Y}(G, \mathfrak{H}^{\pm}).$$

The induced morphism from Y to the fine Siegel moduli scheme pulls back its universal family to the principally polarized abelian surface scheme $\mathcal{A} \to Y$. The self-dual lattice need not be preserved by the whole maximal order: the intersection of levels provides the integral moduli problem, while the rational Hodge structure retains the B-action. Fix this level and family once and for all.

There may be finitely many geometric components of the smooth projective curve Y. The level map $Y \to X$ is unramified by the uniformization and the absence of elliptic

stabilizers on X. Consequently its composite π has ramification index two at every point above $\mathcal{C}$ and index one elsewhere.

Every geometric fiber $A = \mathcal{A}_y$ has left B-multiplication in $\mathrm{End}^0_{\overline{\mathbb{Q}}}(A)$. Left multiplication commutes with the right complex structures on V; the equivalence between abelian varieties up to isogeny and their polarizable rational Hodge structures supplies the endomorphisms. Homomorphisms of abelian varieties defined over $\overline{\mathbb{Q}}$ do not acquire new elements after extension to $\mathbf{C}$ [45].

For non-CM points, the moduli construction of Guitart and Molina gives a related description by compatible isogenies [29, Lemma 1, Corollary 2, and Section 3]. For this fine-level family, the following Hodge argument gives the underlying geometric isogenies for all fibers needed here; compatible choices are made in the descent argument. Changing level or applying the normalizer changes the finite-adelic coordinate, while domain coordinates of fibers above a common point of X^* can be identified up to $G(\mathbb{Q})$. Their rational Hodge structures are therefore isomorphic, so the fibers are geometrically isogenous. In particular, if $\pi(y) = s \in \mathbf{P}^1(\mathbb{Q})$, then $\pi({}^g y) = g(s) = s$ and $\mathcal{A}_y$ is isogenous to each of its Galois conjugates.

Finally, the polarized moduli map of the family is nonconstant on every geometric component. Its complex structures vary in a half plane, whereas a fixed rational Hodge structure has only countably many presentations on the fixed rational space V. Thus a component cannot have a single geometric isomorphism class of fibers.

4.3 Rational fibers and odd-contact ramification

The ramification index two now explains the parity in the theorem. We make precise how even contact with a branch value gives unramified local lifts, and record the good reduction of the resulting surfaces.

Enlarge a fixed finite set S of rational primes whenever necessary in the following construction. All these enlargements depend only on the fixed curves, maps, and family, and occur before choosing s. Include 2, primes at which the branch sections meet, and primes at which the chosen coordinates or the geometric models degenerate. Spread Y, $\mathbf{P}^1$, π, and $\mathcal{A}$ over $\mathbb{Z}[S^{-1}]$ so that the curves are smooth and proper, π is finite and étale away from the branch sections, and

$$\pi^* b = 2D_b$$

as relative Cartier divisors, with D_b the reduced inverse image, finite étale over the ground ring. The abelian scheme also spreads over this proper model of Y.

Lemma 4.1. *For $s \in \mathbf{P}^1(\mathbb{Q}) \setminus \mathcal{C}$, put $M = M_S(s)$. There is a Galois extension $K_0/\mathbb{Q}$ of uniformly bounded degree splitting the fiber $\pi^{-1}(s)$, unramified outside $S \cup \{q : q \mid M\}$. Every fiber $A = \mathcal{A}_y$ above s is defined over K_0 and has good reduction at every place of K_0 outside S.*

Proof. The fiber has at most $\deg \pi$ geometric points, so its splitting field is Galois of degree at most $(\deg \pi)!$. Fix $q \notin S$. If the reduction of s avoids the branch sections, finite étaleness implies that the fiber is unramified at q. Otherwise it meets exactly one branch section b, and its contact order is

$$v_q(L_b(u, v)) - \min\{v_q(u), v_q(v)\}.$$

After a finite unramified local extension, all points of D_b over this reduction are defined, since D_b is finite étale. At each such point choose a formal coordinate z cutting out D_b. If t cuts out b on the base, the divisor equality gives $\pi^* t = z^2 e(z)$ with $e(z)$ a unit. After a further unramified extension its residue has a square root, and Hensel's lemma, since $q \neq 2$,

gives a square root of $e(z)$ in the formal power-series ring. Replacing z by $z\sqrt{e(z)}$ changes the local equation to $t = z^2$. Make this calculation at every point of D_b over the branch reduction, using a common further finite unramified extension.

If the contact order is even, a square root of $t(s)$ is obtained over an unramified extension by extracting the square root of its unit part. Both local lifts at every such point are consequently unramified. By properness, every geometric point of the fiber specializes to one of the points of D_b just treated, so the whole fiber is unramified. This proves that ramification can occur only at the primes in the assertion.

Every K_0-point y extends over the local valuation rings outside S by properness of Y. Pulling back the spread abelian scheme along these sections gives good reduction of $\mathcal{A}_y$. This also applies at places where $K_0/\mathbb{Q}$ is ramified. Thus field ramification may occur at primes dividing M, while good reduction over K_0 holds outside the fixed set S. □

This use of contact multiplicities has a classical specialization precedent. Darmon and Granville work with a fixed Galois cover of the projective line branched at $0, 1, \infty$. In that setting, their Proposition 3.2, credited there to Beckmann, relates divisibility of contact orders by branch indices to the absence of new ramification outside fixed bad primes in specialized fiber fields [11]. For the cover used here, the local equation $t = z^2$ makes the parity of the contact order relevant.

4.4 Comparison of base height and Faltings height

The fiber-field lemma identifies the primes that may ramify when a rational base point is lifted to the family. We also need to recover the height of that base point from the fiber. The next lemma depends only on the fixed family, and applies to every algebraic point of Y.

Lemma 4.2. *For every $y \in Y(\overline{\mathbb{Q}})$ one has*

$$h(\pi(y)) \le C\big(1 + \max\{0, h_{\mathrm{F}}(\mathcal{A}_y)\}\big), \tag{4.2}$$

with a constant independent of y and its field of definition.

Proof. We will use a theta-null map with two properties: its projective height agrees up to a fixed additive constant with Pazuki's theta height h_Θ of the fiber, and the map is nonconstant on each component of a fixed cover of Y. Here h_Θ is the projective height of the theta-null point for the corresponding polarization and level [51, Section 2.3]. The second property makes the pulled-back $\mathcal{O}(1)$ ample on the covering curve. Comparing that ample height with the pullback of the base height gives

$$h(\pi(y)) \ll 1 + h_\Theta(\mathcal{A}_y).$$

Pazuki's comparison of theta and Faltings height will then prove the lemma. We first construct a map with the two stated properties.

For every fixed even integer $r \ge 2$, the required algebraic theta data can be chosen after a finite cover of each geometric component of Y. Over a finite extension of its function field, choose a symmetric ample line bundle $\mathcal{M}$ representing the principal polarization, together with the theta data of [51, Section 2.3]. These choices require only a finite extension: a representative of a geometric principal polarization can be made symmetric by translation, and the torsion points and finite theta data are defined after a finite extension. They spread over a dense open subset of the smooth projective covering curve, where the theta-null map is algebraic. The value of this map at a fiber is the zero-section point $\Theta(0)$ constructed from Pazuki's good choices of theta data; after complex uniformization of that fiber, [51,

Sections 2.3.2–2.3.4] identifies this point, up to the finite coordinate ambiguity among good choices, with the characteristic Nullwerte displayed below. These covers and data may depend on r; we will fix the final level after proving nonconstancy.

We next compare the projective height of this map with Pazuki's theta height, uniformly in the fiber at any fixed r. For a fiber $A = \mathcal{A}_y$ and a period matrix τ for its principal polarization, the characteristic coordinates can be taken to be

$$\theta[a,b](\tau,0) = \sum_{n\in\mathbb{Z}^2} \exp\bigl(\pi i(n+a)^{\mathrm{t}}\tau(n+a) + 2\pi i(n+a)^{\mathrm{t}}b\bigr), \quad a,b \in (r^{-1}\mathbb{Z}^2)/\mathbb{Z}^2. \tag{4.3}$$

Start with a nonzero section of $\mathcal{M}$, whose space of sections has dimension one, pull it back by $[r]$, and use $[r]^*\mathcal{M} \simeq \mathcal{M}^{r^2}$. Translating by the theta group over representatives of $A[r^2]/A[r]$ gives these coordinates. Changing the generating section multiplies all coordinates by one common scalar. Two symmetric representatives differ by a two-torsion element of Pic^0, which shifts the half-characteristic; these shifts are included because r is even.

Put $k = r^2$. At the generic fiber of the chosen cover, let g_x denote the good theta-group lift of a k-torsion point x. These lifts satisfy the good-choice relation $g_x g_y = \zeta_{x,y} g_{x+y}$ with $\zeta_{x,y} \in \mu_k$. It gives $g_0 \in \mu_k$ and, by iteration at $kx = 0$, $g_x^k \in \mu_k$; hence $g_x^{k^2} = 1$. Choose $u_x = c_x g_x$ with $u_x^{2k} = 1$ by taking a $2k$-th root of the central scalar $(g_x^{2k})^{-1}$. Since $2k \mid k^2$, we have $c_x^{k^2} = u_x^{k^2} g_x^{-k^2} = 1$. These finitely many roots can be chosen once after a finite constant extension of the function field and spread with the good data. We use these normalizations only to compare coordinates of the fixed good-choice theta-null map.

In analytic theta coordinates, a lift of translation by $(m + \tau l)/k$ is given by shifting the argument and multiplying by

$$\exp(\pi i l^{\mathrm{t}}\tau l/k + 2\pi i l^{\mathrm{t}} z).$$

Theta automorphy shows that its order divides $2k$ and that it permutes $\theta[a,b](\tau, rz)$, up to bounded-order roots of unity, by $(a,b) \mapsto (a + l/r, b + m/r)$. Two lifts of the same translation whose $2k$-th powers are one differ by a scalar in μ_{2k}. Thus comparison with the algebraic u_x introduces a μ_{2k} factor after the preceding μ_{k^2} change from the good lift. All these coordinate roots of unity preserve projective height, as do permutations and a common scalar.

The customary Fourier change of theta basis is a fixed invertible matrix at this level, so it changes height by $O_r(1)$. Grouping $r(n+a)$ modulo k in (4.3) relates these coordinates blockwise by finite Fourier matrices to $\theta[e/k, 0](k\tau, 0)$, $e \bmod k$. There are r^4 independent sections, as required by $h^0(\mathcal{M}^{r^2}) = r^4$, and their values are not all zero since $\mathcal{M}^{r^2}$ is basepoint-free. Pazuki's exact theta height uses the ℓ^2 projective height [51, Definition 2.6]; its difference from the usual projective height is also bounded at this fixed level. These comparisons give one additive constant depending on r, not on the fiber or its field of definition.

It remains to choose a level at which the theta-null map is nonconstant. The following argument avoids an assumption about generic injectivity of a particular theta map. Take a connected analytic disk with a fixed homology marking on which the period matrix varies. Suppose the projective vectors in (4.3) were constant at every sufficiently divisible even level. At a chosen period τ_0, select rational b_0 with $\theta[0,b_0](\tau_0,0) \neq 0$; such a b_0 exists since the function of real b has constant Fourier coefficient one. Shrink the disk so its denominator stays nonzero. At a common level containing b and b_0, projective constancy implies that

$$\frac{\theta[0,b](\tau,0)}{\theta[0,b_0](\tau,0)}$$

is constant as τ varies, for every rational b. Density of rational b, continuity, and absolute convergence of the theta series show that the functions of real b for τ and τ_0 are proportional. Their constant Fourier coefficients are one, so they are equal. Consequently every $\exp(\pi i n^{\mathrm{t}}\tau n)$ is constant. Taking $n = e_i$ and $n = e_i + e_j$ and using continuity forces τ itself to be constant, a contradiction. A level with a nonconstant vector therefore exists. Nonconstancy persists at multiples of that level because the old characteristic coordinates occur among the new ones. A common even multiple works for the finitely many components.

Fix this common level, and make the preceding choices of covers and theta data at that level. The finitely many covering curves, maps, and open subsets are now fixed, and can all be defined over a fixed number field. On each base component, choose a smaller marked disk avoiding the branch values of the covers and the images of their omitted points, and take a connected lift to each covering curve above it. The selected holomorphic characteristic vector remains nonconstant on the smaller disk. If the theta-null map were constant on a covering curve, the preceding pointwise identification would confine that vector on the connected lift to a fixed finite set. Continuity and connectedness would then make it constant. Thus the theta-null map is nonconstant on each covering curve and extends across the omitted points. Its pullback of $\mathcal{O}(1)$ has positive degree and is ample. A sufficiently large multiple of this line bundle minus the pullback by π of $\mathcal{O}_{\mathbf{P}^1}(1)$ also has positive degree. The height machine and the lower bound for an ample height on a projective curve therefore give

$$h(\pi(y)) \le C(1 + h_\Theta(\mathcal{A}_y)) \tag{4.4}$$

on the open where the theta data are defined [3].

Here h_Θ is Pazuki's exact theta height; the bounded coordinate and projective-height comparisons have already been absorbed in (4.4). Pazuki uses the Deligne normalization

$$h_{\mathrm{F}}^{\mathrm{D}}(A) = h_{\mathrm{F}}(A) + \frac{\dim A}{2}\log\pi;$$

compare [51, Section 2.2] and [24, Section 2.3]. Here h_{F} is the original stable Faltings height. For the present surfaces the difference is the fixed constant $\log\pi$.

For $H_\Theta = \max\{h_\Theta, 1\}$ and $H_{\mathrm{D}} = \max\{h_{\mathrm{F}}^{\mathrm{D}}, 1\}$, [51, Corollary 1.3(2)] gives

$$|H_\Theta - \tfrac{1}{2}H_{\mathrm{D}}| \le C_2 \log(\min\{H_\Theta, H_{\mathrm{D}}\} + 2),$$

where C_2 depends only on dimension and level. Since $\log(H_{\mathrm{D}} + 2) \le 2H_{\mathrm{D}}$ and $H_{\mathrm{D}} \le \max\{h_{\mathrm{F}}, 1\} + \log\pi$ for these surfaces, we obtain

$$h_\Theta \le C(1 + \max\{0, h_{\mathrm{F}}\}).$$

Combining this with (4.4) proves the assertion on the chosen open sets.

The omitted points on the fixed covering curves form a finite geometric set. Their base heights are bounded, and enlarging C handles them. Absolute heights and stable Faltings heights are unchanged by field extension, so passing to points of the fixed covers introduces no dependence on their residue-field degrees. Taking the maximum over the finitely many components proves the lemma. □

Remark 4.3 (An alternative deduction on the compact moduli image)**.** The same numerical comparison can alternatively be deduced by bounding the period term in Pazuki's Theorem 1.1 for this family. That theorem compares $h_\Theta - \frac{1}{2}h_{\mathrm{F}}^{\mathrm{D}}$ with the average of logarithms of determinants of imaginary parts of Siegel-reduced periods, with bounded error at fixed dimension and level [51]. The determinant terms are bounded here. The moduli image is compact because Y is projective, and the minimum nonzero period length in the polarization

metric has a positive lower bound on it. On the subgroup $\mathbb{Z}^2$ of the period lattice the metric has matrix $(\operatorname{Im}\tau)^{-1}$. Minkowski's theorem bounds $\det \operatorname{Im}\tau$ above in terms of that minimum length, and Siegel reduction bounds it below by a positive constant. These bounds hold at every embedding, because every conjugate belongs to the fixed compact moduli image. Theorem 1.1 and the same normalization conversion therefore also give $h_\Theta \le C(1+\max\{0,h_{\mathrm{F}}\})$.

4.5 A field of definition for homomorphisms

We now choose the fiber used in the arithmetic proof. Fix $s \in \mathbf{P}^1(\mathbb{Q}) \setminus \mathcal{C}$, put $M = M_S(s)$ and $T = S \cup \{q : q \mid M\}$, and choose y above s. Let $A = \mathcal{A}_y$ and choose the splitting field K_0 from Lemma 4.1. Since π is defined over $\mathbb{Q}$ and s is rational, $\pi({}^g y) = g(s) = s$ for every $g \in G_{\mathbb{Q}}$. The common-base property of the family therefore makes the conjugates of A geometrically isogenous. The descent argument needs all their homomorphisms over one Galois field. We can define them without losing the bounded degree or introducing ramification at further primes.

Lemma 4.4. *There is a Galois extension $L/\mathbb{Q}$ containing K_0, of uniformly bounded degree, over which all geometric homomorphisms between all Galois conjugates of A are defined. It is unramified outside T.*

Proof. There are boundedly many conjugates because they are defined over K_0. Let $\mathcal{H}$ be the direct sum of the geometric Hom groups indexed by all ordered pairs of these conjugates. This is a free abelian group of bounded rank. Galois transport gives a natural action of $G_{\mathbb{Q}}$ on $\mathcal{H}$, permuting the summands according to its action on the ordered pairs. The action of G_{K_0} on $\mathcal{H}$ has finite image: a finite set of generators of the Hom groups is defined over a finite extension. Finite subgroups of $\mathrm{GL}_r(\mathbb{Z})$ have bounded order for bounded r. Let L be the fixed field of the kernel of this action inside G_{K_0}. The transport action of $G_{\mathbb{Q}}$, together with normality of G_{K_0}, makes that kernel normal in $G_{\mathbb{Q}}$. Thus $L/\mathbb{Q}$ is Galois and its degree is uniformly bounded.

At a place outside T, the extension $K_0/\mathbb{Q}$ is unramified and all the conjugates have good reduction. For a coefficient prime different from the residue characteristic, inertia acts trivially on their Tate modules and therefore on the Hom groups embedded faithfully between those modules. Thus L/K_0 introduces no ramification there. □

All homomorphisms between conjugates, including the prescribed B-action, are now defined over L. Its degree is uniformly bounded, and it is unramified outside T. Extending K_0 to L preserves good reduction of A outside the fixed set S. We now bound $h_{\mathrm{F}}(A)$ in terms of M, beginning with its endomorphism algebra.

4.6 The endomorphism dichotomy and descent up to isogeny

The endomorphism algebra determines which comparison will give a height bound for A. Put $D = \operatorname{End}^0_{\mathbb{Q}}(A)$, and let C_B be the centralizer of the prescribed left B-action in D. On rational homology the centralizer of left B is right B. A Hodge endomorphism in that centralizer must commute, over $\mathbf{R}$, with the complex structure, whose centralizer in $M_2(\mathbf{R})$ is $\mathbf{C}$. Thus C_B is either $\mathbb{Q}$ or an imaginary quadratic field k contained in B. More explicitly, a nonscalar rational element commuting with the complex structure generates such a quadratic field, and its real centralizer has dimension two, leaving no further possibilities.

Both cases use the following quantitative isogeny theorem to transfer a known height bound. If C_1, C_2 are g-dimensional abelian varieties over a number field k_0 of degree d, and

are isogenous over an extension K/k_0, then there are K-isogenies in both directions with degree at most

$$\kappa(C_1) = \left((14g)^{64g^2} d \max\{h_{\mathrm{F}}(C_1), \log d, 1\}^2\right)^{1024g^3}. \tag{4.5}$$

This is [23, Theorem 1.4, with the definition on p. 2058]. It uses the original stable Faltings height and requires neither semistability over k_0 nor a principal polarization. In our applications we take $K = k_0$. The stable height is additive on products and satisfies

$$h_{\mathrm{F}}(C_2) \le h_{\mathrm{F}}(C_1) + \tfrac{1}{2} \log \deg \phi \quad \text{for an isogeny } \phi : C_1 \longrightarrow C_2 \tag{4.6}$$

[20]. We will always apply (4.5) to a source whose height has already been bounded.

If $C_B = k$, the central-simple-algebra double-centralizer calculation gives

$$D = B \otimes_{\mathbb{Q}} k \simeq M_2(k).$$

The last isomorphism holds because an embedded quadratic field is a maximal subfield of B and splits it. Matrix idempotents, all defined over L by Lemma 4.4, give $A \sim_L E^2$ for a CM elliptic curve E/L: take the image of an integer multiple of a rank-one idempotent, and use the matrix units to identify the two elliptic factors up to isogeny. The degree of $j(E)$ is at most $[L : \mathbb{Q}]$, hence is bounded. Complex multiplication theory identifies the degree of a singular modulus with the class number of its imaginary quadratic order. Class-number growth for imaginary quadratic fields and the formula for class numbers of orders imply that only finitely many such j-invariants have bounded degree [10, 65]. Their stable elliptic heights, and therefore $h_{\mathrm{F}}(E^2)$, are bounded. Apply (4.5) to the source E^2, with dimension two, bounded height, and bounded $[L : \mathbb{Q}]$, and then apply (4.6). This bounds $h_{\mathrm{F}}(A)$ uniformly. Lemma 4.2 already proves Theorem 2.4 for these fibers, since $M \ge 1$. The finiteness invoked here need not be effective.

We henceforth treat the other case, $C_B = \mathbb{Q}$. The same double-centralizer calculation gives $D = B$. One can check this calculation after an extension splitting B: an algebra containing the full 2×2 matrix units is a matrix algebra over their centralizer. Because B is division, A is absolutely simple; a proper isogeny factor would give a nontrivial idempotent in D.

Our goal in this case is to find a primitive weight-two newform f of level N, a Galois extension $L'/\mathbb{Q}$ containing L, and an abelian variety D_0/L' such that, for fixed constants $C, a > 0$,

$$[L' : \mathbb{Q}] \le CM, \qquad N \le CM^a, \qquad J_1(N) \sim_{L'} A \times D_0.$$

Once the height of this modular Jacobian has been bounded in terms of N, the factorization will let us transfer the bound to A over a field of controlled degree. We first use the conjugacy isogenies to construct a two-dimensional representation. Their discrepancy under composition will be reduced to a sign and then split with control on both the field degree and inertia.

Compatible-isogeny cocycles and their scalar splittings are part of the theory of $\mathbb{Q}$-curves and building blocks [59, Section 6][56, Chapters 4–5]; compare [57, Section 4]. The scalar cochain constructed below takes values among roots of unity of order $O(M)$, and its inertia restrictions will control the modular level.

Put $\Gamma = \mathrm{Gal}(L/\mathbb{Q})$. For each $g \in \Gamma$, choose a quasi-isogeny $\mu_g : {}^g A \to A$, with $\mu_1 = 1$, satisfying

$$\mu_g \, {}^g\alpha = \alpha \mu_g \qquad (\alpha \in B). \tag{4.7}$$

To obtain it, start with an isogeny between the fibers above the common rational point. Transporting the B-action along this isogeny gives an automorphism of $B = D$, which is

inner by Skolem–Noether. Composing with the corresponding element of $B^\times$ gives (4.7). All these maps are defined over L by Lemma 4.4.

The maps $\mu_g\,{}^g\mu_h$ and μ_{gh} have the same source and target and both intertwine the B-actions. Their quotient therefore lies in the centralizer $\mathbb{Q}^\times$:

$$\mu_g\,{}^g\mu_h = c(g,h)\mu_{gh}, \qquad c(g,h) \in \mathbb{Q}^\times.$$

For a quasi-isogeny, use the degree obtained by extending the geometric degree multiplicatively from isogenies. A rational scalar c on a surface has degree c^4, so the last identity gives

$$|c(g,h)| = \left(\frac{\deg \mu_g \,\deg \mu_h}{\deg \mu_{gh}}\right)^{1/4}.$$

Let $a_g = (\deg \mu_g)^{-1/4}$ be the positive real algebraic fourth root. We extend the coefficient scalars in the isogeny category and put

$$\widetilde{\mu}_g = a_g\mu_g \in \mathrm{Hom}^0_L({}^gA, A) \otimes_{\mathbb{Q}} \overline{\mathbb{Q}}.$$

Here Galois transports the Hom factor and acts trivially on the coefficient factor $\overline{\mathbb{Q}}$. The preceding degree identity now gives

$$\widetilde{\mu}_g\,{}^g\widetilde{\mu}_h = \epsilon(g,h)\widetilde{\mu}_{gh}, \qquad \epsilon(g,h) = \frac{c(g,h)}{|c(g,h)|} \in \{\pm 1\}. \tag{4.8}$$

Associativity makes ϵ a normalized cocycle on Γ; we inflate it to $G_{\mathbb{Q}}$. This is the obstruction in the isogeny version of Weil descent. The normalized discrepancy depends only on the sign of $c(g,h)$, so we will not need bounds for the individual degrees $\deg \mu_g$.

A scalar cochain with coboundary ϵ will cancel this sign. On G_L the inflated cocycle is trivial, so the restriction of such a cochain will be a character. We seek one whose restricted character is conjugacy invariant and has controlled order: its kernel will then define a Galois field over which the corrected action agrees with the Tate action of A. Control on inertia will also be needed to bound the modular level. The following splitting supplies both properties.

Lemma 4.5 (Controlled splitting of the sign obstruction)**.** *There is a finite Dirichlet character χ of order $m \le CM$, unramified outside T, and a continuous cochain $\beta : G_{\mathbb{Q}} \to \mu_{2m}$ such that*

$$\beta^2 = \chi, \qquad \frac{\beta(g)\beta(h)}{\beta(gh)} = \epsilon(g,h). \tag{4.9}$$

The cochain can be chosen trivial on inertia at every prime outside T. Its restriction to G_L is a conjugacy-invariant character. The local conductor of χ at any fixed prime is bounded independently of s, and χ is tame at odd primes.

Proof. Identify $H^2(G_{\mathbb{Q}}, \{\pm 1\}) = \mathrm{Br}(\mathbb{Q})[2]$. At a prime outside T, the cocycle factors through the unramified quotient $\widehat{\mathbb{Z}}$, which has cohomological dimension one on torsion coefficients. Its local class is therefore zero. We use local reciprocity and the Brauer local-global theorem in the form described in [64, 42, 41].

For a finite character χ_v with values in $\overline{\mathbb{Q}}^\times$ with trivial coefficient action, its square-root obstruction $\partial\chi_v$ in $H^2(G_{\mathbb{Q}_v}, \{\pm 1\})$ vanishes if and only if χ_v has a continuous character square root. By reciprocity this is equivalent to its being trivial on the order-two subgroup of $\mathbb{Q}_v^\times$, that is, on -1. To verify the equivalence, use $\mathbb{Z}_p^\times \simeq \mu_{p-1} \times \mathbb{Z}_p$ for odd p, and $\mathbb{Z}_2^\times \simeq \{\pm 1\} \times \mathbb{Z}_2$ at 2. A character of the infinite cyclic valuation factor has a square root by choosing a square root of its value on a uniformizer; the procyclic free factors also admit

character square roots. The only obstruction in the finite cyclic torsion factors is the value on their element of order two. At the real place the same assertion follows directly from $G_{\mathbf{R}} = \mathbb{Z}/2\mathbb{Z}$.

At each odd finite prime with nonzero local class of ϵ, choose an odd character modulo that prime, so its value at -1 is -1. At 2, if necessary, use the nontrivial character modulo 4. Their product is a global Dirichlet character χ. Its local square-root obstruction matches that of ϵ at every finite place. The real invariant matches too, because the sum of local invariants of a Brauer class is zero. Injectivity in the Brauer local-global theorem gives $\partial\chi = [\epsilon]$.

Each odd-prime factor has order at most $q-1$, and the factor at 2 has order two. Thus the order m of their product satisfies

$$m \le 2 \prod_{q \in T,\ q \neq 2} (q-1) \le CM.$$

The conductor at an odd prime divides that prime, and the conductor at 2 divides 4. Factors at other primes are unramified locally; this proves the stated local conductor assertions.

Choose a continuous set-theoretic square root of χ, valued in μ_{2m}. Equality of the obstruction classes allows its coboundary to be corrected to exactly ϵ by a sign-valued continuous cochain. The resulting β satisfies (4.9). On G_L the inflated cocycle is identically one, so β restricts to a character. Its conjugacy invariance follows from the cochain identity: if $u \in G_L$ and $g \in G_{\mathbb{Q}}$, then $\epsilon(g,u) = \epsilon(gug^{-1}, g) = 1$, and comparison of $\beta(gu)$ computed in these two ways gives $\beta(gug^{-1}) = \beta(u)$.

At an odd prime outside T, inertia fixes L, and χ is unramified, so the restriction of β to inertia is a sign character. Only finitely many such restrictions are nontrivial, because a continuous function to the finite set μ_{2m} factors through a finite quotient of $G_{\mathbb{Q}}$. At an odd prime there is a unique nontrivial quadratic character of inertia: wild inertia is pro-odd and the tame quotient has a unique quotient of order two. Let d_β be a squarefree integer whose odd prime factors are precisely these offending primes. Multiplication of β by the quadratic character of $\mathbb{Q}(\sqrt{d_\beta})$ cancels all the unwanted inertia characters. This quadratic character is unramified away from those primes and possibly $2 \in S$. Its square is one and its coboundary is one, so (4.9) and the restriction properties are preserved. Since $\{\pm 1\} \subset \mu_{2m}$, the corrected cochain still takes values in μ_{2m}, whose elements have orders dividing $2m \le CM$. □

4.7 Two-dimensional representations and modularity

We construct representations at coefficient primes 2 and 3. The 2-adic one will yield a modular form. The 3-adic one will measure the conductor at 2, where the conductor comparison requires a coefficient prime different from the residue prime. This requires identifying both representations with the same newform.

Fix coefficient embeddings $\overline{\mathbb{Q}} \hookrightarrow \overline{\mathbb{Q}}_\ell$ for $\ell = 2, 3$. On the covariant rational Tate module of A with these extended coefficients put

$$R_\ell(g) = \beta(g)\widetilde{\mu}_{\bar{g}} \circ g, \qquad g \in G_{\mathbb{Q}}, \quad \bar{g} = g|_L. \tag{4.10}$$

Here g first transports torsion points to ${}^g A$, and the quasi-isogeny returns them to A. In a composition the discrepancy of the maps is $\epsilon(g,h)$, and the discrepancy of the scalar cochain is the same sign. Their product is one. Thus R_ℓ is a genuine continuous representation. All algebraic coefficients and the splitting of B can be taken in one finite extension of $\mathbb{Q}_\ell$ for each ℓ.

Equation (4.7) shows that R_ℓ commutes with B. After coefficient extension, $B \otimes \overline{\mathbb{Q}}_\ell \simeq M_2(\overline{\mathbb{Q}}_\ell)$; the matrix units then give

$$R_\ell \simeq r_\ell \oplus r_\ell \tag{4.11}$$

for a two-dimensional representation r_ℓ. On G_L the quasi-isogeny in (4.10) is $\widetilde{\mu}_1 = 1$, so only the character $\beta|_{G_L}$ separates R_ℓ from the Tate action of A. Define

$$G_{L'} = \ker(\beta|_{G_L}).$$

Conjugacy invariance of this character makes $L'/\mathbb{Q}$ Galois, and its image has order at most $2m$. Consequently

$$[L' : \mathbb{Q}] \leq 2m[L : \mathbb{Q}] \leq CM. \tag{4.12}$$

On $G_{L'}$, the representation R_ℓ is exactly the Tate-module representation of A, with coefficients extended. This field will allow us to recover an abelian factor over L'. Its degree may grow with M; the local conductor argument will use L, whose degree is uniformly bounded.

The modularity input is [49, Theorem 1.1] in its exact form: every continuous irreducible odd two-dimensional 2-adic representation of $G_\mathbb{Q}$, unramified outside finitely many primes and de Rham at 2 with distinct Hodge–Tate weights, is a Tate twist of the representation associated with a classical cuspidal eigenform. There is no residual-image hypothesis in that statement. We now verify its hypotheses for r_2. The same arguments also give the properties of r_3 used below. The commutant and Hodge-decomposition arguments have antecedents in [59, Section 3]; we apply them to the corrected action (4.10).

By Faltings' semisimplicity and endomorphism theorems,

$$\mathrm{End}_{G_{L'}}(V_\ell(A)) = \mathrm{End}_{L'}(A) \otimes \mathbb{Q}_\ell = B \otimes \mathbb{Q}_\ell$$

[20]. Therefore $r_\ell|_{G_{L'}}$ is semisimple with scalar commutant after algebraic coefficient extension, and is absolutely irreducible. So is r_ℓ.

The de Rham comparison theorem for smooth proper varieties, together with its Hodge–Tate decomposition, applies to A over every completion of L' at its coefficient prime [61, Corollary 1.8]; see also the corrected definitions in [62]. Since the covariant Tate module is dual to first étale cohomology, $V_\ell(A)$ is de Rham with cyclotomic exponents $0, 1$, each twice. The de Rham property descends through finite local extensions [21, Section 3.9(ii)]. Thus R_ℓ is de Rham over $\mathbb{Q}_\ell$, and its decomposition (4.11) gives exponents $0, 1$, each once, for r_ℓ. Here the cyclotomic character has exponent 1 and Hodge–Tate weight -1.

For complex conjugation j, the involution $r_\ell(j)$ cannot be scalar. If it were, (4.11) would make $R_\ell(j)$ scalar too. Under comparison with Betti homology, $\mu_{\bar{j}} \circ j$ is conjugation followed by a holomorphic quasi-isogeny. It exchanges the two nonzero Hodge subspaces, whereas a scalar preserves both. Multiplying by the nonzero algebraic coefficient in (4.10) does not change this fact. Hence the two eigenvalues of $r_\ell(j)$ are 1 and -1, and r_ℓ is odd.

At a prime outside $T \cup \{\ell\}$, inertia fixes L, A has good reduction, and the corrected β is trivial. Thus r_ℓ is unramified there. In particular it is ramified at only finitely many primes. All hypotheses of the stated modularity theorem now hold for r_2. Applying it realizes r_2 as a Tate twist of the representation associated with a classical cuspidal eigenform.

We use the arithmetic-Frobenius normalization in which a weight-k form f has good traces $a_q(f)$ and determinant $\psi_f \chi_{\mathrm{cyc}}^{k-1}$. Its cyclotomic exponents are $0, k-1$; the companion fixes this convention in [49, Section 2.1]. The two exponents $0, 1$ of r_2 therefore force weight two and Tate shift zero. Thus r_2 is realized by the arithmetic 2-adic representation of a normalized primitive weight-two newform.

For comparison with a dual presentation, if an intermediate realization is written as $V_{f,2}^{\vee} \otimes \chi_{\mathrm{cyc}}^{n}$, then

$$V_{f,2}^{\vee} \otimes \chi_{\mathrm{cyc}}^{n} \simeq V_{g,2} \otimes \chi_{\mathrm{cyc}}^{n+1-k}, \qquad g = (f \otimes \psi_f^{-1})_{\mathrm{new}}.$$

Here g is the primitive newform of the indicated finite-character twist. The target exponents $0, 1$ give $k = 2$, $n = 1$ in this presentation, and hence arithmetic Tate exponent zero. Changing from arithmetic to geometric Frobenius inverts the evaluated eigenvalues; it does not dualize the representation. After this conversion when necessary, let f denote the normalized primitive newform whose arithmetic 2-adic representation realizes r_2, and let N be its level.

To use r_3 for the dyadic conductor of this level, we must identify it with the same algebraic conjugate of f. A single algebraic trace, obtained from the good special fiber, will compare the two coefficient primes.

Lemma 4.6. *The representation r_3 is the 3-adic representation of the same algebraic conjugate of f that realizes r_2.*

Proof. Fix a rational prime $q \notin T \cup \{2, 3\}$ with $q \nmid N$, a place of $\overline{\mathbb{Q}}$ above q, and an arithmetic Frobenius g in that decomposition group. These primes form a cofinite set, as needed for the Chebotarev argument below. The automorphism induced by g on the residue field $\overline{\mathbf{F}}_q$ is $x \mapsto x^q$. Therefore the reduction of ${}^g A$ is the Frobenius twist $\widetilde{A}^{(q)}$ of the same geometric special fiber $\widetilde{A}$. Specialization of prime-to-q torsion identifies the action of g with the relative Frobenius map $\widetilde{A} \to \widetilde{A}^{(q)}$. A multiple of $\mu_{\bar{g}}$ extends over good reduction, and its reduction, after dividing by that integer, maps $\widetilde{A}^{(q)}$ back to $\widetilde{A}$. The composite is thus an element

$$u_q = \mathrm{red}(\mu_{\bar{g}}) \circ \mathrm{Frob}_q^{\mathrm{rel}} \in \mathrm{End}^0(\widetilde{A}).$$

Here $\mathrm{red}(\mu_{\bar{g}})$ is reduction of the unnormalized quasi-isogeny; the scalar $a_{\bar{g}}$ from degree normalization is inserted below.

The characteristic polynomial of an endomorphism of an abelian variety on rational Tate modules has rational coefficients and is independent of the coefficient prime. This applies to a rational endomorphism after clearing denominators [40, I, Theorem 10.9 and Proposition 10.20]. Consequently

$$\alpha_q = \frac{1}{2}\beta(g) a_{\bar{g}} \,\mathrm{Tr}(u_q) \in \overline{\mathbb{Q}}$$

is a single algebraic number whose images at 2 and 3 are the traces of $r_2(g)$ and $r_3(g)$, respectively. The factor $1/2$ comes from (4.11). This construction is independent of any comparison of coefficient-field degrees.

The identification of r_2 with the representation of f selects an embedding of its Hecke field in $\overline{\mathbb{Q}}$. Injectivity of the fixed embedding $\overline{\mathbb{Q}} \hookrightarrow \overline{\mathbb{Q}}_2$ then gives $\alpha_q = a_q(f)$ as algebraic numbers at every good prime. Under the chosen embedding at 3, the same equalities hold. Chebotarev density and the character criterion for semisimple representations [40, IV, Theorem 3.3 and Lemma 3.9] identify r_3 with that member of the compatible system of f [15]. Semisimplicity on our side follows already from irreducibility. □

4.8 A uniform bound for the modular level

Lemma 4.7. *There are fixed constants $C, a > 0$ such that*

$$N \le C M^a. \tag{4.13}$$

Proof. For a primitive weight-two form, the local newform conductor is the exponent of q in its primitive level. Local–global compatibility, with conductor preservation under the local correspondence, identifies this with the Artin-conductor exponent of its ℓ-adic representation whenever $q \neq \ell$ [5, Theorem A in Section 0.7, with the normalization in Section 0.5]. The conductor includes the Weil–Deligne monodromy contribution. We use r_2 at all primes other than 2, and r_3 at 2, justified by Lemma 4.6. The exponent is zero at $q \notin T$ by the unramified assertion above.

First let $q \in T$ be larger than a fixed constant that includes all primes in S, the primes $2, 3$, and the uniform upper bound for $[L : \mathbb{Q}]$. Wild inertia, a pro-q group, acts trivially on L. The surface has good reduction over L at this prime, so wild inertia also acts trivially on its prime-to-q Tate modules. On this subgroup β is a character, and its square is the tame character χ. Hence its wild restriction is sign-valued and must be trivial, as q is odd. Formula (4.10) makes R_ℓ tame at q, and hence r_ℓ is tame there as well. Its dimension is two, so its conductor exponent is at most two.

It remains to bound the exponents at the fixed finite set of smaller primes. Fix one of them, say q, a place v of L above it, and put $F = L_v$. Let $\ell \neq q$ be the coefficient prime chosen above. Only finitely many embedded extensions $F/\mathbb{Q}_q$ can occur, because their degrees are bounded. We seek a bound u_q, independent of s, such that for every $u > u_q$ the upper ramification group $G^u_{\mathbb{Q}_q}$ lies in G_F, is killed by the character $\beta|_{G_F}$, and acts trivially on $V_\ell(A)$. On such a group $g|_L = 1$, so the quasi-isogeny in $R_\ell(g)$ is also the identity. We obtain this bound by controlling the character and the Tate action over the finitely many possible fields F.

First consider the character. By local reciprocity, the square of $\beta|_{G_F}$ corresponds to $\chi \circ \mathrm{Norm}_{F/\mathbb{Q}_q}$. Although the global order of χ may grow with M, its conductor on $\mathbb{Q}_q$-units is uniformly bounded: of the constructed Dirichlet factors only the factor at q ramifies there. Choose an integer b large enough, uniformly over the finitely many possibilities for F, that this square character is trivial on $U^b_F = 1 + \mathfrak{m}^b_F$ and that the local logarithm identifies U^b_F with $\mathfrak{m}^b_F$. Then $\beta(U^b_F) \subset \{\pm 1\}$. For odd q, squaring maps U^b_F onto itself, and therefore β kills U^b_F. For $q = 2$, writing $e_F = v_F(2)$ gives

$$(U^b_F)^2 = U^{b+e_F}_F$$

in these logarithmic coordinates. Thus β kills a uniformly deeper unit group also at 2. Reciprocity gives a uniform upper ramification cut for this character. The quadratic correction in Lemma 4.5 does not affect the argument, because it leaves $\beta^2 = \chi$ unchanged. An arbitrarily large unramified character order has no effect on a unit-filtration bound.

Next consider the Tate action. Choose a fixed prime $p_0 \geq 3$, different from q, and put

$$F_1 = F(A[p_0]).$$

Its degree satisfies $[F_1 : F] \leq |\mathrm{GL}_4(\mathbb{F}_{p_0})|$. The semistable reduction criterion says that inertia is unipotent over F_1 [66, Theorem 3.5]. Its wild pro-q subgroup has trivial action on $V_\ell(A)$, since a compact unipotent subgroup of $\mathrm{GL}_4(\mathbb{Q}_\ell)$ has no nontrivial pro-q subgroup when $\ell \neq q$. Since both $[F : \mathbb{Q}_q]$ and $[F_1 : F]$ are bounded, the fields F_1 and their normal closures over $\mathbb{Q}_q$ also range over finite sets.

Choose an upper ramification cut over $\mathbb{Q}_q$ exceeding the breaks of all these finite field extensions and whose images under the finitely many relevant Herbrand inverse functions exceed the cuts for $\beta|_{G_F}$. The subgroup and quotient rules for upper ramification groups then place every sufficiently high $G^u_{\mathbb{Q}_q}$ inside G_F, in the kernel of $\beta|_{G_F}$, and in the wild Tate-module kernel over F_1 [64]. On this group $\widetilde{\mu}_{\bar{g}} = \widetilde{\mu}_1 = 1$, so (4.10) makes R_ℓ trivial there. Its decomposition (4.11) then makes r_ℓ trivial there as well. We have therefore obtained the

required uniform upper cut u_q for r_ℓ, independent of s. Its Swan conductor is at most $2u_q$. The remaining contribution from inertia invariants, including Weil–Deligne monodromy, is at most two because r_ℓ has dimension two. This bounds its full conductor exponent at the fixed prime.

The finitely many fixed primes contribute a constant to N. Every remaining ramified prime divides M and has exponent at most two, so their product contributes at most M^2. This proves (4.13), and the proof allows $a = 2$ after altering the fixed constant. The uniform local bounds use the bounded degree of L; (4.12) supplies the degree of the field L' for the global isogeny estimate. □

4.9 The modular factor and the final height bound

For elliptic curves over $\mathbb{Q}$, Murty and Pasten earlier obtained an effective $O(N \log N)$ bound for Faltings height in terms of the conductor, with applications to S-unit equations [46, Theorems 1.1 and 7.1]. The abelian-variety comparison used in this final step follows the modular height strategy of von Känel [70, Sections 1.2 and 5]; see also [1]. The preceding descent and ramification arguments provide the point-dependent fields and the level bound needed to apply that strategy here.

Let $J = J_1(N)$. The weight-two Eichler–Shimura realization places r_2 in $V_2(J) \otimes \overline{\mathbb{Q}}_2$ [58, Theorem 2.1 and the following weight-two realization]. We will turn this shared constituent into a nonzero homomorphism $A \to J$ over L'. The next argument descends a Hom from the extended coefficient field to $\mathbb{Q}_2$, giving the abelian factor over L' itself.

Lemma 4.8. *There is an abelian variety D_0/L' such that*

$$J \sim_{L'} A \times D_0.$$

Proof. Put $V = V_2(A)$, $W = V_2(J)$, and $H = G_{L'}$. After extending coefficients, V is two copies of $r_2|_H$, while W contains at least one copy. Hence $\mathrm{Hom}_H(V \otimes \overline{\mathbb{Q}}_2, W \otimes \overline{\mathbb{Q}}_2) \neq 0$. Invariants commute with coefficient extension in this situation. Indeed, in the finite-dimensional space $\mathrm{Hom}_{\mathbb{Q}_2}(V, W)$ they are the intersection of the kernels of the linear equations expressing commutation with each $h \in H$. Finitely many of these equations cut out that same intersection, by stabilization of dimensions. Kernels of this finite system commute with scalar extension. Thus

$$\mathrm{Hom}_H(V, W) \otimes \overline{\mathbb{Q}}_2 = \mathrm{Hom}_H(V \otimes \overline{\mathbb{Q}}_2, W \otimes \overline{\mathbb{Q}}_2) \neq 0.$$

Faltings' homomorphism theorem identifies $\mathrm{Hom}_H(V, W)$ with $\mathrm{Hom}_{L'}(A, J) \otimes \mathbb{Q}_2$, so an actual nonzero L'-homomorphism $A \to J$ exists. Absolute simplicity makes its kernel finite. Poincare complete reducibility over L' then gives the claimed factor [20, 45]. A single shared constituent was sufficient to produce the Hom; the resulting rational Tate-module injection automatically supplies its full required multiplicity. In particular $\dim J \geq 2$. □

The controlled field, level, and factorization sought in the non-CM case are now established. We turn to the height of the modular Jacobian. Write $g_0 = \dim J$. The modular group index bounds the degree of the map $j : X_1(N) \to \mathbf{P}^1$ by CN^2. It is branched only over $0, 1728, \infty$, so after rescaling the target it is a Belyi map. Riemann–Hurwitz, or the usual genus formula, also gives $g_0 \leq CN^2$. Javanpeykar's theorem [32, Theorem 1.1.1 and Section 2.3] for a smooth projective connected curve X of genus $g \geq 1$ bounds its Jacobian's stable height in his convention by

$$13 \cdot 10^6 g \, \deg_B(X)^5.$$

Javanpeykar's height here equals the original stable Faltings height of the Jacobian used in (4.5). The addition of $(g/2)\log\pi$ in [32, Lemma 2.4.4] converts instead to the Deligne normalization $h_{\mathrm{F}}^{\mathrm{D}}$ of [24, Section 2.3]. Thus no normalization correction is needed here; in either convention the difference is only $O(g)$. Since Lemma 4.8 ensures $g_0 \geq 2$, the genus hypothesis applies. We obtain

$$g_0 \leq CN^2, \qquad \max\{0, h_{\mathrm{F}}(J)\} \leq CN^{12}. \tag{4.14}$$

For the final comparison we choose a principally polarized complement. An abelian variety over a number field has a polarization defined over that field and is consequently isogenous there to its dual. When $D_0 \neq 0$, choose such a polarization over L'. Applying Zarhin's trick to it gives a principal polarization over L' on

$$D_Z = (D_0 \times D_0^{\vee})^4$$

[74, Theorem 1.1(ii) and Remark 1.4]; see also [4, proof of Corollary 3.2]. If $D_0 = 0$, omit this factor and take all its height and dimension contributions below to be zero. In either case this choice of complement gives

$$C_1 = J^8 \sim_{L'} C_2 = A^8 \times D_Z.$$

Both sides have dimension $g = 8g_0$. Their common field has degree $d = [L' : \mathbb{Q}] \leq CM$, and the source has known stable height $h_{\mathrm{F}}(C_1) = 8h_{\mathrm{F}}(J)$. From (4.13) and (4.14),

$$g \leq CM^{2a}, \qquad \max\{0, h_{\mathrm{F}}(C_1)\} \leq CM^{12a}.$$

Applying (4.5) to this source gives an isogeny $\phi : C_1 \to C_2$ whose logarithmic degree is at most

$$\begin{aligned} \log\deg\phi &\leq 1024g^3\bigl(64g^2\log(14g) + \log d + 2\log\max\{h_{\mathrm{F}}(C_1), \log d, 1\}\bigr) \\ &\leq CM^C. \end{aligned} \tag{4.15}$$

The last inequality follows by substituting the preceding bounds: every power of g is a fixed power of M, while the remaining logarithms are bounded by $C(1 + \log M) \leq C'M$ for $M \geq 1$. This is the point at which the explicit dependence on dimension in the isogeny theorem is needed. The theorem is applied with source J^8, whose height is already bounded.

By additivity and (4.6),

$$8h_{\mathrm{F}}(A) + h_{\mathrm{F}}(D_Z) \leq h_{\mathrm{F}}(C_1) + \tfrac{1}{2}\log\deg\phi \leq CM^C.$$

Bost's lower bound applies to every abelian variety over a number field. For the complement D_Z chosen above, the form in [24, Corollary 8.4] gives

$$h_{\mathrm{F}}^{\mathrm{D}}(D_Z) \geq -\tfrac{1}{2}(\dim D_Z)\log(2\pi).$$

The normalization identity $h_{\mathrm{F}}^{\mathrm{D}} = h_{\mathrm{F}} + \frac{1}{2}(\dim D_Z)\log\pi$ therefore yields

$$h_{\mathrm{F}}(D_Z) \geq -\tfrac{1}{2}(\dim D_Z)\log(2\pi^2).$$

Since $\dim D_Z \leq g \leq CM^{2a}$, it follows that

$$h_{\mathrm{F}}(A) \leq CM^C. \tag{4.16}$$

Neither a degree bound for the original μ_g nor a bound for the degree of a field of coefficients has entered this estimate.

Completion of the proof of Theorem 2.4. For the arbitrary point s fixed above, let $A = \mathcal{A}_y$ be the chosen fiber. The CM case gave a uniform upper bound for $h_{\mathrm{F}}(A)$; the other case gave (4.16). Applying Lemma 4.2 gives $h(s) \le HM_S(s)^c$ after enlarging fixed constants H, c. All curves, levels, theta data, exceptional primes, and comparison constants were fixed independently of s. This proves the theorem. Since $M_S(s) \ge 1$, the exponent c may, if desired, be increased to a positive integer. □

Remark 4.9 (Related height bounds)**.** Let U be a nonempty open subscheme of $\operatorname{Spec} \mathcal{O}_K$ for a number field K. For fixed relative dimension and $[K : \mathbb{Q}]$, related work bounds stable Faltings height polynomially in the radical of the product of the absolute discriminant of K and the norms of the prime ideals outside U. Von Känel proves such a bound for abelian schemes over U whose geometric generic fiber is simple, non-CM, and $\mathbb{Q}$-virtual of GL_2-type, with conjugate isogenies compatible with all geometric endomorphisms [69, Proposition 9.9]. Von Känel and Kret treat the product-GL_2-type class with $G_{\mathbb{Q}}$-isogenies and also abelian schemes with geometric CM [71, Theorem 7.1].

5 The rank-one elliptic group and its local parameters

We prove Propositions 2.1 and 2.2, the two elliptic inputs used in Section 2.1. Recall the fixed data of (2.1):

$$F = \mathbb{Q}(\sqrt{2}), \qquad d = 75 - 53\sqrt{2}, \qquad E:\ y^2 = x^3 - d^2 x,$$

with σ the nontrivial automorphism of $F/\mathbb{Q}$. The algebraic integer d is positive at both real embeddings and has norm 7. In particular, $\mathfrak{d} = (d)$ is a prime ideal with residue field $\mathbb{F}_7$.

In Section 6, we will choose split primes of good reduction for E and its conjugate E^σ at which 7 is a nonsquare. The identity $dd^\sigma = 7$ then makes their reductions quadratic twists with opposite Frobenius traces.

Proof of Proposition 2.1. The ring of integers is $\mathcal{O}_F = \mathbb{Z}[\sqrt{2}]$. Minkowski's bound is $\sqrt{8}/2 = \sqrt{2} < 2$, so every ideal class contains an ideal of norm one. Thus $\mathcal{O}_F$ is a principal ideal domain. The unit theorem gives $|\mathcal{O}_F^*/\mathcal{O}_F^{*2}| = 4$. The units -1 and $1 + \sqrt{2}$ realize all four sign patterns at the real embeddings, so a totally positive unit is a square.

Consider the two-isogenous curve

$$E':\ y^2 = x^3 + 4d^2 x.$$

We use the usual two-isogeny descent, with isogenies $\varphi : E \to E'$ and $\widehat{\varphi} : E' \to E$ whose composite is [2]; see [67, Chapter X, Proposition 4.7 and Example 4.9, pp. 336–337]. Its homomorphisms

$$\alpha : E(F) \longrightarrow F^*/F^{*2}, \qquad \alpha' : E'(F) \longrightarrow F^*/F^{*2}$$

send a point with nonzero x to the class of x, send the origin to 1, and send $(0, 0)$ to the class of the linear coefficient in the equation. Their kernels are $\widehat{\varphi}E'(F)$ and $\varphi E(F)$, respectively. In the present case

$$|\operatorname{im}\alpha|\,|\operatorname{im}\alpha'| = 4\,2^{\operatorname{rank} E(F)}. \tag{5.1}$$

For completeness, E has full rational two-torsion, and either point $(\pm d, 0)$ maps under φ to the nonidentity point of $\ker\widehat{\varphi}$. Hence the kernel of $E'(F)/\varphi E(F) \to \widehat{\varphi}E'(F)/2E(F)$ is trivial. Taking indices in $2E(F) \subset \widehat{\varphi}E'(F) \subset E(F)$ gives the left side of (5.1) as $|E(F)/2E(F)|$, which is $2^{\operatorname{rank} E(F)+2}$ by the Mordell–Weil Theorem.

If the linear coefficient of an equation $y^2 = x^3 + cx$ is a unit at a finite place, then a nonzero x has even valuation there. Indeed, when $v(x) < 0$ the term x^3 has strictly smaller valuation than cx, while when $v(x) > 0$ the term cx has strictly smaller valuation; an odd $v(x)$ would give an odd valuation of y^2 in either case. Applied to E, this shows that every class in $\operatorname{im}\alpha$ is supported, apart from its unit class, only at $\mathfrak{d}$. The class number calculation therefore bounds the image by $4 \cdot 2 = 8$. The torsion points already give the four distinct classes

$$1, -1, d, -d.$$

Let $w = -5 + 4\sqrt{2} = d(1+\sqrt{2})^3$. Direct calculation gives $w^3 - w = -8d$, and consequently

$$Q = (-dw,\, 2\sqrt{2}\, d^2) \in E(F).$$

The element $-dw$ has opposite signs at the two real embeddings, whereas the four torsion classes have equal signs at both embeddings. Thus $\alpha(Q)$ is outside those four classes, proving $|\operatorname{im}\alpha| = 8$.

We next prove $\operatorname{im}\alpha' = \{1\}$. At both real embeddings, $x^3 + 4d^2x$ has the sign of x, so every nonzero x on E' is totally positive. The unit-coefficient argument excludes odd valuations away from the prime above 2 and $\mathfrak{d}$. Normalize the valuation at the prime above 2 to have value group $\mathbb{Z}$. Then $v(4d^2) = 4$. If $v(x)$ were odd, the valuations $3v(x)$ and $4 + v(x)$ would be unequal and their minimum would be odd, again impossible for y^2. At $\mathfrak{d}$ the coefficient has valuation 2. An odd $v(x)$ is impossible unless $v(x) = 1$, by the same comparison. In that remaining case write $x = du$ with u a local unit. The equation becomes

$$y^2 = d^3u(u^2 + 4).$$

An even valuation would require $u^2 + 4 \equiv 0 \pmod{\mathfrak{d}}$. But $-4 = 3$ is not a square in $\mathbb{F}_7$. Thus all finite valuations of x are even. Since $\mathcal{O}_F$ is principal, x is a square times a unit; total positivity makes that unit a square. The exceptional point $(0, 0)$ also gives the trivial class $4d^2$, and the origin does likewise. This proves the assertion about α'. Equation (5.1) now gives $\operatorname{rank} E(F) = 1$.

Finally, choose m divisible by the exponent of the finite torsion subgroup of $E(F)$. The finitely generated group $mE(F)$ is torsion-free of rank one and therefore has a generator P as claimed. □

Recall the parameter $z = -x/y$, with $z(O) = 0$, and the formal logarithm and its degree-K truncation

$$\ell(Z) = Z + \sum_{t \geq 2} c_t Z^t, \qquad \ell_K(Z) = Z + \sum_{2 \leq t \leq K} c_t Z^t.$$

Reading multiplication indices from elliptic coordinates has a precedent in Poonen's rank-one Diophantine construction. His Lemma 9 reads divisibility of indices from denominators of elliptic multiples, and Lemma 11 uses the formal group to approximate an integer square by a quotient of coordinates [52]. Proposition 2.2 uses a truncated formal logarithm to approximate a quotient of indices. The paired local tests in the reduction use this comparison together with the conjugate curve.

The effective computation of the truncations in Section 2.1 can be made explicit. With $w = -1/y$ one has

$$x = z/w, \qquad w = z^3 - d^2zw^2.$$

Successively solving this identity determines w as a formal power series in z; substitution in $dx/(2y)$ and termwise integration determines ℓ. The formal group and its invariant differential have integral coefficients at every place of good integral reduction [67, Chapter IV].

Proof of Proposition 2.2. Include 2, the primes of bad reduction of either equation, and the ramified primes of $F/\mathbb{Q}$ in S_E. Let $p \notin S_E$, let v be a place above p normalized by $v(p) = 1$, and let $n, n' \in \mathbb{Z}$, with $n' \neq 0$ and both $nP, n'P$ reducing to O, as in the proposition. At this allowed place $F_v/\mathbb{Q}_p$ is unramified, so its normalized value group is $\mathbb{Z}$. The kernel of reduction is the formal group on the maximal ideal. For a finite point on this integral good model, membership in that kernel is equivalent to $v(x) < 0$; its parameter then has $v(z) > 0$.

Write the invariant differential as $(1 + \sum_{j\geq 1} a_j Z^j)dZ$, with a_j integral at v. Thus $v(c_t) \geq -v_p(t)$. If $v(Z) \geq 1$, the term of degree $t \geq 2$ in the logarithm has valuation at least $t\,v(Z) - v_p(t) > v(Z)$ for $p > 2$, and these valuations tend to infinity. The logarithm consequently converges on the entire formal kernel, is a homomorphism there, and satisfies

$$v(\ell(Z)) = v(Z).$$

These are also the standard formal-logarithm assertions of [67, Chapter IV, §6]; the displayed coefficient bound specifies the uniformity we require.

Set $U = nP$ and $V = n'P$. The identity $[n']U = [n]V$, entirely within the formal kernel, gives $n'\ell(z(U)) = n\ell(z(V))$. The point V is nonzero, hence $\ell(z(V)) \neq 0$. Valuation preservation proves (2.6), including $U = O$.

Now fix $K \geq 1$ and suppose $p > K + 1$ and $v_p(n/n') \geq 0$. For every $t > K$,

$$t - 1 - v_p(t) \geq K. \tag{5.2}$$

For $t < p$ this is immediate. For $t \geq p$, if $j = v_p(t) > 0$ then $t-1-j \geq p^j-1-j \geq p-2 \geq K$; if $j = 0$ it is again immediate. Therefore, for $v(Z) \geq 1$,

$$v(\ell(Z) - \ell_K(Z)) \geq v(Z) + K. \tag{5.3}$$

The inequality also holds at $Z = 0$ with the stated convention. Put $q = n/n'$. The hypothesis $v_p(q) \geq 0$ and (2.6) give $v(z(U)) \geq v(z(V))$. The identity $\ell(z(U)) = q\ell(z(V))$ and (5.3) imply

$$v\big(\ell_K(z(U)) - q\ell_K(z(V))\big) \geq v(z(V)) + K.$$

The same tail estimate gives $v(\ell_K(z(V))) = v(z(V))$, in particular a nonzero denominator. Division proves (2.7). The proof applies verbatim to the conjugate equation and the polynomial ℓ_K^σ. □

The quantifiers in Proposition 2.2 are those used in Lemma 2.3: for each fixed K, a single finite set of excluded rational primes works for all integer indices and all unramified places above the remaining primes. The transfer there concerns integer multiples of the fixed F-rational points P, P^σ and finite polynomial evaluations. No infinite series is evaluated in the elementary extension.

6 Prime patterns for the two conjugate local tests

This section proves Corollary 2.6, the ordinary coverage input used in the recursive theory. The raw prime-pattern result below supplies the reduction orders; its first consequence constructs the paired integer witnesses by the Chinese Remainder Theorem.

Retain the five points $\mathcal{C}$ and the primitive integral forms L_b of Theorem 2.4, choosing $L_\infty(U, V) = V$. For integers A, u_1, u_2, v with $v \neq 0$, setting $u_3 = Av - u_1 - u_2$ gives

$$A = \frac{u_1}{v} + \frac{u_2}{v} + \frac{u_3}{v}.$$

Evaluating the five contact forms on (u_i, v) gives the common value v and four further values for each slope. We seek to express each of these thirteen entries as the product of a sign, a multiplier from one fixed finite set, and a positive prime at which the reduction orders of E and E^σ permit simultaneous local tests.

Lemma 6.1 (Prime patterns)**.** *There is a finite set $\mathcal{K} \subset \mathbb{Z}_{>0}$ such that, for every $A \in \mathbb{Z}$ and every real $B > 0$, there exist integers u_1, u_2, v, with $v \neq 0$, for which, on setting $u_3 = Av - u_1 - u_2$, each of the thirteen numbers*

$$v, \qquad L_b(u_j, v) \quad (1 \le j \le 3,\ b \in \mathcal{C} \setminus \{\infty\}) \tag{6.1}$$

has the form $\epsilon k r$, where $\epsilon \in \{1, -1\}$, $k \in \mathcal{K}$, and $r > B$ is a positive rational prime. Every such prime splits in F and is a prime of good reduction for both E and E^σ. At either place of F above r, if d_1, d_2 are the cardinalities of the reduction groups of E, E^σ over $\mathbb{F}_r$, then

$$\gcd(d_1, d_2) \mid 4, \qquad r \nmid d_1 d_2. \tag{6.2}$$

The set $\mathcal{K}$ is independent of A and B.

The listed values are nonzero, so the slopes $(u_i : v)$ lie outside $\mathcal{C}$. A prime counted by one of their contact radicals divides a listed contact value, because the corresponding normalized contact depth is positive and the coordinates are integral; unless it divides a member of $\mathcal{K}$, it is therefore one of the selected primes r.

Before proving the lemma, we record the consequence that explains its splitting and reduction-order conditions. Together with Proposition 2.2, they give two local tests at each place, with one pair of auxiliary integers at that place serving every integer a.

Corollary 6.2 (Simultaneous integer tests)**.** *Fix an integer $K \ge 1$ and a prime r supplied by Lemma 6.1, large enough for Proposition 2.2 at precision K. At each of the two places v above r there are nonzero integers h, j_0, chosen independently of $a \in \mathbb{Z}$, such that, for every $a \in \mathbb{Z}$:*

1. *hP and haP reduce to O at v, $\ell_K(z(hP)) \neq 0$, and*
$$v\left(\frac{\ell_K(z(haP))}{\ell_K(z(hP))} - a\right) \ge K;$$
2. *$j_0 P^\sigma$ and $(h-4)aP^\sigma$ reduce to O at the same place v, and*
$$v\left(\frac{z((h-4)aP^\sigma)}{z(j_0 P^\sigma)}\right) \ge K.$$

The choices of h, j_0 may differ at the two places.

Proof. Let d_1, d_2 be the two reduction orders at the chosen place. Choose $j_0 = d_2$. Since $r \nmid d_1$ and $\gcd(d_1, d_2) \mid 4$, the pair of integer congruences

$$h \equiv 0 \pmod{d_1}, \qquad h \equiv 4 \pmod{d_2 r^K}$$

is consistent. Choose a positive solution h. Reduction orders show that all four indicated points are in the respective formal kernels, uniformly for $a \in \mathbb{Z}$. The logarithm congruence follows from Proposition 2.2 with indices ha, h, whose ratio is a. For the conjugate ratio use (2.6) on the conjugate curve:

$$v\left(\frac{z((h-4)aP^\sigma)}{z(j_0 P^\sigma)}\right) = v_r\left(\frac{(h-4)a}{d_2}\right) \ge K.$$

The nonzero denominator follows because $j_0 P^\sigma$ is nontorsion. Zero numerator indices are covered by the infinite-valuation convention. Finally, at these split unramified primes the valuation ring on F is the localization of $\mathcal{O}_F = \mathbb{Z}[\sqrt{2}]$ at the chosen prime ideal, and r is a uniformizer. Thus every displayed inequality is exactly membership in $r^K(\mathcal{O}_F)_\mathfrak{p}$. For a finite formal point, the condition $v(x) < 0$ equivalently places $1/x$ in $r(\mathcal{O}_F)_\mathfrak{p}$, as used in the ring-language tests of Section 2.3. □

Proof of Corollary 2.6. Fix $K \geq 1$, $A \in \mathbb{Z}$, and $Y_0 > 0$. Choose a real threshold larger than Y_0, $K+1$, and every prime in the fixed finite set S_E. Lemma 6.1 gives the thirteen contact values with primes above this threshold, using its one finite set $\mathcal{K}$. These primes split in F and are good for both curves. At each of the two places above each occurrence, Corollary 6.2 gives nonzero h, j_0, fixed for every integer a, with its two formal-kernel conditions and valuation comparisons. In the identity embedding with $T_a = aP$, these are exactly conditions (i) and (ii) of Lemma 2.3. The final localization calculation in the proof of Corollary 6.2 gives the stated membership and finite-point interpretations. The choices are made separately at the different occurrences and places, as allowed. The fixed set $\mathcal{K}$ is independent of K, A, Y_0 because the prime-pattern lemma makes it independent of A and of the threshold. □

We now prove Lemma 6.1. Constants with a subscript A may depend on A, a subsequently fixed real box, and the fixed contact forms, but not on the small-prime cutoff L. The order of choices is: the finite multiplier set; then A, a box and one initial congruence class; then L; and finally a large dilation parameter X.

6.1 Reduction orders and their possible common factors

Restrict attention to primes

$$r \equiv 1 \pmod 8, \qquad r \equiv 3 \pmod 7. \tag{6.3}$$

Quadratic reciprocity gives $(2/r) = 1$, so r splits in F, and $(7/r) = (r/7) = -1$. At either chosen place above r, the two nonzero reductions of d and d^σ have product 7. They therefore represent opposite square classes in $\mathbb{F}_r^*$. The curves E and E^σ are the quadratic twists by these classes of $E_0 : y^2 = x^3 - x$, so their traces of Frobenius are opposite.

Here is the specific description of that trace which we need. Since $r \equiv 1 \pmod 4$, the automorphism $(x, y) \mapsto (-x, iy)$ of E_0 is defined over $\mathbb{F}_r$ and has square $[-1]$. Its centralizer in the rational geometric endomorphism algebra is $\mathbb{Q}(i)$. Indeed, the faithful action of endomorphisms on an auxiliary two-dimensional rational Tate module bounds this centralizer by dimension two, while it already contains $\mathbb{Q}(i)$. More explicitly one may use the 3-adic Tate module for $r > 3$, on which the minimal polynomial $T^2 + 1$ is irreducible and its matrix centralizer has dimension two. The faithful action remains injective after extending coefficients [67, Theorem III.7.4].

Frobenius commutes with this automorphism. Its characteristic polynomial is integral, and its determinant is r; these are the usual degree and trace identities for elliptic endomorphisms [67, Proposition III.8.6 and Chapter V, §2]. It is thus an algebraic integer in $\mathbb{Q}(i)$, hence equals $a + bi$ for $a, b \in \mathbb{Z}$, with

$$a^2 + b^2 = r, \qquad a \neq 0. \tag{6.4}$$

The last condition follows because a rational prime is not a square. Consequently, in either order,

$$d_1, d_2 = r + 1 \pm 2a. \tag{6.5}$$

Both orders are divisible by four, since both curves have full rational two-torsion over $\mathbb{F}_r$. Their sum is $2(r+1)$, whose 2-adic valuation is exactly two by (6.3). Thus the 2-part of their greatest common divisor is exactly four. For sufficiently large r, the Hasse bound gives $0 < d_i < 2r$. Since $4 \mid d_i$ and r is odd, $d_i \neq r$, proving $r \nmid d_i$.

If an odd prime l divides both orders, their sum and difference show that $l \mid r+1$ and $l \mid a$. Equation (6.4) then gives

$$a \equiv 0 \pmod{l}, \qquad b^2 \equiv -1 \pmod{l}. \tag{6.6}$$

In particular such an l is 1 modulo 4. We shall first exclude small l by congruences on $r+1$, then use an upper sieve to remove the remaining primes r satisfying these two local conditions.

6.2 A finite multiplier set and the initial residue conditions

Write $L_b(U,V) = \alpha_b U + \gamma_b V$ for $b \neq \infty$. Here $\alpha_b \neq 0$, $\gcd(\alpha_b, \gamma_b) = 1$, and $\alpha_b\gamma_c - \alpha_c\gamma_b \neq 0$ for $b \neq c$. For an integer A, put $T = (u_1, u_2, v)$ and denote the thirteen homogeneous forms in (6.1) by $F_\nu(T)$, $1 \leq \nu \leq t = 13$. Their coefficient vectors, other than $(0,0,1)$, are

$$(\alpha_b, 0, \gamma_b), \quad (0, \alpha_b, \gamma_b), \quad (-\alpha_b, -\alpha_b, \alpha_b A + \gamma_b).$$

Every vector is primitive, including the last one because $\gcd(\alpha_b, \alpha_b A + \gamma_b) = 1$. There is a fixed threshold p_0, independent of A, above which any two vectors are linearly independent over $\mathbb{F}_p$. Within one of the three families this follows from the fixed nonzero determinants of the L_b; between different families it follows from the three distinct directions $(1,0), (0,1), (-1,-1)$ in the first two coordinates. The form v is independent of each remaining form whenever $p \nmid \alpha_b$.

Choose once and for all a finite set S_0 consisting of all primes up to a sufficiently large threshold, including $2, 7$, the exceptions just described, and all fixed bad primes needed for the elliptic curves. Increase the threshold so that $p > 2t$ and the later local-density estimates hold for every $p \notin S_0$. Choose a fixed integer H_0 with $t2^{-H_0} < 1$. For any A and any $p \in S_0$, each primitive F_ν maps $(\mathbb{Z}/p^{H_0}\mathbb{Z})^3$ onto $\mathbb{Z}/p^{H_0}\mathbb{Z}$. The proportion of vectors for which $p^{H_0} \mid F_\nu(T)$ is therefore p^{-H_0}. Since $tp^{-H_0} < 1$, there is a vector class modulo p^{H_0} on which none of the F_ν is divisible by p^{H_0}. Every lift of this class modulo p^{H_0+3} has

$$v_p(F_\nu(T)) < H_0 \qquad (1 \leq \nu \leq t).$$

This proves the existence of the small-prime classes that we will select for a particular A.

We next choose a finite catalogue of multipliers, without selecting A or any of those classes. Put $m_2 = 8$ and $m_p = p$ for odd $p \in S_0$. Define units $c_2 = 1 \pmod 8$, $c_7 = 3 \pmod 7$, and $c_p = 1 \pmod p$ for the remaining $p \in S_0$. These are the desired residues of the positive primes in Lemma 6.1. A catalogue datum consists of integers and unit residues

$$e_{p\nu} \in \{0, \ldots, H_0 - 1\}, \qquad \xi_{p\nu} \in (\mathbb{Z}/m_p\mathbb{Z})^* \quad (p \in S_0,\ 1 \leq \nu \leq t).$$

For such a datum set $k_{0\nu} = \prod_{p \in S_0} p^{e_{p\nu}}$. The Chinese Remainder Theorem and Dirichlet's Theorem [12, §4] allow us to choose distinct primes $h_1, \ldots, h_t \notin S_0$, hence above the fixed threshold, with

$$h_\nu \equiv c_p^{-1}\xi_{p\nu} \pmod{m_p} \qquad (p \in S_0). \tag{6.7}$$

When the datum later comes from a chosen vector class, $\xi_{p\nu}$ will be the signed unit remaining after the factors from S_0 have been removed. The congruence gives $\xi_{p\nu}h_\nu^{-1} \equiv c_p \pmod{m_p}$, so division by the additional factor h_ν gives the desired residue. There are only finitely

many catalogue data. Fix one such tuple of primes for each datum and let $\mathcal{K}$ be the union of the resulting integers $k_{0\nu}h_\nu$. Thus $\mathcal{K}$ is finite and is fixed before A, the cutoff L, and the dilation X.

Now fix the integer A for which the lemma is to be proved. Choose a bounded open rectangular box $\mathcal{B}$ of positive volume whose closure avoids the zero planes of all F_ν, and let $\epsilon_\nu \in \{1,-1\}$ be the sign of F_ν on this box. For each $p \in S_0$, choose one of the vector classes modulo p^{H_0+3} supplied above and put $e_{p\nu} = v_p(F_\nu(T))$ on that class. These values do not depend on its representative. They define $k_{0\nu}$ as above. The class also determines the unit

$$\xi_{p\nu} = \epsilon_\nu \frac{F_\nu(T)}{p^{e_{p\nu}}} \left(\frac{k_{0\nu}}{p^{e_{p\nu}}}\right)^{-1} \pmod{m_p}.$$

Here the second factor is a unit modulo m_p, and the first quotient is known modulo $p^{H_0+3-e_{p\nu}}$, a modulus containing at least four powers of p. In particular it determines the required residue even when $p = 2$ and $m_p = 8$. Select the tuple h_ν already fixed for this datum, and put $k_\nu = k_{0\nu}h_\nu$.

At $p = h_\nu$, impose $F_\nu(T) = 0 \pmod p$. On this plane, each other F_j is nonconstant. Each condition $\epsilon_j F_j(T) = 0$ or $-k_j$ removes at most one line, and at most $2(t-1)p < p^2$ points are removed. There is therefore a point satisfying

$$\epsilon_j F_j(T) \neq 0, -k_j \pmod{h_\nu} \quad (j \neq \nu).$$

Lift that point modulo p^2 so that

$$\epsilon_\nu F_\nu(T)/p \neq 0, -k_{0\nu} \pmod p.$$

Such a lift exists because F_ν is primitive, so varying the lift changes $F_\nu(T)/p$ through all residues modulo p. The Chinese Remainder Theorem combines these choices with the chosen classes at S_0 into one vector class modulo

$$W_0 = \prod_{p \in S_0} p^{H_0+3} \prod_{\nu=1}^{t} h_\nu^2. \tag{6.8}$$

On this class all

$$\Psi_\nu(T) = \epsilon_\nu F_\nu(T)/k_\nu$$

are integers and are nonzero modulo every prime dividing W_0. At a prime in S_0, the definition of $\xi_{p\nu}$ and (6.7) give $\Psi_\nu \equiv c_p \pmod{m_p}$. The conditions at each h_j give the remaining exclusions. Thus the divided forms equal 1 modulo 8 and 3 modulo 7, and avoid -1 modulo every odd prime dividing W_0. Notice also that W_0/k_ν is divisible by every prime dividing W_0. This extra divisibility will control the local factors of the prime-counting theorem.

Fix a smaller rectangular box $\mathcal{B}'$ of positive volume with $\overline{\mathcal{B}'} \subset \mathcal{B}$. Since the signs agree with the forms on $\mathcal{B}$ and its closure avoids their zero planes, there are constants $0 < c_A < C_A$ such that

$$c_A X \leq \Psi_\nu(T) \leq C_A X \quad (T \in X\mathcal{B},\ 1 \leq \nu \leq t). \tag{6.9}$$

All choices for this A, including W_0, are now fixed before choosing L. The catalogue was finite, so W_0 ranges in a fixed finite set as A varies.

6.3 A lower bound for the prime patterns

Let L be a fixed cutoff larger than every prime dividing W_0, and put

$$W = W_0 \prod_{\substack{p \le L \\ p \nmid W_0}} p.$$

At each new prime $p \le L$ allow all residue classes of T satisfying

$$\epsilon_\nu F_\nu(T) \neq 0, -k_\nu \pmod p \quad (1 \le \nu \le t). \tag{6.10}$$

At most $2t$ affine planes, each of p^2 points, are excluded. Thus if $\mathcal{R}_W$ is the set of all allowed vector classes modulo W, including the one fixed class modulo W_0, then

$$\frac{|\mathcal{R}_W|}{W^3} \ge W_0^{-3} \prod_{\substack{p \le L \\ p \nmid W_0}} (1 - 2t/p) \ge c_0 (\log L)^{-C_0}. \tag{6.11}$$

The last bound follows by taking logarithms and using Mertens' estimate $\sum_{p \le L} p^{-1} = \log\log L + O(1)$ [73, equation (1.1) and Theorem 1]. The constants can be fixed independently of L; since W_0 ranges in a fixed finite set they could also be made independent of A. Whenever $\Psi_\nu(T) = r$ is prime, these conditions and the initial conditions imply that no odd $l \le L$ divides $r + 1$. Such an l cannot divide both orders in (6.5).

Let $N_{A,L}(X)$ be the number of $T \in X\mathcal{B} \cap \mathbb{Z}^3$ for which every $\Psi_\nu(T)$ is prime and all imposed congruences hold. We claim that

$$N_{A,L}(X) \ge c_A (\log L)^{-C_0} \frac{X^3}{(\log X)^t} \tag{6.12}$$

for sufficiently large X depending on A, L, with $c_A > 0$ independent of L.

We use the following precise form of the Green–Tao–Ziegler theorem. For a fixed finite collection of integer affine-linear forms $\psi_1, \dots, \psi_t$ on $\mathbb{Z}^d$ whose nonzero linear parts are pairwise nonproportional, and a fixed dilating box where all values are positive, the von Mangoldt weighted count is

$$\sum_{n \in \mathcal{D}_X \cap \mathbb{Z}^d} \prod_{\nu=1}^{t} \Lambda(\psi_\nu(n)) = \mathrm{vol}(\mathcal{D}_X) \prod_p \beta_p + o(X^d), \tag{6.13}$$

where

$$\beta_p = (1 - 1/p)^{-t} p^{-d} \#\{n \in \mathbb{F}_p^d : \psi_\nu(n) \neq 0 \text{ for every } \nu\}.$$

Here all coefficients and the box are fixed before $X \to \infty$. This is the finite-complexity case of the Main Theorem of [28], combined with the Möbius–nilsequence theorem [25, Theorem 1.1] and the inverse theorem [26, Theorem 1.3], with its correction [27]. Pairwise nonproportional linear parts give finite complexity: for any chosen form, the other forms can be partitioned into singletons, none spanning it. The 2024 erratum corrects intermediate factorization and filtration statements and leaves the stated inverse theorem, and hence (6.13), unchanged.

For each $T_0 \in \mathcal{R}_W$, use its representative in $[0, W)^3$ and write $T = T_0 + Wn$. The forms

$$\psi_{\nu,T_0}(n) = \epsilon_\nu F_\nu(T_0 + Wn)/k_\nu$$

have integer coefficients and pairwise nonproportional linear parts. Use $n \in (X/W)\mathcal{B}'$. For all sufficiently large X, depending on A, L, this entails $T \in X\mathcal{B}$, uniformly over the finite set of representatives T_0. The volume in (6.13) is then $\mathrm{vol}(\mathcal{B}')X^3/W^3$.

At $p \mid W$, the coefficients of the variable parts of every divided form vanish modulo p, whereas their constant values are nonzero. Consequently $\beta_p = (1-1/p)^{-t} \geq 1$. At $p \nmid W$, multiplication by W and by k_ν^{-1} is invertible. The zero sets of the forms are distinct affine hyperplanes, and each pair has an intersection of codimension two. Inclusion–exclusion, using only the first two terms for a uniform error, gives

$$p^{-3}\#\{n : \psi_{\nu,T_0}(n) \neq 0 \ (\forall \nu)\} = 1 - t/p + O_t(p^{-2}).$$

Thus $\beta_p = 1 + O_t(p^{-2})$. Taking the original fixed threshold large enough, the product over any subset of these primes is bounded below by one fixed positive convergent product. This lower bound is independent of L and T_0.

Prime powers in (6.13) are negligible. On a box of scale X, there are $O(X^{1/2}\log X)$ proper prime powers in the range of each nonconstant form, and at most $O_{A,L}(X^2)$ lattice points for each prescribed value of that form, by solving for one coordinate. Their total weighted contribution is therefore $O_{A,L}(X^{5/2}(\log X)^{t+1}) = o_{A,L}(X^3)$. For genuine prime values the weight is at most $C_A(\log X)^t$, by (6.9).

Sum (6.13) over the finitely many allowed progressions. For each fixed L their errors still sum to $o_{A,L}(X^3)$, while the W^{-3} volume factor cancels the W^3 in the number of progressions in (6.11). The lower bound for the local-factor product, the removal of prime powers, and the bound on the prime weights therefore give (6.12), with a positive coefficient independent of L. We require no uniformity in the little-oh error as A or L varies.

We must now remove the patterns for which some pair of reduction orders has an odd common factor $l > L$. For one form, we will bound the number of possible bad prime values by $O_A(X/(L\log X))$ plus terms of smaller order as $X \to \infty$ for fixed L. For each fixed bad value, a second count will bound the points on its affine fiber for which the other $t-1$ forms remain prime by $O_A(X^2/(\log X)^{t-1})$. The contribution from the term $X/(L\log X)$ is therefore $O_A(X^3/(L(\log X)^t))$. This retains the factor $(\log X)^{-t}$ in (6.12) and introduces the saving $1/L$. The first count sieves the pairs (a,b) in (6.4); the second sieves the two free coordinates on a fiber. Both use the following two-dimensional estimate.

6.4 An elementary upper sieve with uniform errors

Let $\mathcal{P}_Z$ be a set of primes at most Z, with forbidden subsets $\Omega_p \subset \mathbb{F}_p^2$ having densities $g(p) = |\Omega_p|/p^2$ in $(0, 1/2)$. Extend g multiplicatively to squarefree integers supported on $\mathcal{P}_Z$. Let $\mathbf{x}$ range over a square in $\mathbb{Z}^2$ with n integer choices in each coordinate. The simultaneous forbidden conditions at the primes dividing a squarefree integer e hold at

$$n^2 g(e) + O(ne + e^2) \tag{6.14}$$

points, uniformly in the location of the square. Indeed the Chinese Remainder Theorem gives $e^2 g(e)$ residue classes, and the count in each differs from n^2/e^2 by $O(n/e+1)$.

The following estimate uses Selberg's λ^2 method [63]; compare [44, Section 3.2]. We give the weights and error term needed for these two-dimensional boxes.

Lemma 6.3. *The number of points in this square avoiding every Ω_p is at most*

$$\frac{n^2}{G(Z)} + O(nZ^6 + Z^8), \qquad G(Z) = \sum_{\substack{e \leq Z \\ p\mid e \Rightarrow p \in \mathcal{P}_Z}} \ \prod_{\substack{squarefree \\ p \mid e}} \frac{g(p)}{1-g(p)}. \tag{6.15}$$

The error constant is absolute.

Proof. All integers in the sums that follow are squarefree and supported on $\mathcal{P}_Z$. Put $g_*(p) = (1-g(p))/g(p)$ and extend it multiplicatively, with $g(1) = g_*(1) = 1$. For real weights λ_e supported on $e \le Z$ with $\lambda_1 = 1$, the square

$$\left(\sum_{\substack{e \le Z \\ \mathbf{x} \bmod p \in \Omega_p \ (p|e)}} \lambda_e \right)^2$$

majorizes the indicator of avoiding all forbidden sets: at such a point only $e = 1$ contributes. Upon summing and applying (6.14), its main term divided by n^2 is $\sum_{e,f} \lambda_e \lambda_f g([e,f])$. The identity $1/g(d) = \sum_{h|d} g_*(h)$ gives the diagonalization

$$\sum_{e,f} \lambda_e \lambda_f g([e,f]) = \sum_h g_*(h) \left(\sum_{\substack{e \le Z \\ h|e}} \lambda_e g(e) \right)^2. \tag{6.16}$$

Prescribe the inner sums to equal $\mu(h)/(G(Z)g_*(h))$. Finite Möbius inversion gives

$$\lambda_e = \frac{\mu(e)}{G(Z)g(e)g_*(e)} \sum_{\substack{k \le Z/e \\ (k,e)=1}} \frac{\mu(k)^2}{g_*(k)}. \tag{6.17}$$

In particular $\lambda_1 = 1$ and

$$|\lambda_e| \le \frac{1}{g(e)g_*(e)} = \prod_{p|e} (1-g(p))^{-1} \le 2^{\omega(e)} \le e.$$

Substitution in (6.16) gives $1/G(Z)$. For the error at $[e,f]$, use $[e,f] \le ef$ and $|\lambda_e \lambda_f| \le ef$. Its sum is bounded by

$$O\left(n \sum_{e,f \le Z} e^2 f^2 + \sum_{e,f \le Z} e^3 f^3 \right) = O(nZ^6 + Z^8),$$

as required. □

We shall always set $Z = X^{1/100}$. Two consequences for $G(Z)$ will be useful. If $\mathcal{P}_Z$ contains all primes above a fixed threshold up to Z, and $g(p) = s/p + O_s(p^{-2})$ for a fixed positive integer s, then

$$G(Z) \gg (\log X)^s. \tag{6.18}$$

If instead it contains the primes $p \equiv 1 \pmod 4$ above a fixed threshold up to Z, with at most one additional prime omitted, and $g(p) = 2/p + O(p^{-2})$, then

$$G(Z) \gg \log X, \tag{6.19}$$

uniformly in that omitted prime. In these two estimates the implicit constants may depend on the fixed threshold and the fixed density estimates, as well as on s, but not on X or the omitted prime.

Here are details of the truncation behind both claims. For $y = Z^\theta$, where $\theta > 0$ will be fixed sufficiently small, the Euler product

$$D(y) = \prod_{\substack{p \le y \\ p \in \mathcal{P}_Z}} (1 + g_*(p)^{-1})$$

is respectively $\gg (\log y)^s$ or $\gg \log y$, by the ordinary and progression Mertens estimates [73, equation (1.1) and Theorem 1]. In the latter case removing one prime loses at most a factor two because $g(p) < 1/2$. In the expansion of $D(y)$ give a squarefree e weight $1/g_*(e)$. The weighted mean of $\log e$ equals

$$\sum_{\substack{p\le y\\ p\in\mathcal{P}_Z}} g(p)\log p \ll \log y.$$

This uses $\sum_{p\le y}(\log p)/p \ll \log y$, and the bound is uniform when a prime is omitted. For a small enough fixed θ, Markov's inequality shows that at least half the weight has $e \le Z$. Those terms occur in $G(Z)$, proving both assertions.

6.5 Counting bad prime values

Fix one index ν. All its possible prime values lie in the interval (6.9). Let $B_{A,L}(X)$ be the number of primes in that interval satisfying (6.3) for which the two orders have a common odd factor $l > L$. We claim

$$B_{A,L}(X) \ll_A \frac{X}{L\log X} + X^{7/8} + \sqrt{X}\log X. \tag{6.20}$$

It suffices to count pairs $(a,b) \in \mathbb{Z}^2$ with $|a|, |b| \le C_A\sqrt{X}$, $a \ne 0$, for which a^2+b^2 is a prime in the interval and (6.6) holds for some prime $l > L$. Every bad prime supplies such a pair. Overcounting pairs, and counting a pair several times if necessary, gives an upper bound.

First consider $L < l \le X^{1/8}$. If -1 has no square root modulo l, there are no pairs. Otherwise take each of its two roots b_0 and write

$$a = lU, \qquad b = b_0 + lV.$$

The relevant U, V lie in an encompassing square with

$$n \asymp_A \sqrt{X}/l \gg_A X^{3/8} \tag{6.21}$$

integer choices in each coordinate, uniformly for these l and for least nonnegative representatives b_0. Sieve by all sufficiently large primes $p \equiv 1 \pmod 4$, $p \ne l$, up to $Z = X^{1/100}$. Modulo p the substitution $(U,V) \mapsto (lU, b_0 + lV)$ is an affine bijection, so the forbidden equation $a^2+b^2 = 0$ has exactly $2p-1$ solutions: it is the union of two distinct lines meeting at one point. Thus

$$g(p) = 2/p - 1/p^2,$$

independently of l, b_0. Since the prime value a^2+b^2 is $\asymp_A X > Z$, it avoids all these forbidden sets. Lemma 6.3 and (6.19) give $O(n^2/\log X + nZ^6 + Z^8)$ pairs. At the shortest possible squares,

$$\frac{nZ^6+Z^8}{n^2/\log X} \ll_A (\log X)(X^{6/100-3/8} + X^{8/100-3/4}) = o_A(1).$$

The error is therefore absorbed uniformly before summing over l. Each l contributes $O_A(X/(l^2\log X))$, and summing over even all integers $l > L$ gives $O_A(X/(L\log X))$.

For $l > X^{1/8}$ use a lattice count without a sieve. Since $a \ne 0$ and $l \mid a$, necessarily $l \le C_A\sqrt{X}$, and the number of possible a is $O_A(\sqrt{X}/l)$, with no additive constant. There are at most two classes for b, giving $O_A(\sqrt{X}/l + 1)$ possibilities. The count for one l is thus

$$O_A(X/l^2 + \sqrt{X}/l).$$

To sum it explicitly, split $X^{1/8} < l \le C_A\sqrt{X}$ into dyadic blocks $D < l \le 2D$. Even if every integer in a block were allowed, its contribution would be $O_A(X/D + \sqrt{X})$. The first terms form a geometric sum $O_A(X^{7/8})$, while there are $O_A(\log X)$ blocks. Their total is $O_A(X^{7/8} + \sqrt{X}\log X)$. This proves (6.20) with constants independent of L.

6.6 A uniform sieve on each affine fiber

For a fixed prime value $r \asymp_A X$ of the νth form, consider

$$F_\nu(T) = \epsilon_\nu k_\nu r. \tag{6.22}$$

We claim that the number of $T \in X\mathcal{B} \cap \mathbb{Z}^3$ on this fiber for which all other $\Psi_j(T)$ are primes is

$$\ll_A \frac{X^2}{(\log X)^{t-1}}, \tag{6.23}$$

uniformly in r and independently of L. In this upper bound we discard all the congruence conditions used for the lower bound.

Since F_ν is primitive, an integral unimodular coordinate change makes it the first coordinate. The other two coordinates of every point of the fiber inside $X\mathcal{B}$ lie in an encompassing square with $n \asymp_A X$ integer choices in each direction. The coordinate change depends on A, ν but not on r, L. Sieve by all primes up to $Z = X^{1/100}$ above a fixed sufficiently large threshold, excluding every prime dividing any multiplier. For all sufficiently large X, every sieving prime differs from r, so the fiber constant $c = \epsilon_\nu k_\nu r$ is nonzero modulo each such prime p.

On the affine plane $F_\nu = c$ over $\mathbb{F}_p$, every other F_j restricts to a nonconstant affine function, because F_j is not proportional to F_ν. Furthermore the $t-1$ zero lines of these restrictions are distinct. If two coincided, their affine functions would be scalar multiples on that plane. For some $\lambda \in \mathbb{F}_p^*$ and $\mu \in \mathbb{F}_p$ this would give the identity

$$F_j - \lambda F_k = \mu(F_\nu - c).$$

The left side is homogeneous, so comparison of constant terms gives $\mu c = 0$. Since $c \neq 0$, one has $\mu = 0$, contradicting the pairwise independence of F_j, F_k. This argument also covers the case that three ambient forms are jointly dependent: a nonzero constant shift then gives distinct parallel lines, rather than coincident ones.

The union of the $t-1$ distinct zero lines has $(t-1)p + O_t(1)$ points. All must be avoided, since the other prime values have size $\asymp_A X$ and exceed the sieving primes; the multiplier is invertible modulo every sieving prime. Thus $g(p) = (t-1)/p + O_t(p^{-2})$, uniformly in the fiber. Lemma 6.3 and (6.18) give

$$O_A\left(\frac{X^2}{(\log X)^{t-1}} + XZ^6 + Z^8\right) = O_A\left(\frac{X^2}{(\log X)^{t-1}}\right),$$

which proves (6.23). All excluded sieving primes were fixed in advance of r, L, and residue-count errors are uniform in the position of the encompassing square. This proves the claimed uniformity.

6.7 Completing the prime-pattern construction

For each of the t forms, multiply the number of its possible bad prime values in (6.20) by the uniform fiber bound (6.23). A union bound shows that at most

$$\frac{C_A X^3}{L(\log X)^t} + o_A\left(\frac{X^3}{(\log X)^t}\right) \tag{6.24}$$

of the prime patterns can have a common odd factor $l > L$ in one of their pairs of reduction orders. More explicitly, the ratios of the two other terms to $X^3/(\log X)^t$ are $O_A(X^{-1/8}\log X)$ and $O_A(X^{-1/2}(\log X)^2)$; both tend to zero. None of the constants in this upper estimate depends on L.

Choose L, after all data depending on A have been fixed, large enough that

$$C_A/L < \tfrac{1}{2}c_A(\log L)^{-C_0}.$$

This is possible because $(\log L)^{C_0}/L \to 0$. Then choose X sufficiently large for (6.12), for all the upper estimates, and for the residual little-oh in (6.24) to be smaller than the remaining lower-bound coefficient. At least one prime pattern survives. Small common odd factors were already excluded by the CRT conditions, and all larger ones have just been removed. The 2-part and the prime-to-r assertions proved from (6.5) now give (6.2). The primes have size at least $c_A X$, so increasing X makes all of them exceed any prescribed B. This proves Lemma 6.1.

Together with Corollary 6.2, the prime-pattern lemma proves Corollary 2.6, completing the ordinary integer witnesses used in Section 2. The analysis there combines the tested representations with the parity, height, and elliptic inputs to turn success of all finite tests into an ordinary integer zero.

7 Turing degree and a quartic normal form

The finite-test reduction determines the Turing degree of rational solvability. A separate arithmetic-circuit conversion gives the same degree for two restricted forms of the input.

Corollary 7.1. *With the usual effective coding, let* $\mathrm{H10}(\mathbb{Q})$ *be the set of integral polynomials having a rational zero, with the number of variables part of the input. Then* $\mathrm{H10}(\mathbb{Q})$ *has Turing degree* $0'$*, the degree of the halting problem. The same is true for each of the following restricted input problems:*

(i) rational solvability for integral polynomials of total degree at most four;

(ii) given a nonempty finite list $q_1, \dots, q_s \in \mathbb{Z}[X_1, \dots, X_N]$*, each of total degree at most two, whether* $\sum_{i=1}^{s} q_i^2$ *has a zero in* $\mathbb{Q}^N$*.*

No bound is imposed on the number of variables, the number of summands in (ii), or the coefficient heights.

Proof. Let $\mathrm{H10}(\mathbb{Z})$ denote the analogous set of integral polynomials having an integer zero. The parallel search in the proof of Theorem 1.1 in Section 2 can be run with an $\mathrm{H10}(\mathbb{Q})$ oracle in place of the hypothetical rational decision procedure. The finite tests of §2.4 are uniformly computable from f and their indices once the fixed arithmetic data have been chosen, and each test uses finitely many oracle queries. The search therefore halts on every input, proving

$$\mathrm{H10}(\mathbb{Z}) \leq_{\mathrm{T}} \mathrm{H10}(\mathbb{Q}),$$

where $\leq_{\mathrm{T}}$ denotes Turing reducibility. The fixed choices used to define $\mathcal{T}$ are independent of f, so this construction gives a single oracle machine.

The Davis–Putnam–Robinson–Matiyasevich theorem, in the form [13, Theorem 8.1], applies to a positive-integer coding of the halting set. Its positive-integer witnesses can be replaced by $1 + u_1^2 + u_2^2 + u_3^2 + u_4^2$, using the four-square theorem as in [13, p. 234]. Specializing the parameter in the resulting fixed polynomial gives a computable many-one

reduction of the halting set to H10($\mathbb{Z}$). Integer solvability is itself computably enumerable by enumeration of integer tuples, so it has Turing degree $0'$. Rational solvability is also computably enumerable: enumerate rational tuples and evaluate the input polynomial exactly. Every computably enumerable set is Turing reducible to the halting set. Together with the displayed oracle reduction, this proves that H10($\mathbb{Q}$) has degree $0'$.

For the degree restriction, we apply over $\mathbb{Q}$ the algebraic normalization used in [13, Theorem 7.5]. Given $f \in \mathbb{Z}[X_1, \ldots, X_n]$, effectively construct an arithmetic circuit evaluating f, with input nodes, integral constant nodes, and addition and multiplication gates. Give every node v a new variable Y_v. For an input node X_j and a constant node c, impose respectively

$$Y_v - X_j = 0, \qquad Y_v - c = 0.$$

For a sum or product node v with incoming nodes u, w, impose respectively

$$Y_v - Y_u - Y_w = 0, \qquad Y_v - Y_u Y_w = 0.$$

Finally impose $Y_{\text{out}} = 0$ at the output node. This gives a nonempty finite list $q_1, \ldots, q_s$ of integral polynomials of total degree at most two in $N = n+m$ common variables $\mathbf{Z} = (\mathbf{X}, \mathbf{Y})$, where m is the number of circuit nodes. It also handles constant inputs. Induction through the circuit shows that each rational input tuple has a unique extension satisfying the gate equations, with output value $f(\mathbf{X})$. Since squares of rationals are nonnegative, it follows that

$$\begin{aligned} & (\exists \mathbf{x} \in \mathbb{Q}^n)\ f(\mathbf{x}) = 0 \\ \iff \quad & (\exists \mathbf{z} \in \mathbb{Q}^N)\ \bigwedge_{i=1}^{s} (q_i(\mathbf{z}) = 0) \\ \iff \quad & (\exists \mathbf{z} \in \mathbb{Q}^N)\ \sum_{i=1}^{s} q_i(\mathbf{z})^2 = 0. \end{aligned}$$

Let S_2 denote the presented list problem in (ii), and let $H_{\leq 4}$ denote the problem in (i). Sending f to the list $(q_1, \ldots, q_s)$, and then expanding its squared sum, gives computable many-one reductions

$$\text{H10}(\mathbb{Q}) \leq_{\text{m}} S_2 \leq_{\text{m}} H_{\leq 4}.$$

The expanded polynomial has integral coefficients and total degree at most four. Both target sets are computably enumerable by rational-tuple enumeration, so the reductions and the first part of the proof give Turing degree $0'$ for both.

The many-one reduction to the restricted problems starts from H10($\mathbb{Q}$); the reduction from H10($\mathbb{Z}$) remains a Turing reduction. This argument does not establish many-one completeness of H10($\mathbb{Q}$). The list in (ii) is supplied as part of the input, so no recognition of sum-of-squares presentations is needed. The construction supplies no bound on the number of variables or coefficient heights, and it gives no single existential definition of $\mathbb{Z}$ in $\mathbb{Q}$. □

References

[1] Levent Alpöge. Modularity and effective Mordell I, 2021. arXiv:2109.07917v2, November 24, 2021. Source.

[2] Levent Alpöge, Manjul Bhargava, Wei Ho, and Ari Shnidman. Rank stability in quadratic extensions and Hilbert's tenth problem for the ring of integers of a number field. *Inventiones Mathematicae*, 243:1129–1139, 2026. Version of record published online December 1, 2025. doi:10.1007/s00222-025-01392-3.

[3] Enrico Bombieri and Walter Gubler. *Heights in Diophantine Geometry*, volume 4 of *New Mathematical Monographs*. Cambridge University Press, 2006. doi:10.1017/CBO9780511542879.

[4] Jean-Benoît Bost. Périodes et isogénies des variétés abéliennes sur les corps de nombres (d'après D. Masser et G. Wüstholz). In *Séminaire Bourbaki, Vol. 1994/95, Exposés 790–804*, number 237 in Astérisque, pages 115–161. Société Mathématique de France, 1996. Exposé 795. Source.

[5] Henri Carayol. Sur les représentations l-adiques associées aux formes modulaires de Hilbert. *Annales scientifiques de l'École Normale Supérieure*, 19(3):409–468, 1986. doi:10.24033/asens.1512.

[6] J. W. S. Cassels. Arithmetic on curves of genus 1. IV. Proof of the Hauptvermutung. *Journal für die reine und angewandte Mathematik*, 211:95–112, 1962. doi:10.1515/crll.1962.211.95.

[7] Gunther Cornelissen and Alexandra Shlapentokh. Defining the integers in large rings of a number field using one universal quantifier. *Journal of Mathematical Sciences (New York)*, 158(5):713–726, 2009. Originally published in Zapiski Nauchnykh Seminarov POMI 358 (2008), 199–223; author version arXiv:0708.3075v2, February 14, 2008. Author version. doi:10.1007/s10958-009-9404-4.

[8] Gunther Cornelissen and Karim Zahidi. Topology of Diophantine sets: remarks on Mazur's conjectures. In Jan Denef, Leonard Lipshitz, Thanases Pheidas, and Jan Van Geel, editors, *Hilbert's Tenth Problem: Relations with Arithmetic and Algebraic Geometry*, volume 270 of *Contemporary Mathematics*, pages 253–260. American Mathematical Society, Providence, RI, 2000. Author preprint dated June 20, 2000. doi:10.1090/conm/270/04377.

[9] Gunther Cornelissen and Karim Zahidi. Elliptic divisibility sequences and undecidable problems about rational points. *Journal für die reine und angewandte Mathematik*, 613:1–33, 2007. Author manuscript, version 3, June 23, 2006. Source.

[10] David A. Cox. *Primes of the Form x^2+ny^2: Fermat, Class Field Theory, and Complex Multiplication*. John Wiley & Sons, second edition, 2013. doi:10.1002/9781118400722.

[11] Henri Darmon and Andrew Granville. On the equations $z^m = F(x,y)$ and $Ax^p + By^q = Cz^r$. *Bulletin of the London Mathematical Society*, 27(6):513–543, 1995. doi:10.1112/blms/27.6.513.

[12] Harold Davenport. *Multiplicative Number Theory*, volume 74 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, third edition, 2000. Revised and with a preface by Hugh L. Montgomery. Source.

[13] Martin Davis. Hilbert's tenth problem is unsolvable. *The American Mathematical Monthly*, 80(3):233–269, 1973. Source.

[14] Martin Davis, Hilary Putnam, and Julia Robinson. The decision problem for exponential Diophantine equations. *Annals of Mathematics*, 74(3):425–436, 1961. doi:10.2307/1970289.

[15] Pierre Deligne. Formes modulaires et représentations ℓ-adiques. In *Séminaire Bourbaki, Vol. 1968/69, Exposés 347–363*, volume 179 of *Lecture Notes in Mathematics*, pages 139–172. Springer-Verlag, 1971. Exposé 355. Source.

[16] Pierre Deligne. Travaux de Shimura. In *Séminaire Bourbaki, Vol. 1970/71, Exposés 382–399*, volume 244 of *Lecture Notes in Mathematics*, pages 123–165. Springer-Verlag, 1971. Exposé 389. Source.

[17] Pierre Deligne. Variétés de Shimura: interprétation modulaire, et techniques de construction de modèles canoniques. In *Automorphic Forms, Representations and L-Functions, Part 2*, volume 33 of *Proceedings of Symposia in Pure Mathematics*, pages 247–289. American Mathematical Society, 1979. Source.

[18] Jeroen Demeyer and Jan Van Geel. An existential divisibility lemma for global fields. *Monatshefte für Mathematik*, 147(4):293–308, 2006. Author manuscript dated July 11, 2005. doi:10.1007/s00605-005-0342-z.

[19] Kirsten Eisenträger, Russell Miller, Jennifer Park, and Alexandra Shlapentokh. As easy as $\mathbb{Q}$: Hilbert's tenth problem for subrings of the rationals and number fields. *Transactions of the American Mathematical Society*, 369(11):8291–8315, 2017. doi:10.1090/tran/7075.

[20] Gerd Faltings. Endlichkeitssätze für abelsche Varietäten über Zahlkörpern. *Inventiones Mathematicae*, 73(3):349–366, 1983. doi:10.1007/BF01388432. Erratum: ibid. 75 (1984), 381, doi:10.1007/BF01388572.

[21] Jean-Marc Fontaine. Représentations p-adiques semi-stables. In Jean-Marc Fontaine, editor, *Périodes p-adiques: Séminaire de Bures, 1988*, number 223 in Astérisque, pages 113–184. Société Mathématique de France, 1994. Exposé III. Source.

[22] Natalia Garcia-Fritz, Hector Pasten, and Xavier Vidaux. Effectivity for existence of rational points is undecidable. *Journal of Number Theory*, 276:81–97, 2025. doi:10.1016/j.jnt.2025.01.023.

[23] Éric Gaudron and Gaël Rémond. Polarisations et isogénies. *Duke Mathematical Journal*, 163(11):2057–2108, 2014. doi:10.1215/00127094-2782528.

[24] Éric Gaudron and Gaël Rémond. Théorème des périodes et degrés minimaux d'isogénies. *Commentarii Mathematici Helvetici*, 89(2):343–403, 2014. doi:10.4171/CMH/322.

[25] Ben Green and Terence Tao. The Möbius function is strongly orthogonal to nilsequences. *Annals of Mathematics*, 175(2):541–566, 2012. doi:10.4007/annals.2012.175.2.3.

[26] Ben Green, Terence Tao, and Tamar Ziegler. An inverse theorem for the Gowers $U^{s+1}[N]$-norm. *Annals of Mathematics*, 176(2):1231–1372, 2012. doi:10.4007/annals.2012.176.2.11.

[27] Ben Green, Terence Tao, and Tamar Ziegler. Erratum for An inverse theorem for the Gowers $U^{s+1}[N]$-norm, April 2024. Six-page author erratum. Source.

[28] Benjamin Green and Terence Tao. Linear equations in primes. *Annals of Mathematics*, 171(3):1753–1850, 2010. doi:10.4007/annals.2010.171.1753.

[29] Xavier Guitart and Santiago Molina. Parametrization of Abelian K-surfaces with quaternionic multiplication. *Comptes Rendus. Mathématique*, 347(23–24):1325–1330, 2009. doi:10.1016/j.crma.2009.09.025.

[30] D. R. Heath-Brown. The size of Selmer groups for the congruent number problem, II. *Inventiones Mathematicae*, 118:331–370, 1994. With an appendix by Paul Monsky, pp. 365–370. doi:10.1007/BF01231536.

[31] David Hilbert. Mathematische Probleme. Vortrag, gehalten auf dem internationalen Mathematiker-Kongreß zu Paris 1900. *Nachrichten von der Königlichen Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1900(3):253–297, 1900. Source.

[32] Ariyan Javanpeykar. Polynomial bounds for Arakelov invariants of Belyi curves. *Algebra & Number Theory*, 8(1):89–140, 2014. With an appendix by Peter Bruin. doi:10.2140/ant.2014.8.89.

[33] Jochen Koenigsmann. Defining $\mathbb{Z}$ in $\mathbb{Q}$. *Annals of Mathematics*, 183(1):73–93, 2016. doi:10.4007/annals.2016.183.1.2.

[34] Peter Koymans and Carlo Pagano. Hilbert's tenth problem via additive combinatorics. *Journal of the American Mathematical Society*, 2026. Published electronically August 10, 2026; arXiv:2412.01768v3 submitted November 24, 2025. doi:10.1090/jams/1081.

[35] D. D. Long, C. Maclachlan, and A. W. Reid. Arithmetic Fuchsian groups of genus zero. *Pure and Applied Mathematics Quarterly*, 2(2):569–599, 2006. doi:10.4310/PAMQ.2006.v2.n2.a9.

[36] David Marker. *Model Theory: An Introduction*, volume 217 of *Graduate Texts in Mathematics*. Springer, New York, 2002. doi:10.1007/b98860.

[37] Yu. V. Matiyasevich. Diofantovost' perechislimykh mnozhestv. *Doklady Akademii Nauk SSSR*, 191(2):279–282, 1970. English title: The Diophantineness of enumerable sets. Source.

[38] Barry Mazur. Speculations about the topology of rational points: an up-date. *Astérisque*, (228):165–181, 1995. Source.

[39] Barry Mazur and Karl Rubin. Ranks of twists of elliptic curves and Hilbert's tenth problem. *Inventiones Mathematicae*, 181(3):541–575, 2010. doi:10.1007/s00222-010-0252-0.

[40] James S. Milne. Abelian Varieties, 2008. Course notes, version 2.0, March 16. Source.

[41] James S. Milne. Lectures on Étale Cohomology, 2013. Course notes, version 2.21, March 22. Source.

[42] James S. Milne. Class Field Theory, 2020. Course notes, version 4.03, August 6. Source.

[43] Youcef Mokrani. Adaptation of Monsky matrices for θ-congruent numbers. *International Journal of Number Theory*, 16(2):377–396, 2020. Published online September 5, 2019. doi:10.1142/S1793042120500207.

[44] Hugh L. Montgomery and Robert C. Vaughan. *Multiplicative Number Theory I: Classical Theory*, volume 97 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 2007. doi:10.1017/CBO9780511618314.

[45] David Mumford. *Abelian Varieties*, volume 5 of *Tata Institute of Fundamental Research Studies in Mathematics*. Oxford University Press, London, 1970.

[46] M. Ram Murty and Hector Pasten. Modular forms and effective Diophantine approximation. *Journal of Number Theory*, 133(11):3739–3754, 2013. doi:10.1016/j.jnt.2013.05.006.

[47] Joan Nualart Riera. *On the Hyperbolic Uniformization of Shimura Curves with an Atkin–Lehner Quotient of Genus 0*. PhD thesis, Universitat de Barcelona, October 2015. Title-page date October 2015; repository release February 3, 2016. Source.

[48] OpenAI. A pointwise 2-converse for elliptic curves with rational two-torsion. OpenAI Math Release preprint OAI:A-pointwise-2-converse-for-elliptic-curves-with-rational-two-torsion-October-7-2026, 2026. Theorem 1.1.

[49] OpenAI. Fontaine–Mazur modularity at the prime 2. OpenAI Math Release preprint OAI:Fontaine-Mazur-modularity-at-the-prime-2-October-6-2026, 2026. Theorem 1.1 and Section 2.1.

[50] OpenAI. Goldfeld's analytic density conjecture and the 2-converse for elliptic curves. OpenAI Math Release preprint OAI:Goldfelds-analytic-density-conjecture-and-the-2-converse-for-elliptic-curves-October-7-2026, 2026. Theorem 1.1.

[51] Fabien Pazuki. Theta height and Faltings height. *Bulletin de la Société Mathématique de France*, 140(1):19–49, 2012. doi:10.24033/bsmf.2623.

[52] Bjorn Poonen. Using elliptic curves of rank one towards the undecidability of Hilbert's tenth problem over rings of algebraic integers. In Claus Fieker and David R. Kohel, editors, *Algorithmic Number Theory: 5th International Symposium, ANTS-V, Sydney, Australia, July 2002, Proceedings*, volume 2369 of *Lecture Notes in Computer Science*, pages 33–42. Springer-Verlag, Berlin, 2002. Source.

[53] Bjorn Poonen. Hilbert's tenth problem and Mazur's conjecture for large subrings of $\mathbb{Q}$. *Journal of the American Mathematical Society*, 16(4):981–990, 2003. doi:10.1090/S0894-0347-03-00433-8.

[54] Bjorn Poonen. Characterizing integers among rational numbers with a universal-existential formula. *American Journal of Mathematics*, 131(3):675–682, 2009. doi:10.1353/ajm.0.0057.

[55] Bjorn Poonen and Michael Stoll. The Cassels–Tate pairing on polarized abelian varieties. *Annals of Mathematics*, 150(3):1109–1149, 1999. doi:10.2307/121064.

[56] Elisabeth Eve Pyle. *Abelian Varieties over* $\mathbf{Q}$ *with Large Endomorphism Algebras and Their Simple Components over* $\overline{\mathbf{Q}}$. PhD thesis, University of California, Berkeley, 1995.

[57] Jordi Quer. Fields of definition of $\mathbf{Q}$-curves. *Journal de Théorie des Nombres de Bordeaux*, 13(1):275–285, 2001. doi:10.5802/jtnb.321.

[58] Kenneth A. Ribet. Galois representations attached to eigenforms with nebentypus. In Jean-Pierre Serre and Don Bernard Zagier, editors, *Modular Functions of One Variable V*, volume 601 of *Lecture Notes in Mathematics*, pages 18–52. Springer, Berlin, Heidelberg, 1977. doi:10.1007/BFb0063943.

[59] Kenneth A. Ribet. Abelian varieties over $\mathbf{Q}$ and modular forms. In *Algebra and Topology 1992*, pages 53–79. Korea Advanced Institute of Science and Technology, Taejon, 1992. Preprint arXiv:alg-geom/9208002. Source.

[60] Julia Robinson. Definability and decision problems in arithmetic. *The Journal of Symbolic Logic*, 14(2):98–114, 1949. doi:10.2307/2266510.

[61] Peter Scholze. p-adic Hodge theory for rigid-analytic varieties. *Forum of Mathematics, Pi*, 1:e1, 1–77, 2013. Corollary 1.8; see also the 2016 corrigendum. doi:10.1017/fmp.2013.1.

[62] Peter Scholze. p-adic Hodge theory for rigid-analytic varieties—corrigendum. *Forum of Mathematics, Pi*, 4:e6, 2016. doi:10.1017/fmp.2016.4.

[63] Atle Selberg. On an elementary method in the theory of primes. *Norske Videnskabers Selskab, Forhandlinger*, 19(18):64–67, 1947.

[64] Jean-Pierre Serre. *Local Fields*, volume 67 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1979. Translated from the French by Marvin Jay Greenberg. doi:10.1007/978-1-4757-5673-9.

[65] Carl Ludwig Siegel. Über die Classenzahl quadratischer Zahlkörper. *Acta Arithmetica*, 1(1):83–86, 1935. doi:10.4064/aa-1-1-83-86.

[66] Alice Silverberg and Yuri G. Zarhin. Semistable reduction and torsion subgroups of abelian varieties. *Annales de l'Institut Fourier*, 45(2):403–420, 1995. doi:10.5802/aif.1459.

[67] Joseph H. Silverman. *The Arithmetic of Elliptic Curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, New York, second edition, 2009. doi:10.1007/978-0-387-09494-6.

[68] John Voight. *Quaternion Algebras*, volume 288 of *Graduate Texts in Mathematics*. Springer, Cham, 2021. doi:10.1007/978-3-030-56694-4.

[69] Rafael von Känel. Modularity and integral points on moduli schemes, 2013. arXiv:1310.7263v1, October 27, 2013; revised version 2, March 23, 2014. Source.

[70] Rafael von Känel. The effective Shafarevich conjecture for abelian varieties of GL_2-type. *Forum of Mathematics, Sigma*, 9:e39, 1–29, 2021. doi:10.1017/fms.2021.29.

[71] Rafael von Känel and Arno Kret. Integral points on coarse Hilbert moduli schemes, 2023. arXiv:2307.06944v1, July 13, 2023. Source.

[72] Tao Wei and Xuejun Guo. The rank of 2-Selmer group associate to θ-congruent numbers, 2022. arXiv:2210.01678v1, October 4, 2022. Source.

[73] Kenneth S. Williams. Mertens' theorem for arithmetic progressions. *Journal of Number Theory*, 6:353–359, 1974. doi:10.1016/0022-314X(74)90032-8.

[74] Yuri G. Zarhin. Abelian varieties, quaternion trick and endomorphisms. In Ivan Cheltsov, Xiuxiong Chen, Ludmil Katzarkov, and Jihun Park, editors, *Birational Geometry, Kähler–Einstein Metrics and Degenerations*, volume 409 of *Springer Proceedings in Mathematics and Statistics*, pages 857–864. Springer, Cham, 2023. Author version arXiv:2010.07995v5. doi:10.1007/978-3-031-17859-7_42.