ECC RESEARCH PROGRAM · P3.1

Four Analytic Leaves, One Removed: The GRH Dependency of the Supersingular Isogeny-Path / Endomorphism-Ring Equivalence

A complete lemma-level usage map, an unconditional rank-four prime sampler replacing the fixed-form leaf, and exact obstructions to four shortcuts

math.st [1] **@math__street** [1]

with GPT 5.6 Sol · Claude Fable 5

[1]Independent research collective, math.st · July 2026

ABSTRACT

The foundational theorem of isogeny-based cryptography — that finding smooth-degree isogenies between supersingular curves (Problem A) is equivalent to computing their endomorphism rings (Problem B) — is proved, in its strongest known form, only under the generalized Riemann hypothesis. The 2026 unconditional equivalence of Herlédan Le Merdy and Wesolowski covers unrestricted isogeny degree and explicitly retains GRH for the smooth-degree problem ℓ-`IsogenyPath`. We ask exactly where GRH lives in the smooth-degree proof and how much of it can be removed. First, a source-checked lemma-level audit of Wesolowski's FOCS 2021 proof shows that precisely four direct analytic leaves carry every GRH-qualified result: D1 (a polylogarithmic Frobenius prime), D2 (prime density of one fixed quadratic form), D3 (a power-of-x Titchmarsh modulus range), and D4 (polylogarithmic class-group expansion); one further appeal to ordinary RH is removable by the unconditional prime number theorem. Second, we remove D2: a direct rank-four residue-lattice sampler, combining Rouse's effective quaternary representability, explicit primes in arithmetic progressions, and fixed-rank CVP, replaces the fixed-binary-form sampler of Proposition 3.8 unconditionally in expected time polynomial in $\log p$ and the numerical value of ℓ. Third, four tempting shortcuts fail for precise structural reasons — Brandt mixing lacks a pushforward anti-concentration bound, ramification at p cannot be hidden from any orthogonal $2+2$ norm split, ellipsoid rejection for one fixed target loses a factor of the target, and the 2026 flexible dictionary does not reach the small-discriminant core of the old norm-equation solver. The equivalence therefore remains conditional through D1, D3, and D4, and we isolate the single strongest replacement target: a constructive fixed-target quaternary norm solver on arbitrary maximal-order lattices. Toy-scale experiments validate the sampler's residue and density mechanism (22 cases, zero reciprocity violations) and a right-order-aware Deuring round trip (30/30 seeds at $p = 11$); a five-size timing fit (exponent 1.68, $R^2 = 0.98$) supplies a first measured component of the concrete security-loss function while demonstrating precisely why no security-bit loss can yet be reported.

Keywords. supersingular elliptic curves · endomorphism rings · Deuring correspondence · GRH · quaternion algebras · quadratic forms · isogeny-based cryptography

1 Introduction

Supersingular elliptic curves over $\mathbb{F}_{p^2}$ have endomorphism algebras isomorphic to the quaternion algebra $B_{p,\infty}$ ramified exactly at p and ∞; this is Deuring's correspondence [1], and its basic statement requires no analytic hypothesis. CITED Isogeny-based cryptography rests on two computational problems built on that correspondence: **Problem A** (ℓ-`IsogenyPath`), finding an isogeny of smooth degree between two given supersingular curves, and **Problem B** (`EndRing` / `MaxOrder`), computing the endomorphism ring of a given supersingular curve as an explicit maximal order. The load-bearing security theorem is that the two are equivalent. CITED Wesolowski proved this equivalence in expected polynomial time — but under GRH [2]. CITED In 2026, Herlédan Le Merdy and Wesolowski made the **unrestricted-degree** equivalence unconditional, and stated explicitly that the reduction from ℓ-`IsogenyPath` to `EndRing` remains GRH-conditional [3].

Problem P3.1 asks two things: (Task 1) prove the smooth-degree equivalence without GRH, or establish a precise necessity result; and (Task 2) quantify the concrete security loss of the reductions as a function of p. This paper reports the state of both tasks after a two-session research effort. Neither task is resolved; what we obtained instead is a complete map of the obstruction, the unconditional removal of one of its four analytic supports, exact post-mortems for four natural shortcuts, and a validated toy-scale experimental base for the loss function. We report these results with their original epistemic tags.

Main findings (honest summary). **(1)** PROVED The audited equivalence proof [2] has exactly four direct GRH leaves, D1–D4 (§3); every GRH-qualified lemma inherits from them, and one extra appeal to ordinary RH is removable by the prime number theorem. **(2)** PROVED Leaf D2 — prime sampling from one fixed quadratic-form class — is removed unconditionally: a direct rank-four residue-lattice sampler replaces Wesolowski's Proposition 3.8 in expected time polynomial in $\log p$ and the numerical value of ℓ (§5). This is a repo-internal candidate proof, not a published resolution. **(3)** PROVED Leaves D1, D3, D4 remain. Four shortcuts (Brandt-average transfer, basis change to hide p, direct quaternary rejection for one target, a flexible-dictionary hybrid) fail with exact structural obstructions (§6). **(4)** EMPIRICAL: p <= 71, ell = 3 The sampler's residue and density mechanism and a right-order-aware local Deuring round trip are validated at toy scale; the recorded timing fit does **not** identify a security-bit loss, and we prove exactly why not (§7–8).

1.1 Contributions and honest scope

We contribute (i) a source-checked GRH usage map with a complete dependency closure (§3, Table 1, Figure 1); (ii) an analysis of the strongest checked unconditional substitutes and the exact point at which each fails to preserve polynomial time (§4); (iii) an unconditional replacement for the fixed-form prime sampler, with its proof architecture and remaining obligations stated openly (§5); (iv) four precise negative results (§6); (v) toy-scale validation with real data for every empirical claim (§7); and (vi) a non-identifiability statement for the concrete loss function, together with the counters a valid measurement must record (§8).

The scope boundary is strict. Task 1 is **not** solved: the smooth-degree equivalence remains conditional in the checked literature and in this work. PROVED Task 2 is **not** solved: the recorded

experiments contain no oracle call, so the bit-loss proxy is not yet evaluable. All empirical results live at toy scale ($p \leq 71$) and are labeled with their exact parameter ranges.

2 Setting and notation

Fix a prime $p > 3$. A supersingular elliptic curve $E/\mathbb{F}_{p^2}$ has $\mathrm{End}(E)$ isomorphic to a maximal order in the quaternion algebra $B_{p,\infty}$ ramified exactly at p and ∞ [1]. For a left ideal I of a maximal order $\mathcal{O}$, we write Nrd for the reduced norm, $\mathrm{Nrd}(I)$ for the ideal norm, and $q_I(\alpha) = \mathrm{Nrd}(\alpha)/\mathrm{Nrd}(I)$ for the normalized norm form, a primitive positive-definite integral quadratic form of rank four and discriminant p^2 [2], [4].

Definition 1 (the two problems). **Problem A** (ℓ-`IsogenyPath`): given supersingular $E_1, E_2/\mathbb{F}_{p^2}$ and a prime $\ell \neq p$, find an isogeny $E_1 \to E_2$ of smooth (ℓ-power) degree. **Problem B** (`EndRing`, with variant `MaxOrder`): given supersingular $E/\mathbb{F}_{p^2}$, compute $\mathrm{End}(E)$ as an explicit maximal order in $B_{p,\infty}$. The audited equivalence results are Theorems 7.2, 7.4, 8.1, and 8.3 of [2], all conditional on GRH.

Definition 2 (usage-map conventions). A **direct GRH use** is a lemma whose own proof invokes a GRH-conditional analytic theorem. An **inherited GRH dependency** is a result that calls a direct-use lemma without adding another analytic assumption. A reduction is **polynomial-time** only when its cost is polynomial in the bit length $\log p$ and in the explicit input/output lengths — a bound that is polynomial in the **numerical value** of a size-$p^{O(1)}$ parameter is exponential in the input length and does not qualify.

The distinction in Definition 2 is the recurring theme of this paper: every failed unconditional substitute below fails by conflating a polynomial numerical bound with a polynomial bit-length bound.

Throughout, "the 2026 route" means the unconditional equivalence of unrestricted `Isogeny`, `EndRing`, and `MaxOrder` in [3], which we cite as a boundary: it does not produce smooth-degree paths, and its authors say so (Section 1.1 of [3]). CITED

3 The GRH usage map

The first deliverable of P3.1 was a lemma-level audit of [2] (arXiv:2111.01481v1; theorem numbers below refer to that version), with every GRH-qualified result expanded to its direct analytic inputs and every input source-checked in the primary literature.

Theorem 3 (four-leaf usage map). PROVED The proof of the GRH-conditional equivalence in [2] has exactly four direct GRH leaves: **D1** — construction of a small auxiliary Frobenius prime $q_p = O((\log p)^2)$ defining a special quaternion model; **D2** — a prime-counting theorem with square-root error for the values of one fixed primitive binary quadratic form; **D3** — a uniform Titchmarsh divisor estimate whose auxiliary moduli reach a fixed power of the sampled integer; **D4** — expansion of imaginary-quadratic class groups from prime ideals of polylogarithmic norm. Every GRH-qualified result in the paper is a direct instance of D1–D4 or inherits its assumption from them; Table 1 gives the dependency closure. One further appeal, to ordinary RH inside Theorem 6.4, is removable (Proposition 4 below).

The proof of Theorem 3 is the audit itself: an exhaustive expansion of every occurrence of GRH in the 31-page source, followed by dependency closure. We summarize the four leaves, then record the one removable use.

D1 (small special quaternion model). CONDITIONAL: GRH For $p \equiv 1 \bmod 8$, Lemma 2.2 of [2] chooses the least prime $q_p \equiv 3 \bmod 4$ with $\left(p/q_p\right) = -1$ and uses $q_p = O\big((\log p)^2\big)$, derived from effective Chebotarev in $\mathbb{Q}(\sqrt{p}, i)$ (Proposition 1 of [5]). Lemma 2.3 uses the same bound to make both the index $[\mathcal{O}_0 : R + Rj]$ and $|\mathrm{disc}\, R|$ of order $O\big((\log p)^2\big)$; Lemmas 2.5 and 2.6 compute the special curve and its quaternion/endomorphism dictionary by enumerating degree-q_p isogenies, so polynomial time **requires** $q_p = \mathrm{poly}(\log p)$.

D2 (primes represented by one fixed form). CONDITIONAL: GRH Theorem 3.1 of [2] states, for a primitive positive-definite binary form f of discriminant D,

$$\pi_f(\rho) = \frac{\delta\rho}{h(D)\log\rho} + O\big(\rho^{1/2}\log(|D|\rho)\big), \quad \delta \in \{1/2, 1\}, \tag{1}$$

via effective Chebotarev. Proposition 3.4 takes $\rho = O_\varepsilon\big(|D|^{1+\varepsilon}\big)$ so that uniform lattice sampling hits a prime with inverse-polylogarithmic probability; Propositions 3.5–3.6 reduce higher rank to this binary case; Theorem 3.7 applies it to the norm form q_I of discriminant p^2, and Proposition 3.8 adds a norm interval and the condition that ℓ be a quadratic nonresidue modulo the sampled prime.

D3 (uniform Titchmarsh range). CONDITIONAL: GRH Theorem 4.4 of [2] adapts Theorem 2.1 of [6] in the range $b, c, d \le x^\delta$ with power-saving error $O\big(x^{1-\delta}\big)$; Theorem 4.2 and Corollary 4.3 use it to count, uniformly over a genus, representations of n as a prime plus a binary-form value; Theorem 5.1 turns the count into an expected-polynomial-time solver for

$$\det(\gamma)^2 f(s,t) + b f_\gamma(x,y) = n \tag{2}$$

with only $\log n \ge c \log b$ required of the parameter b.

D4 (small-prime class-group expansion). CONDITIONAL: GRH Lemma 5.3 of [2] invokes [7] to use prime ideals of norm at most

$$C = O_\varepsilon\big((\log|d|)^{2+\varepsilon} + \omega(m)^{1+\varepsilon}\big) \tag{3}$$

as an expanding Cayley multiset mixing in $O(\log|d|)$ steps, producing a common integer B, coprime to md, of bit length $O_\varepsilon\big(\log|d| \cdot \big((\log|d|)^{2+\varepsilon} + \omega(m)^{1+\varepsilon}\big)\big)$ whose divisors are represented across all form classes. Lemma 5.4 uses the same small generators with explicit class-group computation. Theorem 5.1 uses both to randomize the form f_γ within its class or genus.

> **Proposition 4** (the RH appeal in Theorem 6.4 is removable). PROVED The proof of Theorem 6.4 of [2] cites the prime number theorem under RH to bound from below the product of the primes in the window $(\log p)^{c_0} < \ell < (\log p)^{c_0+\delta}$. The unconditional prime number theorem suffices.

Proof. Set $X = (\log p)^{c_0}$ and $Y = (\log p)^{c_0+\delta}$. The unconditional PNT gives $\theta(Y) - \theta(X) = Y - X + o(Y) \sim Y$ for the Chebyshev function θ, so for any fixed $\delta > 0$ the product of primes in (X, Y) eventually exceeds $e^{(\log p)^{c_0}}$, and the only property used in Theorem 6.4 is that this product exceeds its powersmoothness target of size $\exp((\log p)^{c_0})$. Removing this appeal does not make Theorem 6.4 unconditional: it still calls Theorem 6.3, which inherits D1–D4. □

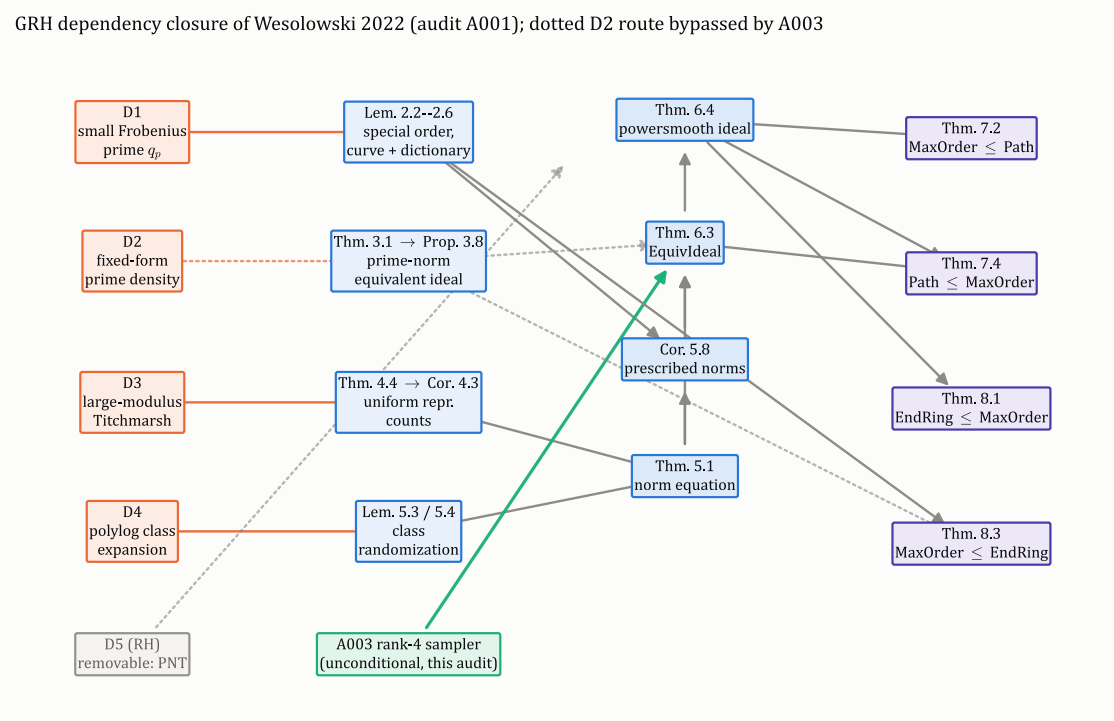

Figure 1: The dependency closure of the four GRH leaves through [2], following the audit `GRH_USAGE_MAP.md`. Orange: direct analytic leaves. Blue: conditional results of the quaternion core. Violet: the four final equivalence theorems. The dotted D2 route (Theorem 3.1 → Proposition 3.8) is the one this work replaces unconditionally (green, §5); D5 marks the removable RH appeal of Proposition 4. Translation lemmas 7.1/7.3 (inheriting D1 via Lemma 2.6) are omitted for legibility.

4 The strongest checked unconditional substitutes, and where they fail

For each leaf we checked the strongest unconditional replacement in the primary literature. Each fails at the boundary drawn in Definition 2.

Proposition 5 (D1: existence bounds do not search). PROVED Unconditional effective Chebotarev (Theorem 1.1 of [8]) gives $q_p \le p^{O(1)}$ for the D1 prime. This does not restore Lemma 2.2: enumerating primes up to p^C is exponential in the input length $\log p$, and the polynomial-size guarantees of Lemmas 2.3–2.6 fail with any bound weaker than $q_p \le (\log p)^C$ plus a polynomial-time search.

Proposition 6 (D2: least primes and averages are not densities). PROVED Theorem 1.2 of [8] shows every primitive positive-definite binary form of discriminant D represents a prime $\ll |D|^{694}$; this controls a least prime but not the inverse-polylogarithmic sampling density that Proposition 3.4 needs, and naive use is exponential in $\log|D|$ because the relevant ellipsoid holds $\Theta\left(\rho/\sqrt{|D|}\right)$ lattice points. The average-over-classes theorems of [9] (Theorem 1.1, Corollary 1.3) and [10] (Theorems 1.1–1.2, Corollary 1.8) certify a positive proportion of classes, not the fixed adversarial form q_I handed to Theorem 3.7.

Proposition 7 (D3: the unconditional range forces exponential outputs). PROVED Theorem 2.1 of [6] is unconditional when its auxiliary moduli are bounded by $(\log x)^C$; GRH is what enlarges

Table 1: PROVED Exhaustive audit of the GRH-qualified results in [2] (theorem numbers of arXiv:2111.01481v1) and their dependency closure. Deuring correspondence, lattice reduction, Brandt-graph mixing, quaternion ideal arithmetic, and the formal correctness arguments do not themselves invoke GRH; GRH enters exactly where the algorithms need small or statistically plentiful arithmetic representatives.

Result	Role	Leaves	Unconditional status (this audit)
Lem. 2.2	least q_p, model $(-q_p, -p)$	D1	fixed-power Chebotarev only; no polylog search
Lem. 2.3	small special order and discriminant	D1	fails with only $q_p \le p^{O(1)}$
Lem. 2.5	special curve + dictionary	D1	2026 route avoids the global dictionary
Lem. 2.6	powersmooth ideal-to-isogeny	D1	modern translation covers unrestricted degree only
Thm. 3.1	fixed-form prime count	D2	least-prime / average results lack sampling density
Prop. 3.4–3.6	binary and higher-rank sampling	D2	bypassed by A003 for rank-4 quaternion lattices
Thm. 3.7	equivalent prime-norm ideal	D2	A003 supplies the special-case sampler directly
Prop. 3.8	prime norm in interval, residue cond.	D2	**replaced unconditionally** (Theorem 9, §5)
Thm. 4.4	uniform Titchmarsh estimate	D3	unconditional ABL range is polylog moduli only
Thm. 4.2 / Cor. 4.3	uniform representation counts	D3	exponential output length unconditionally
Lem. 5.3	polylog-time class randomization	D4	no unconditional polylog expansion [7]
Lem. 5.4	small-disc. class enumeration	D4	runtime already poly in $\lvert d\rvert$, not $\log\lvert d\rvert$
Thm. 5.1	quaternion norm-equation solver	D3+D4	both analytic inputs essential to its proof
Cor. 5.8	prescribed norms in $\mathcal{O}_0$	D1+D3+D4	uses the small special form and Thm. 5.1
Thm. 6.3	`EquivIdeal` core	D1–D4	calls Prop. 3.8, Cor. 5.8, Thm. 5.1
Thm. 6.4	powersmooth equivalent ideal	D1–D4	PNT removes only the extra RH appeal (D5)
Lem. 7.1 / 7.3	isogeny/ideal translation	D1; D1–D4	enumeration calls Lem. 2.6
Thm. 7.2	`MaxOrder` reduces to ℓ-`IsogenyPath`	D1–D4	no additional analytic input
Thm. 7.4	ℓ-`IsogenyPath` reduces to `MaxOrder`	D1–D4	the direction left conditional in [3]
Thm. 8.1	`EndRing` reduces to `MaxOrder`	D1–D4	no additional analytic input
Thm. 8.3	`MaxOrder` reduces to `EndRing`	D1+D2	Step 3 calls Prop. 3.5; model uses small q_p

them to $x^{\delta'}$. In the quaternion application the modulus parameter is $b = p$, so the unconditional range forces $\log x \geq p^{1/C}$. Since x is on the scale of the represented integer n, the output bit length $\log n$ becomes exponential in the input length $\log p$, and the reduction is no longer polynomial-time.

Proposition 8 (D4: Minkowski generation neither shrinks nor mixes). PROVED Unconditionally, an imaginary-quadratic class group is generated by prime ideals of norm $O\left(\sqrt{|d|}\right)$ (choose in each class an integral ideal of Minkowski-bounded norm and factor it). This cutoff is exponential in $\log|d|$, and generation alone supplies no rapid-mixing bound, so substituting it in Lemma 5.3 makes the common integer B exponentially long. The authors of [7] record (their Sections 1 and 7.2) that even quadratic characters meet the least-prime-nonresidue obstruction, and offer only typical-modulus unconditional substitutes.

The pattern is uniform: unconditional analytic number theory currently supplies **existence at fixed-power distance** or **density on average**, while the reduction needs **density at a fixed instance within polylogarithmic bit-distance**. GRH is used exactly to close that gap, four times.

5 Removing D2: a direct quaternary prime sampler

The one leaf we could remove is D2, and the removal works precisely by refusing the binary restriction: the rank-four norm lattice has enough represented primes **unconditionally**, even though no fixed binary subform is known to.

Theorem 9 (unconditional replacement for Proposition 3.8). PROVED Let I be a left ideal of a maximal order in $B_{p,\infty}$ with normalized norm form $Q = q_I$, and let $\ell \neq p$ be prime. There is an algorithm that outputs an equivalent ideal of prime norm q in a prescribed interval with $(\ell/q) = -1$, with success probability at least inverse polynomial per sample, in expected time polynomial in $\log p$ and in the numerical value of ℓ. In particular D2 is not a necessary assumption for the smooth-path proof of [2].

Remark 10 (status). PROVED This is a written repo-internal candidate proof (`attempts/A003-direct-quaternary-prime-sampler.md`), checked against its cited sources and subjected to a toy refutation grid (§7.1); it is not a published or independently refereed resolution, and we treat it accordingly. Its complexity convention — polynomial in the **numerical value** of ℓ — matches the convention of the proposition it replaces.

5.1 Construction

For odd ℓ, set $M = 8\ell$. Choose a primitive $x \in I$ whose normalized norm $a = Q(x)$ satisfies $a \equiv 1 \bmod 8$, $a \not\equiv 0 \bmod p$, and a a quadratic nonresidue modulo ℓ. PROVED Such x is found by enumerating four coordinates modulo M — polynomial in the numerical value of ℓ because the normalized lattice is locally the determinant form at every prime dividing M — and the condition modulo p is enforced without enumerating modulo p: add kMv for a vector v on which the reduced quadratic polynomial is nonzero and try three values of k, since a nonzero quadratic polynomial has at most two roots. For $\ell = 2$, take $M = 8$ and $a \equiv 3$ or $5 \bmod 8$.

Define the residue sublattice

$$\Lambda = \mathbb{Z}x + MI, \quad [I : \Lambda] = M^3, \tag{4}$$

(the index when x extends to a basis of I). PROVED The restricted form Q_Λ is primitive, its discriminant is $D_\Lambda = M^6 p^2$ (Gram determinants scale by the square of the index), and if the level of Q divides a fixed multiple of p then the level of Q_Λ divides the same fixed multiple of $M^2 p$.

Lemma 11 (strong local solubility). PROVED Every prime $q \equiv a \bmod M$ with $q \neq p$ is strongly locally represented by Q_Λ in the sense of [11].

Proof. For $r \nmid M$ the lattices $\Lambda \otimes \mathbb{Z}_r$ and $I \otimes \mathbb{Z}_r$ agree. If additionally $r \neq p$, the normalized local form is equivalent to the determinant form on $M_2(\mathbb{Z}_r)$, and a diagonal matrix gives a good-type representation of every r-adic integer. At $r = p$, the reduced norm on the maximal order of the local division algebra maps units onto $\mathbb{Z}_p^\times$, and the trace pairing at a unit supplies a nonzero derivative, so every p-adic unit is strongly represented. For $r \mid M$, every prime $q \equiv a \bmod M$ satisfies $q/a \in (\mathbb{Z}_r^\times)^2$ — congruence modulo 8 handles $r = 2$ and congruence modulo ℓ handles $r = \ell$ — so kx with $k^2 a = q$ is a good-type local representation. □

5.2 Density, sampling, and reciprocity

CITED Rouse's effective strong-local theorem (Theorem 1 of [11], proved by splitting the theta coefficient $r_Q(n) = a_E(n) + a_C(n)$ into Eisenstein and cusp parts) gives a crossover $X_0 \leq (N(Q_\Lambda)D(Q_\Lambda))^C$, for an absolute constant C, beyond which $|a_C(q)| \leq a_E(q)/2$ after a constant-factor increase. PROVED Above the crossover, the local-density lower bounds give

$$r_{Q_\Lambda}(q) \geq c \cdot \frac{q}{\sqrt{D_\Lambda}} \cdot \frac{\varphi(N(Q_\Lambda))}{N(Q_\Lambda)}, \tag{5}$$

with effective absolute $c > 0$, and $\varphi(N)/N \gg 1/\log\log(N+3)$ makes the loss polylogarithmic. PROVED The level/discriminant-uniform strong-local case is the required input here: the sharper explicit bounds for **prime-discriminant** quaternary forms [12] do not directly cover these discriminant-p^2 norm lattices or their index-$(8\ell)^3$ residue sublattices. CITED The Siegel–Walfisz theorem supplies $\gg X/(\varphi(M)\log X)$ primes $q \in [X/2, X]$ with $q \equiv a \bmod M$ once $M \leq (\log X)^A$; taking $\log X \geq M^2$ puts the progression in the unconditional range while keeping the output length polynomial in the numerical value of ℓ, and the explicit progression bounds of [13] (valid threshold at most $\exp\left(0.03\sqrt{M}\log^3 M\right)$ for $M > 10^5$) make the estimate effective. PROVED Summing the representation bound over these primes yields at least

$$\gg \frac{X^2}{\sqrt{D_\Lambda}\,\varphi(M)\log X \log\log(N(Q_\Lambda)+3)} \tag{6}$$

prime-valued vectors in the radius-X ellipsoid, against $O\left(X^2/\sqrt{D_\Lambda}\right)$ total lattice vectors once X exceeds the squared covering radius. A uniform vector in the ellipsoid is therefore prime-valued with probability at least $1/\operatorname{poly}(\ell, \log X, \log p)$.

PROVED Sampling uniformly is done exactly as in Wesolowski's rank-two Voronoi rejection, extended verbatim to rank four using fixed-rank exact CVP ([14] gives $n^{O(n)}$ arithmetic operations, hence polynomial in fixed rank): sampling the ellipsoid enlarged by the covering radius μ and rounding to the nearest lattice point assigns the full equal-volume Voronoi cell to every accepted point, with acceptance probability at least $((r-\mu)/(r+\mu))^4$ at sampling radius r; finitely many random bits change the output by negligible statistical distance.

PROVED Finally, the residue design forces the reciprocity condition that Proposition 3.8 needs. For odd ℓ, every value of Q_Λ modulo ℓ is a square multiple of the nonresidue a, so a prime output q has $(q/\ell) = -1$; since $q \equiv 1 \bmod 4$, quadratic reciprocity gives $(\ell/q) = -1$. For $\ell = 2$, outputs satisfy $q \equiv 3$ or $5 \bmod 8$, so the supplementary law gives $(2/q) = -1$.

5.3 What the removal does and does not achieve

PROVED Theorem 9 removes D2 from the audited proof: the equivalent prime-norm ideal with interval and residue constraints — the only role D2 played — is now produced unconditionally. It does **not** prove the equivalence: the fixed-target norm equation of Algorithm 2, Step 9, still depends on D3 and D4, and the special-model construction still depends on D1. The current analytic frontier is the D3+D4 prescribed-norm combination; merely improving a least-prime exponent does not address it.

6 Four shortcuts and their exact obstructions

Each of the following attempts looked plausible and failed for a reason that is now a precise statement. We record all four; they constrain any future attack.

Proposition 12 (A002: Brandt mixing does not transfer averages). PROVED Wesolowski's Algorithm 2 first randomizes the input ideal class by a Brandt-graph walk (nearly uniform by the Ramanujan bound), then calls Proposition 3.8. The average-over-classes prime theorems of [9] and [10] cannot be substituted after this randomization: Proposition 3.5 selects its binary subform $g_{u,v}(x, y) = q_{I'}(xu + yv)$ by a deterministic LLL-and-gcd rule, the discriminant and class of $g_{u,v}$ vary with I', and no symmetry makes the selection equivariant under binary class-group composition. The sampled object is a quaternion ideal class, not a uniform binary class of fixed discriminant, and the checked sources contain no anti-concentration bound for the pushforward $[I'] \mapsto [g_{u,v}]$. Knowing that a positive proportion of binary classes is good does not help if the deterministic map may concentrate on the complement.

Proposition 13 (A004: ramification cannot be hidden from a 2+2 split). PROVED No integral orthogonal $2 + 2$ decomposition of a maximal-order norm lattice in $B_{p,\infty}$ makes all parameters entering Theorem 5.1 polylogarithmic in p. For odd p, a Hilbert symbol $(a, b)_{\mathbb{Q}_p}$ with both entries p-adic units is split, so a presentation of the division algebra must put odd p-adic valuation into a defining parameter: if the quadratic subfield K has discriminant coprime to p, the decomposition $B = K \oplus Kj$ places p in j^2, producing a block coefficient divisible by p; if instead the coefficient is coprime to p, then K is ramified at p and the binary discriminant is divisible by p. Integrally, for a block form $F_1 \oplus bF_2$ the rank-four Gram determinant is $b^2 \det(F_1) \det(F_2)$ up to normalization and the square of a lattice index; since the lattice discriminant is p^2, at least one of b, $\det(F_1)$, $\det(F_2)$, or the index is divisible by p — and each location reappears as a modulus or numerically-processed parameter in the Theorem 5.1 architecture, forcing $\log x \geq p^{1/C}$ as in Proposition 7.

Proposition 14 (A005: one fixed target defeats rejection sampling). PROVED Algorithm 2, Step 9, must solve

$$N^2 f(s, t) + p f_\Gamma(x, y) = n_2 \ell^e \tag{7}$$

for one fixed right-hand side. Rouse's theorem [11] can certify that a representation **exists** once the target is strongly locally soluble and above the fixed-power crossover — without D3

or D4 — but it does not locate one. A rank-four ellipsoid $Q(v) \le n$ contains on the order of $n^2/\sqrt{D(Q)}$ vectors while the single level set $Q(v) = n$ has $n^{1+o(1)}/\sqrt{D(Q)}$ in the theta-series regime, so uniform ellipsoid rejection succeeds with probability at most $n^{-1+o(1)}$ — exponentially small in $\log n$. The prime-valued sampler of Theorem 9 escapes this because it accepts $\sim X/\log X$ distinct target values; Step 9 permits exactly one.

Proposition 15 (A006: the flexible dictionary does not reach the smooth core). PROVED The 2026 machinery removes the special-curve **dictionary** manifestation of D1: Proposition 13 of [3] builds a flexible quaternion model and explicit dictionary from any endomorphism-ring basis, unconditionally. But the old smooth-equivalent-ideal route does not fail only at the dictionary layer: Corollary 5.8 and Algorithm 2 of [2] solve the smooth norm equation inside the special order $\mathcal{O}_0$, whose quadratic suborder has discriminant $O((\log p)^2)$ only through D1, and the proof of Theorem 5.1 performs an exhaustive search modulo that binary discriminant — its stated complexity is polynomial in the **numerical** discriminant. A flexible model found unconditionally has parameters of size $p^{O(1)}$: polynomial bit length, but a loop polynomial in the numerical value is exponential in $\log p$. The hybrid therefore leaves the frontier at D1+D3+D4. What survives: any future smooth-path construction may use the flexible dictionary of [3] to avoid the special curve, **provided** its ideal-smoothing step works directly on an arbitrary maximal-order norm lattice.

The remaining frontier. PROVED The strongest single replacement target is a **constructive polynomial-time fixed-target solver** for the structured quaternary norm equation on an arbitrary maximal-order lattice, with sufficient output entropy. Such a solver would bypass the special-order dependency D1 **and** the D3+D4 sampling architecture simultaneously (Propositions 14 and 15 isolate exactly this gap). Rouse-type existence above a polynomial threshold is provably insufficient (Proposition 14). Least-prime improvements alone address none of the remaining leaves.

7 Toy-scale empirical validation

All experiments are exhaustive toy-scale fixtures ($p \le 71$) intended to **refute** the mechanisms above if they were wrong, not to demonstrate asymptotic behavior. Every number in this section is stored in a seeded CSV in the problem's `data/` directory.

7.1 The sampler's residue and density mechanism

EMPIRICAL: p <= 31, ell in {3,5}, X = 3000 Across 22 maximal-order or seeded prime-ideal lattices with $3 \le p \le 31$, $p \equiv 3 \bmod 4$, the residue-lattice discriminant always equaled the predicted $(8\ell)^6 p^2$, every prime-valued vector satisfied $(\ell/q) = -1$, and no reciprocity violation occurred. The scaled prime density $\Pr[Q_\Lambda(v) \text{ prime}] \cdot \log X$ ranged from 0.832 to 1.815 with mean 1.276 (Figure 2) — consistent with the $1/\log X$ prime-density scale predicted before implementation. Coverage of admissible progression primes ranged from 0.10 to 0.778, confirming that these small cutoffs sit below a uniform all-primes representation regime, as expected below Rouse's crossover.

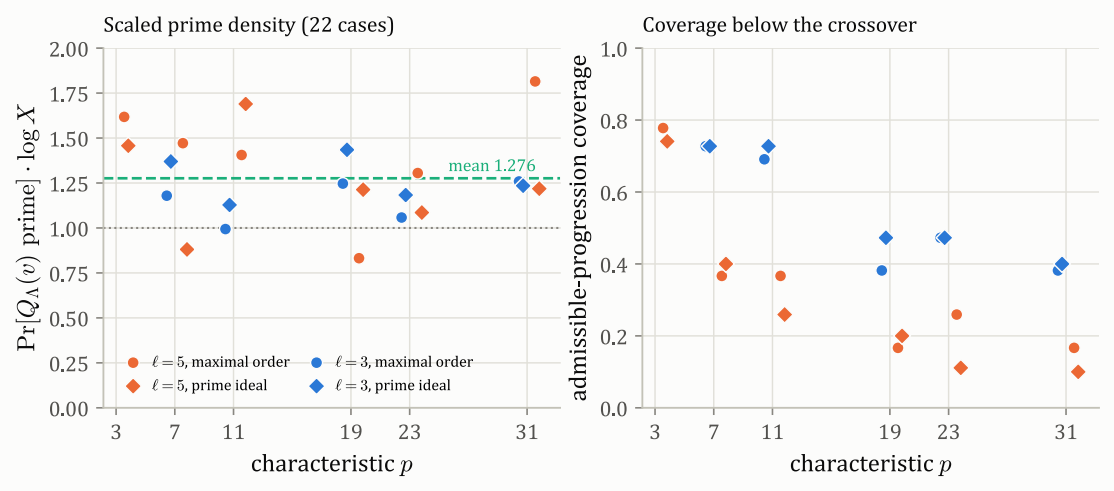

Figure 2: EMPIRICAL: p <= 31, ell in {3,5}, X = 3000 The 22-case A003 refutation grid. Left: scaled prime density per case (circles: maximal orders; diamonds: norm-ℓ' prime ideals), with the grid mean 1.276 (dashed) and the reference value 1 (dotted). Right: fraction of admissible progression primes $q \equiv a \bmod 8\ell$, $q \le 3000$, actually represented — incomplete at these cutoffs, exactly as the pre-registered prediction expected. Zero reciprocity violations occurred in any case. Data: `measure_quaternary_prime_sampler_p31_ells3-5_x3000_20260711.csv`.

EMPIRICAL: p in {7,11}, ell = 3, 250 <= X <= 4000 A five-cutoff scan of the four (p, ideal) cases shows the scaled density approaching a value near one: at $X = 4000$ it lies between 0.991 and 1.079, while admissible-prime coverage rises to between 0.75 and 0.824 (Figure 3).

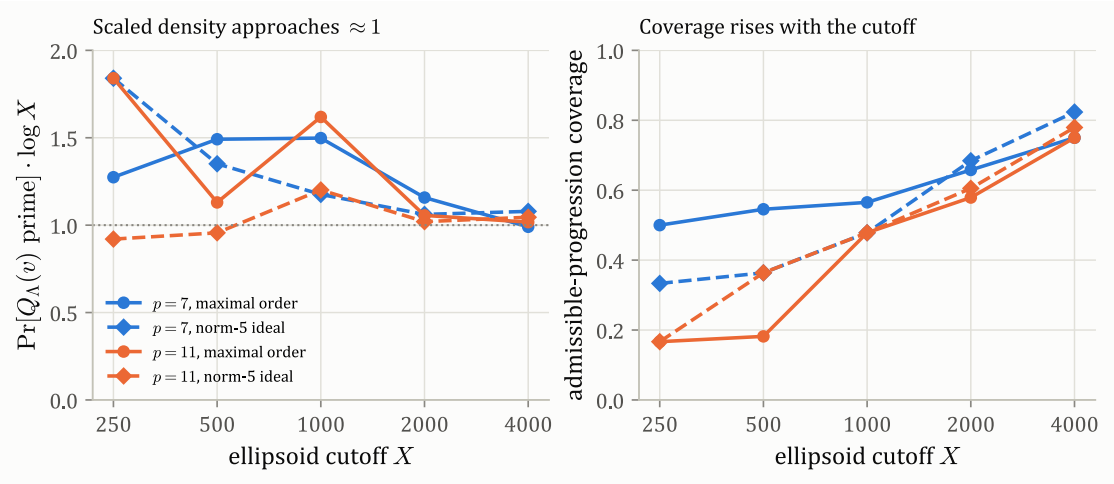

Figure 3: EMPIRICAL: p in {7,11}, ell = 3, 250 <= X <= 4000 Cutoff dependence of the sampler statistics for the four order/ideal cases. Left: scaled prime density; right: admissible-progression coverage. Data: the five files `measure_quaternary_prime_sampler_p11_ells3_x{250,500,1000,2000,4000}_20260711.csv`.

PROVED These computations attempt to falsify the residue and density mechanism of Theorem 9 and fail to do so; they do **not** test Rouse's asymptotic crossover and do not prove a worst-case sampler.

7.2 The right-order-aware Deuring round trip

The algebraic side of any future oracle reduction needs a validated local dictionary between ideals and isogenies. The fixture realizes the special-order generators i and j as the automorphism $\iota : (x, y) \mapsto \left(-x, \sqrt{-1}y\right)$ and Frobenius on $E : y^2 = x^3 - x$ over $\mathbb{F}_{p^2}$, $p \equiv 3 \bmod 4$, converts a norm-ℓ left ideal to its torsion kernel, applies a Vélu step, and returns. Endpoint equality uses the key $\texttt{deuring\_key}(E) = \min\big(\text{curve\_key}(E), \text{curve\_key}\big(E^{(p)}\big)\big)$, where `curve_key` is the lexicographically least coefficient pair in the scaling orbit $(A, B) \mapsto (c^4 A, c^6 B)$: Deuring correspondence identifies maximal-order classes with curve classes **modulo the Galois action**, so the Frobenius quotient is the correct acceptance rule for a round trip that starts from an abstract maximal order. CITED Distinct quaternion ideal lattices can be isospectral (Example 3.1 of [4]), so no acceptance test in the fixture relies on theta-series prefixes.

EMPIRICAL: p = 11, ell = 3, 30 independent seeds At the validated fixture $p = 11$, $\ell = 3$ — where $\#E(\mathbb{F}_{11^2}) = 144$, the nonzero rational 3-torsion has eight points, and exactly four norm-3 neighbor ideals occur — all 30 independently seeded trials passed every check: the ideal–kernel–ideal round trip recovered the ideal; all four embedded neighbor right orders appeared, each with trace discriminant $11^2 = 121$; every dual product satisfied $I\bar{I} = 3\mathcal{O}$; and each dual two-step chain (degree product 9) returned the source `deuring_key` through the computed dual kernel. The Wilson 95% interval for the success probability is approximately $[0.886, 1]$. Each trial enumerated 80 ideal candidates and applied 2 Vélu steps; the batch recorded $\texttt{oracle\_queries} = 0$.

7.3 The five-size cost fit

EMPIRICAL: p in {11,23,47,59,71}, ell = 3, n = 5 One timed exhaustive round trip per characteristic took 0.193, 0.458, 1.663, 2.634, and 4.702 seconds respectively. An ordinary least-squares fit in log–log coordinates gives

$$T(p) = 0.002926 \cdot p^{1.6796} \text{ seconds}, \quad R^2 = 0.9798, \tag{8}$$

with classical 95% exponent interval $[1.2366, 2.1226]$ (Figure 4); the five stored residuals are part of the released data.

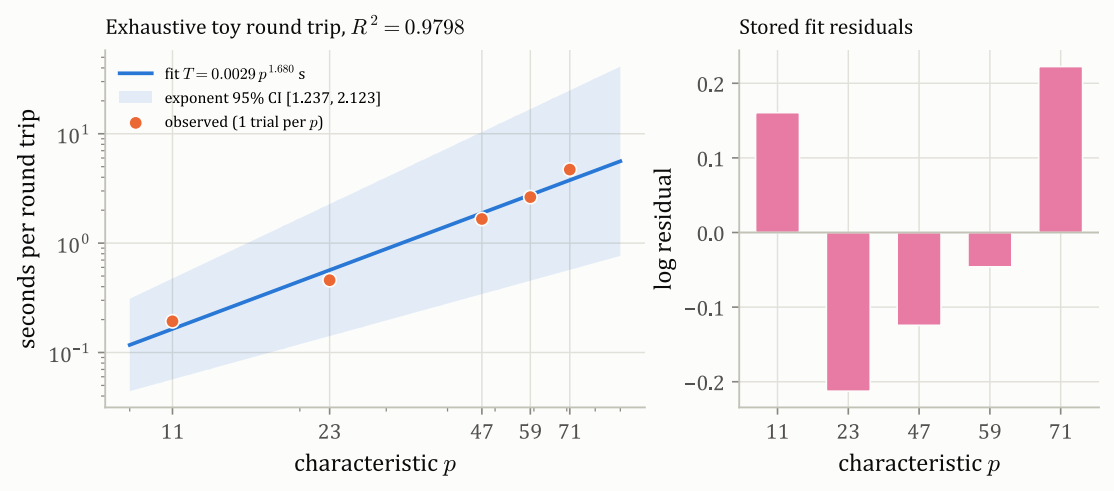

Figure 4: EMPIRICAL: p in {11,23,47,59,71}, ell = 3, n = 5 Left: observed seconds per exhaustive right-order-aware dual round trip against the power-law fit (shaded: the 95% exponent interval). Right: stored logarithmic residuals. The fit validates the fit **machinery** (checked against the exact synthetic relation $t = 2p^{3/2}$); it does not measure the cryptographic reduction (§8). Data: `fit_roundtrip_cost_ell3_p11-71_20260711.csv`.

Table 2: Empirical summary. Every value is stored in a seeded CSV named in the Reproducibility note; intervals are reported exactly as recorded in the research notes, and no row is extrapolated beyond its stated range.

Fixture	Parameters	Measured quantity	Value
Sampler grid	$p \le 31$, $\ell \in \{3,5\}$, $X = 3000$, 22 cases	scaled density $\Pr \cdot \log X$	0.832–1.815, mean 1.276
		admissible-progression coverage	0.10–0.778
		reciprocity violations	0 of 22 cases
		residue-lattice discriminant	$(8\ell)^6 p^2$ in all cases
Cutoff scan	$p \in \{7,11\}$, $\ell = 3$, $X \le 4000$	scaled density at $X = 4000$	0.991–1.079
		coverage at $X = 4000$	0.750–0.824
Round trip	$p = 11$, $\ell = 3$, 30 seeds	full-trial success (Wilson 95%)	30/30 ($[0.886, 1]$)
		right-order trace discriminant	$121 = 11^2$ (all trials)
		dual product $I\bar{I} = 3\mathcal{O}$	30/30
		two-step degree product	9 (all trials)
		oracle queries	0
Cost fit	$p \in \{11, \ldots, 71\}$, $\ell = 3$, $n = 5$	exponent (classical 95% CI)	1.6796 ($[1.2366, 2.1226]$)
		R^2; prefactor	0.9798; 2.93×10^{-3} s

8 Why no security-bit loss can yet be reported

Task 2 asks for a concrete loss function. The research defined the measurement before attempting it, and the definition is what shows the current data cannot supply the answer.

Definition 16 (concrete bit-loss proxy). PROVED For a reduction making $Q(p)$ oracle queries with success probability $s(p)$, define the time expansion

$$R_T(p) = \frac{T_{\text{reduction}}(p) + Q(p)\,T_{\text{oracle}}(p)}{T_{\text{oracle}}(p)} \tag{9}$$

and the bit-loss proxy

$$\Delta(p) = \log_2 Q(p) + \log_2(1/s(p)) + \log_2 \max(1, R_T(p)). \tag{10}$$

This proxy is an accounting convention, not a theorem equating runtime with advantage; reports must show its three summands separately, and a success probability enters only from at least 200 seeds with a Wilson interval excluding zero.

Proposition 17 (non-identifiability of the loss from the current data). PROVED The recorded round-trip path has $Q(p) = 0$: it validates the local Deuring dictionary and never calls a Problem-A or Problem-B oracle. Consequently $\Delta(p)$ is undefined on all current rows — its $\log_2 Q(p)$ and oracle-time terms require a genuine reduction with at least one query — and the fitted exponent 1.6796 describes exhaustive toy enumeration ($\mathbb{F}_{p^2}$ point enumeration, all norm-ℓ neighbor ideals, curve-scaling orbits), operations excluded from any cryptographic-size reduction. Assigning a numerical security-bit loss from these timings would conflate validation overhead with reduction loss.

PROVED A valid first bit-loss row must instrument one complete oracle reduction, record $Q(p) \geq 1$, separate oracle from non-oracle time, and estimate success from at least 200 seeds. The conditional Wesolowski route and the unconditional unrestricted-degree route of [3] must remain separate rows, because only the former targets prescribed smooth degree. CONJECTURE A third row would use the Theorem 9 sampler as an unconditional candidate replacement for D2, if its remaining proof obligations survive external audit; a counterexample to any obligation removes that row.

9 Limitations and open problems

Task 1 remains open, in both directions. The smooth-degree equivalence is neither proved unconditional here nor shown to require GRH; the audit establishes which four analytic inputs a resolution must replace (or prove necessary), not that they cannot be replaced. CITED The latest checked primary source (v2 of [3], checked with a literature search through 2026-06-29) leaves ℓ `-IsogenyPath` to `EndRing` conditional; a 2024 preprint claiming a factoring-oracle replacement is reported by the authors of [3] to contain a mistake, and no published correction was found.

Theorem 9 is a candidate proof. Its toy grid refutes nothing and confirms the residue and density mechanism, but the experiments run far below Rouse's crossover, test no worst-case instance, and the proof has not been externally refereed. PROVED It also does not touch the harder fixed-target equation: the quaternary representability input that powers Theorem 9 provably cannot power Step 9 by rejection (Proposition 14).

The frontier is constructive, not analytic-density. PROVED After the removal of D2, the remaining core is the D3+D4 prescribed-norm architecture plus the D1 special order it runs in; the single strongest target is the constructive fixed-target quaternary solver of §6. Another least-prime or average-density theorem would not move the frontier.

Open experimental threads. SG-03f: propagate the explicit dictionary across a non-backtracking second ideal and identify the terminal right order **without** an exhaustive curve-key lookup — the current lookup is the main exhaustive component of the toy cost. SG-05b: instrument one complete oracle reduction (at least 200 seeds) and evaluate the three summands of $\Delta(p)$ separately; until then Task 2 has a measured component but no loss function.

10 Conclusion

We set out to make the supersingular smooth-isogeny-path / endomorphism-ring equivalence unconditional and to price its concrete security loss. Neither goal is reached, and the paper says so plainly. What the effort produced is, we believe, the correct next-best object: a complete, source-checked map of exactly where GRH enters (four leaves, one removable RH appeal), the unconditional removal of one leaf by moving from binary to quaternary sampling, four sharp negative results that close the tempting shortcuts, a precisely isolated constructive target — the fixed-target quaternary norm solver on arbitrary maximal orders — whose solution would collapse the remaining three leaves to zero, and a validated toy experimental base whose own limitations are proved rather than glossed. The distance between "the equivalence is true under GRH" and "the equivalence is true" is now a single, named algorithmic problem.

Reproducibility

All experiments run from the problem directory `problems/P3.1-endomorphism-ring-equivalence/`. The sampler grid and cutoff scans are produced by `code/measure_quaternary_prime_sampler.py` (a `--smoke` mode runs $p \in \{7, 11\}$, $\ell = 3$ in under ten seconds); the recorded 22-case grid completed in 27.84 seconds after replacing symbolic matrix arithmetic by integer arithmetic. The round trip and timing rows are produced by `code/toy_deuring_roundtrip.py` and fitted by `code/fit_roundtrip_cost.py` (validated against the exact synthetic relation $t = 2p^{3/2}$). Data files used in this paper: `measure_quaternary_prime_sampler_p31_ells3-5_x3000_20260711.csv`, the five `measure_quaternary_prime_sampler_p11_ells3_x{250,...,4000}_20260711.csv` files, `toy_deuring_roundtrip_p11_ell3_trials30_20260711.csv`, the per-size `toy_deuring_roundtrip_p{11,23,47,59,71}_ell3_20260711.csv` files, and `fit_roundtrip_cost_ell3_p11-71_20260711.csv`. Shared libraries: `lib/quaternion.py` (exact ideal/order arithmetic), `lib/isogeny.py` (quadratic-extension curves, `deuring_key`, Vélu steps), `lib/finite_fields.py`. The final repository test suite passed 299 tests and 3 subtests on Python 3.13.4. The full audit table, dependency graph, and per-attempt post-mortems are `GRH_USAGE_MAP.md`, `TIGHTNESS_STATUS.md`, `REDUCTION_COST_SPEC.md`, `DEURING_ROUNDTRIP_SPEC.md`, and `attempts/A001–A006`. Every mathematical claim above carries one of the epistemic tags PROVED, CITED, CONDITIONAL: GRH, EMPIRICAL: range, CONJECTURE as used in the research log; untagged sentences are exposition, not claims.

References

[1] M. Deuring, "Die Typen der Multiplikatorenringe elliptischer Funktionenkörper," *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, vol. 14, pp. 197–272, 1941.

[2] B. Wesolowski, "The supersingular isogeny path and endomorphism ring problems are equivalent," in *62nd IEEE Symposium on Foundations of Computer Science (FOCS 2021)*, 2022, pp. 1100–1111.

[3] A. Herlédan Le Merdy and B. Wesolowski, "Unconditional foundations for supersingular isogeny-based cryptography," in *Theory of Cryptography Conference (TCC 2025)*, 2025, pp. 266–297.

[4] E. Z. Goren and J. R. Love, "On elements of prescribed norm in maximal orders of a quaternion algebra," *Canadian Journal of Mathematics*, vol. 77, no. 6, pp. 1938–1965, 2025.

[5] K. Eisenträger, S. Hallgren, K. E. Lauter, T. Morrison, and C. Petit, "Supersingular isogeny graphs and endomorphism rings: reductions and solutions," in *EUROCRYPT 2018, Part III*, in LNCS, vol. 10822. 2018, pp. 329–368.

[6] E. Assing, V. Blomer, and J. Li, "Uniform Titchmarsh divisor problems," *arXiv:2005.13915*, 2020.

[7] D. Jao, S. D. Miller, and R. Venkatesan, "Expander graphs based on GRH with an application to elliptic curve cryptography," *Journal of Number Theory*, vol. 129, no. 6, pp. 1491–1504, 2009.

[8] J. Thorner and A. Zaman, "An explicit bound for the least prime ideal in the Chebotarev density theorem," *Algebra & Number Theory*, vol. 11, no. 5, pp. 1135–1197, 2017.

[9] N. T. Sardari, "The least prime number represented by a binary quadratic form," *arXiv:1803.03218*, 2019.

[10] J. J. Ditchen, "On the average distribution of primes represented by binary quadratic forms," *arXiv:1312.1502*, 2018.

[11] J. Rouse, "Integers represented by positive-definite quadratic forms and Petersson inner products," *arXiv:1802.03437*, 2018.

[12] J. Rouse and K. Thompson, "Quaternary quadratic forms with prime discriminant," *arXiv:2206.00412*, 2022.

[13] M. A. Bennett, G. Martin, K. O'Bryant, and A. Rechnitzer, "Explicit bounds for primes in arithmetic progressions," *Illinois Journal of Mathematics*, vol. 62, no. 1–4, pp. 427–532, 2018.

[14] R. Kannan, "Minkowski's convex body theorem and integer programming," *Mathematics of Operations Research*, vol. 12, no. 3, pp. 415–440, 1987.