Constructing finite fields of prime degree from small cyclotomic cofactors and at fixed valuation after Bach, von zur Gathen, Lenstra, and Pila

23 November 2026

hexagon:2611.00987v2

Abstract

Added in version 2. OpenAI's preprint *Deterministic polynomial factorization over prime fields* (4 October 2026), with Shoup's reduction, constructs every finite field $\mathbb{F}_{p^r}$ in deterministic time polynomial in r and $\log p$ with no GRH assumption, and so settles the constructions of this paper in outcome. What remains this paper's own is its routes and their bounds: a cofactor construction with no analytic input, constructions whose only analytic input is rational Dirichlet GRH, average results, and the proved limits of the other routes.

Bach, von zur Gathen and Lenstra obtained nonresidues from smooth cyclotomic values. Pila used Frobenius on abelian varieties to construct roots of unity of fixed prime order. We use these ideas to construct a monic irreducible polynomial of odd prime degree over a prime field in two ways. The first is algebraic. Its bit cost is polynomial in the degree, the logarithm of the characteristic, a prime cutoff and the remaining cyclotomic cofactor. For degree r, every prime $\ell \neq r$ up to the cutoff that divides the cyclotomic value and satisfies $\ell \not\equiv 1 \pmod{r}$ can be removed to its full valuation without a root of unity as input. The second uses Alpöge's theorem on rational Dirichlet L-functions. For each fixed degree and fixed valuation of the cyclotomic value at that degree, it runs in time polynomial in the logarithm of the characteristic. For fixed odd prime degree r and fixed t, the same analytic input also removes the full powers of odd primes $\ell \neq r$ up to the cutoff with $v_r(\ell - 1) \leq t$. In degree three with $t = 1$, this includes all eligible primes up to the cutoff except those congruent to one modulo nine. The exponent in this extension may depend on r, t. Tang obtains the field presentation for every degree and characteristic using OpenAI's prime-field factorization algorithm and the Cureton–Thorner Dedekind zeta estimate. The constructions here give different proofs on the stated input classes. We also bound the exceptional characteristics for short auxiliary-prime and cubic searches. GPT-6 Astra (ChatGPT Pro mode) produced the mathematics and text. Claude Opus 5.5 at effort high made targeted checks; GPT-6.2 Astra at reasoning effort ultra (Codex) made targeted checks and edited this text. No human has checked it.

1 Statements, credit, and scope

Given a prime p and a prime degree r, how can one construct an irreducible polynomial of degree r over $\mathbb{F}_p$? Such a polynomial supplies the multiplication law for the field of order p^r. Shoup (1990, Theorem 3.1) reduced the construction to polynomial factorization over $\mathbb{F}_p$. With the uniform prime-field theorem of OpenAI (2026a, Theorem 1.1), this already gives a deterministic bound polynomial in r and $\log p$; that theorem uses its companion Hecke zero-free estimate. Tang (2026a, Theorem 1.1 and Proposition 4.1) gives the field presentation explicitly for arbitrary prime powers, using a different analytic estimate for the auxiliary primes. Our question is which inputs admit a construction by cyclotomic algebra alone and which admit one using the Riemann hypothesis for rational Dirichlet characters as the sole new analytic input.

Write A for the assertion that every nontrivial zero of every Dirichlet L-function has real part $1/2$, including the Riemann zeta function. We take Alpöge's theorem in precisely this form (Alpöge, 2026, p. 1). Statements described as consequences of A require no additional Riemann hypothesis. Alpöge's announced theorem is an input to this paper. The distinction between this input and RH for Hecke L-functions over other number fields is essential.

We separately denote by CT the uniform zero-free statement asserted in Cureton and Thorner's 2026 preprint (Cureton and Thorner, 2026, Theorem 1.1, p. 3): every Dedekind zeta function has no nontrivial zero with real part greater than $31/32$. Version 2, Hexagon 2610.00200v2, required a replacement of the proof of Lemma 8.1; Tang posted that replacement (Tang, 2026a,b, Remark 1.3, p. 2). We mark conclusions that accept CT explicitly. Neither that preprint theorem nor OpenAI's prime-field theorem is independently proved here. They are distinct inputs to Tang's route; A remains the sole new analytic input to our fixed-valuation construction.

The constructions are known in outcome (added in version 2). OpenAI (2026a, Theorem 1.1, p. 2), dated 4 October 2026, gives a uniform deterministic algorithm which, from a prime p in binary and a nonzero polynomial $f \in \mathbb{F}_p[x]$ of degree n given by its dense coefficient list, computes the leading coefficient of f and its distinct monic irreducible factors with their multiplicities in

$$O(((n+1)\lceil \log_2 p \rceil)^{10^{12}})$$

bit operations. It uses no randomness, no integer-factorization or primitive-root oracle, and no GRH assumption. Its analytic input is the zero-free strip $\Re s > 1 - 10^{-6}$ for finite-order Hecke L-functions over cyclotomic fields containing μ_{12}, uniform in conductor and height, which it quotes from OpenAI (2026b, Theorem 1.2, p. 3). From that strip it derives, for every prime $q \le n$ and every prime $p > B^{200000}$ with $B = 20 + (n+1)(\lceil \log_2 p \rceil + 1)$, a prime $\ell \neq p$ such that

$$\ell \equiv 1 \pmod{12q}, \qquad p^{(\ell-1)/q} \not\equiv 1 \pmod{\ell}, \qquad \ell \le c_0 B^{20000000},$$

with an absolute constant c_0 (OpenAI, 2026a, Proposition 11.3, p. 46). Shoup's reduction of the construction of irreducible polynomials to factorization over the prime field (Shoup, 1990, Theorem 3.1, author p. 2) then gives, for every prime p and every degree r, a monic irreducible polynomial of degree r over $\mathbb{F}_p$ in deterministic time polynomial in r and $\log p$.

This settles the constructions of this paper in outcome. The fields of Theorems 1.1, 1.2 and 1.5 and of Corollary 1.4 are constructed by that algorithm for all primes $p \neq r$, with no condition on a cofactor or on a valuation, and with an exponent independent of r. Theorem 1.7 converts each such field into a unit of maximal r-power order with the Teichmüller equivariance; the units U of Theorems 1.1, 1.2 and 1.3 are powers of it. The remaining outputs, namely the r-th roots of Theorem 1.2, the residue map of Theorem 1.3, and the root and the element of order nine of Theorem 1.5, are obtained by factoring explicit polynomials over $\mathbb{F}_p$ and over presented extensions of it, where factorization reduces to the prime field. With that reduction, the same theorem gives a polynomial bound for the factorization problem of Theorem 1.10, whose bound is quasipolynomial and uses further hypotheses.

What remains this paper's own is the routes and what they prove.

1. Theorem 1.1 and its general form, Theorem 5.1, construct the field and the unit with no analytic input of any kind, in time polynomial in r, $\log p$, the listed primes and the numerical residual cofactor. The case $t = 0$ of Theorem 1.2 and the interconversion of Theorem 1.7 are equally free of analytic input. The factorization theorem rests on a zero-free strip; these proofs use none.

2. Theorems 1.2, 1.3 and 1.5 and Corollary 1.4 have A, Dirichlet GRH over $\mathbb{Q}$, as their only analytic input, in place of a zero-free strip for Hecke L-functions. Their bounds are

polynomial in $\log p$ and the stated numerical quantities, with exponents that depend on the fixed parameters (r, t) or (r, b) and are absolute in the cubic case.

3. Theorem 1.6, Theorem 18.5, Theorem 1.10 and Table 1 record which zero statements the classical arguments use beyond A. The equivalence in Theorem 1.6 and its bound $\ell \ll r^2 \log^2(rp)$ are not given by the factorization theorem as stated. Nor is the bound of Theorem 1.8, which is uniform in the prime $r \mid p-1$, whereas factoring a polynomial of degree r takes time polynomial in r.

4. The average results, Theorems 1.11, 1.12, 1.13 and 1.14, bound from A alone the set of characteristics at which a short search fails.

5. Part II identifies, route by route, the additional inputs, exact reformulations and limits of the analytic, elliptic, class-field and Jacobi-sum constructions.

The uniform problem that stays open is the one these methods face: to construct $\mathbb{F}_{p^r}$ in time polynomial in r and $\log p$, with an exponent independent of the valuation $v_r(p^{\mathrm{ord}_r(p)} - 1)$, from A alone or with no analytic input. No result of version 1 is withdrawn.

All time bounds count bit operations. A presented field means either $\mathbb{F}_p[X]/(h)$ with h irreducible, or a basis and its multiplication table. These representations are interconvertible in polynomial time (Lenstra, 1991, Theorem (1.1), pp. 329–330). Constructing a presentation from an integer promised to be a prime power is a different input problem from factoring a polynomial whose coefficient field is already presented. Throughout, log denotes the natural logarithm.

For a positive integer n, Φ_n is its cyclotomic polynomial, and $v_\ell(n)$ is the exponent of the prime ℓ in n. A cofactor construction and a fixed-valuation construction address the remaining inputs. Bach, von zur Gathen and Lenstra construct nonresidues from a smooth cyclotomic value, given a presented field and primitive prime-order roots for every prime dividing that value (Bach et al., 2001, Theorem 3, printed p. 7; proof, pp. 20–24). Their cost is polynomial in the largest such prime and $d \log p$. With the notation of Theorem 1.1 below, for the numerical cofactor $M_2 = \Phi_d(p)/(2^{v_2(\Phi_d(p))} r^a)$, that largest prime is at most $\max(2, r, M_2)$; their theorem already gives a polynomial numerical-cofactor bound with those additional inputs. Their Theorem 2, printed p. 7 (proof, pp. 16–20), constructs the roots when irreducible polynomials of the specified prime degrees dividing $\ell - 1$ are given. Their Theorem 1, printed p. 6, keeps those construction inputs in the factorization reduction.

The following theorem constructs the degree-r field from p, r and a numerical cutoff alone. It uses the separating elements, rational-function count and odd-prime commutator of Bach et al. (2001, Proposition (5.2), pp. 20–21; Sections 5.7–5.11, pp. 22–23), together with simultaneous arithmetic in Lenstra's formal cyclotomic algebras. Primes not congruent to one modulo r can be removed by a full-algebra norm. The congruence condition already appears in the smooth-$p-1$ factoring theorem of Ivanyos et al. (2012, Theorem 1.5, author p. 6, journal p. 498). Section 5 proves the precise cyclotomic-value bound and compares the inputs statement by statement.

Theorem 1.1. *Let $p \neq r$ be primes, with r odd, and let $L \geq 2$ be an integer. Put*

$$d = \mathrm{ord}_r(p), \qquad a = v_r(p^d - 1), \qquad N = \Phi_d(p).$$

Let $\mathcal{T}_L$ consist of the odd primes $\ell \leq L$ such that $\ell \mid N$, $\ell \neq r$, and $r \nmid \ell - 1$, and define

$$C_L = \frac{N}{2^{v_2(N)} r^a \prod_{\ell \in \mathcal{T}_L} \ell^{v_\ell(N)}}.$$

There is an unconditional deterministic algorithm which, from the binary integers p, r, L, constructs a monic irreducible polynomial of degree r over $\mathbb{F}_p$ in

$$(r + L + C_L + \log p)^c$$

bit operations, for one absolute constant c. In the formal algebra $\mathcal{D} = \mathbb{F}_p[Z]/\Phi_r(Z)$, it also constructs an element U satisfying

$$U^{r^{a-1}} = Z, \qquad \sigma_b(U) = U^{\omega_a(b)} \quad (b \in \mathbb{F}_r^\times),$$

whose image has exact order r^a in every field component. Here $\sigma_b(Z) = Z^b$, and $\omega_a(b)$ is the unique lift of b modulo r^a whose order divides $r - 1$. No field presentation, cyclotomic factor, root of unity, or factorization of C_L is an additional input. The parameters L and C_L enter through their numerical values.

The proof is in Section 5; its more general form is Theorem 5.1. For a fixed constant B, promises $L, C_L \le (r + 1 + \log p)^B$ give polynomial time in $r, \log p$, with no bound on a. This does not cover arbitrary residual cofactors. Taking $L = 2$ gives the two-primary cofactor special case, stated and proved as Theorem 5.3 below.

The same norm construction can remove more prime factors under A, provided the part of $\ell - 1$ supported on the fixed target prime has bounded valuation. The additional ingredient is a quantitative version of the fixed-degree CM residue-map construction: its arithmetic data may vary, and its cost is polynomial in their numerical sizes.

Theorem 1.2 (A cofactor bound with bounded congruence valuation)**.** *Fix an odd prime r and an integer $t \ge 0$. For an input prime $p \ne r$ and an integer $L \ge 2$, put*

$$d = \mathrm{ord}_r(p), \quad a = v_r(p^d - 1), \quad N = \Phi_d(p), \quad s = v_2(N),$$

$$\mathcal{T}_{r,t,L}(p) = \{\ell \le L : \ \ell \text{ odd prime}, \ \ell \ne r, \ \ell \mid N, \ v_r(\ell - 1) \le t\},$$

$$C_{r,t,L}(p) = \frac{N}{2^s r^a \prod_{\ell \in \mathcal{T}_{r,t,L}(p)} \ell^{v_\ell(N)}}.$$

Under rational Dirichlet GRH, a monic irreducible polynomial of degree r over $\mathbb{F}_p$ can be constructed deterministically in bit time

$$c_{r,t}\big(L + C_{r,t,L}(p) + 1 + \log p\big)^{C_{r,t}},$$

with effective constants depending only on the fixed pair (r, t). The algorithm computes the displayed prime set and valuations itself; no factorization of the residual cofactor is supplied.

It also constructs $U \in D = \mathbb{F}_p[Z]/\Phi_r(Z)$ such that

$$U^{r^{a-1}} = Z, \qquad \sigma_c(U) = U^{\omega_a(c)} \quad (c \in \mathbb{F}_r^*),$$

and U has exact order r^a in every field component. Here $\sigma_c(Z) = Z^c$ and $\omega_a(c)$ is the tame lift modulo r^a. The simultaneous root interface also holds: for every monic nonconstant $g \mid \Phi_r$ and every unit $z \in A_g = \mathbb{F}_p[Z]/g$ satisfying $z^{(p^d-1)/r} = 1$, an element $y \in A_g$ with $y^r = z$ is constructed within the same bound. The algebras are constructed from the input; no component field, root of unity, or number-field residue map is supplied. For $t = 0$ all these assertions hold unconditionally.

For each fixed r, t, this theorem solves Lenstra's construction problem from A alone whenever its displayed residual and cutoff are polynomial in $\log p$, even when $v_r(p^d - 1)$ is unbounded. In degree three with $t = 1$, it removes all small odd primes other than three that are not one

modulo nine; the required CM varieties always have dimension three. The proof in Section 16.1 keeps the fixed-degree dependence explicit and gives no exponent uniform in growing r, t.

The fixed-valuation result covers, from A alone, every fixed-parameter stratum of the field-construction problem. Pila used Frobenius polynomials of abelian varieties to find roots of unity of fixed prime order in prime fields (Pila, 1990, Theorem D, author p. 2; Section 5, pp. 14–18). We combine that idea with Mumford's theta equations (Mumford, 1967, Section 6, pp. 83–84) and Lenstra's cyclotomic descent (Lenstra, 1991, Theorems (3.1), (5.1), and (5.2), pp. 334–340). The root now has a prescribed prime-power order and lies in an explicit finite algebra. The prime degree and its valuation are fixed parameters; the characteristic varies without exception.

Boyer also develops this fixed-polynomial CM programme (Boyer, 2014, Chapter III, p. 50), allowing precomputation independent of the characteristic. His even-order cyclotomic cases retain the displayed congruence conditions and, in the broader forms, Conjectures 2.7 and 2.11 (Boyer, 2014, pp. 65–72). His discussion of further extensions identifies certification of numerically constructed CM models as an obstacle (Boyer, 2014, §3.4, p. 81). Here the imaginary quadratic field varies with the characteristic. The degree and height estimates, certified theta recognition, marked reduction, and complete addition atlas below establish the required uniform model-construction bound.

Theorem 1.3. *Assume* A*, Dirichlet GRH over* $\mathbb{Q}$*. Fix an odd prime* r *and an integer* $b \geq 2$*. Put*

$$g = r^{b-1}, \qquad K_b = \mathbb{Q}(\zeta_{r^b}), \qquad T_r = \mathbb{Q}(\zeta_r),$$

and let R_b *be the unique degree-*g *subfield of* K_b*. In the precomputation for* (r, b)*, choose an integral basis of* R_b *by a fixed deterministic algorithm. In order of increasing maximum absolute coefficient, breaking ties lexicographically, choose the first integral linear combination* θ *of degree* g*. Write* $f \in \mathbb{Z}[X]$ *for its minimal polynomial. This fixes the index* $[\mathcal{O}_{R_b} : \mathbb{Z}[\theta]]$ *as a function of* (r, b)*. There is a deterministic algorithm, uniform in the prime* p*, with bit cost* $C_{r,b}(1 + \log p)^{c_{r,b}}$*, which applies when*

$$p \neq r, \qquad d = \operatorname{ord}_r(p), \qquad v_r(p^d - 1) \geq b.$$

It computes a unital $\mathbb{F}_p$*-algebra homomorphism* $\mathcal{O}_{R_b}/p\mathcal{O}_{R_b} \longrightarrow \mathbb{F}_p$*, and an element*

$$U \in \mathcal{D}_p := \mathbb{F}_p[Z]/(\Phi_r(Z))$$

having order exactly r^b *in every field component of* $\mathcal{D}_p$*, with* $U^{r^{b-1}} = Z$*. Moreover, for each* $c \in (\mathbb{Z}/r\mathbb{Z})^\times$*,*

$$\sigma_c(U) = U^{\omega_b(c)}, \qquad \sigma_c(Z) = Z^c, \tag{1}$$

where $\omega_b(c)$ *is the unique lift modulo* r^b *of* c *whose order divides* $r - 1$*. The constants are effective for each fixed pair* (r, b)*. If* $v_r(p^d - 1) = b$*, each component of* U *generates the entire* r*-primary subgroup of its multiplicative group. The algorithm does not require a factor of* Φ_r *over* $\mathbb{F}_p$*.*

The theorem computes a residue embedding of a fixed real cyclotomic field, then turns it into a compatible prime-power root in the formal coefficient algebra. This extends Pila's fixed-prime-order programme (Pila, 1990, Theorem D, author p. 2; printed p. 746) to prescribed wild layers and all tame residue degrees; the construction needed for Lenstra's problem follows when that layer is maximal.

Corollary 1.4. *Assume* A*, and fix an odd prime r and $b \geq 2$. There is a deterministic algorithm which, for every prime $p \neq r$ satisfying*

$$v_r(p^{\mathrm{ord}_r(p)} - 1) = b,$$

constructs a monic irreducible polynomial of degree r over $\mathbb{F}_p$, and hence a presentation of $\mathbb{F}_{p^r}$, in bit time $C_{r,b}(1 + \log p)^{c_{r,b}}$.

In words, fixing the prime degree and the valuation removes the exceptional characteristics from the construction. This answers that stratum of Lenstra's finite-field construction problem using only the rational Dirichlet Riemann hypotheses. The earlier small-order cases are covered by Theorem 1.8 and Theorem 1.7. A bound polynomial in r and $\log p$, with an exponent independent of the valuation, is not proved here by this route. Such a bound is known in outcome from OpenAI (2026a, Theorem 1.1, p. 2) with Shoup's reduction, as the version-2 note above records; from A alone it is open. The explicit CM dimension is r^{b-1}, and its theta coordinates grow with that dimension.

The proof chooses an imaginary quadratic field with conductor $O(\log^2 p)$ in which p splits. Its compositum with the fixed real wild cyclotomic field has a primitive CM type and an explicit principal polarization. Exact theta coordinates provide equations and a complete addition atlas. Reducing the marked model over all residue fields simultaneously lets Pila's algorithm compute a Frobenius polynomial without selecting a residue-field factor. The translated CM supports separate the real embeddings, and the cyclotomic tensor product turns one such embedding into the required equivariant root. The proof and its dimension-dependent bit bound are in Section 14. The CM algebraicity and reciprocity theorems, the Colmez and Pazuki height comparisons, and deterministic algebraic-number recognition carry the intervening steps; their hypotheses are stated at the point of use.

The smallest hard case gives a particularly short output polynomial. It is useful both as an example of the general result and as an explicit construction in any presented quadratic field.

Theorem 1.5 (A cubic construction after Pila)**.** *Assume* A*, rational Dirichlet GRH. There is a uniform deterministic algorithm which, on input a prime $p \equiv -1 \pmod 9$, computes:*

1. *a root $u \in \mathbb{F}_p$ of $f(X) = X^3 - 3X + 1$;*
2. *an element z of exact order 9 in $\mathbb{F}_p[W]/(W^2 + 3)$, or in any given presentation of $\mathbb{F}_{p^2}$.*

Its deterministic bit cost is $O((1 + \log p)^{C_*})$, for an absolute effectively computable constant C_*, independent of p. If $v_3(p + 1) = 2$, it also outputs the explicit monic irreducible polynomial

$$H(Y) = Y^3 - 3Y - u \in \mathbb{F}_p[Y],$$

and hence a presentation of $\mathbb{F}_{p^3}$, within the same bound. The only zero hypotheses used are for Dirichlet L-functions. The auxiliary-prime step uses conductors dividing $4p$. One may use the sharper CM-height bound with conductors at most $9q$, where $q = O((1 + \log p)^2)$, or the unconditional polynomial height bound of Proposition 8.2.

Thus a prime in the two residue classes $8, 17 \pmod{27}$ admits a deterministic cubic-extension construction from A alone. These include $p = 17$ (Example 13.1 and Section 21.5). The theorem settles this branch of Lenstra's finite-field construction problem. The preceding fixed-parameter theorem covers the other fixed valuations and prime degrees. These are complexity theorems with large fixed exponents.

For $r = 3$, the general theorem gives every fixed three-adic valuation. The explicit trace polynomial and the distinction between a nonmaximal and a maximal root of unity are developed in Section 15.

Table 1: Classical consequences, with the analytic family and the strength of the conclusion kept explicit. A statement marked "from A" retains any displayed algebraic input promises.

Problem	Result and algorithmic meaning	Required zero statement	Source and status
Miller test	All bases below $2(\log n)^2$ suffice; deterministic polynomial bit time. A perfect-power precheck permits the short proof given below.	Quadratic Dirichlet characters suffice.	From A. Schoof (2008, Thm. 2.2, pp. 105–106); Miller (1976, Thm. 2, p. 301).
Prime-field nonresidue	Given prime $r \mid p-1$, find an rth nonresidue by a scan of $O(\log^2 p)$ candidates.	Dirichlet characters of order r modulo p.	From A, uniformly in r. Lamzouri et al. (2015, pp. 1–2).
Primitive root with factored $p-1$	Construct a generator of $\mathbb{F}_p^*$ in polynomial time from the complete factorization of $p-1$.	The preceding prime-order Dirichlet characters.	From A with the given factorization; the construction below proves the transfer.
Least primitive root	$g(p) \ll t^4(1+\log t)^4 \log^2 p \ll \log^6 p$, $t = \omega(p-1)$. This gives a short candidate list.	Rational Dirichlet characters modulo p.	From A. Shoup (1992, Thm. 3, author p. 4; Prop. 9, p. 10; proof p. 13).
Solovay–Strassen	One Euler–Jacobi trial has composite-input error at most $1/2$; independent repetitions amplify it.	None for the randomized test. A separate derandomization uses a Dirichlet character annihilating its proper liar subgroup.	Randomized result already unconditional. Solovay and Strassen (1977, pp. 84–85); correction p. 118.
Order from factored exponent	Given a factored multiple M of an element's order, repeated prime removal computes the order.	None.	Already unconditional; neither M nor its factorization is provided by this statement.
Integer-modulus order	Order computation and integer factoring are polynomial-time Turing equivalent under the stated RH.	Rational Dirichlet L-functions in the small-nonresidue step.	Equivalence from A; no deterministic polynomial-time factoring algorithm follows. Miller (1976, Thm. 4, pp. 311–312; appendix pp. 315–316).
Artin primitive roots	Hooley's positive-density asymptotic for a fixed integer $a \neq -1$ that is not a square.	$\mathrm{RH}(\zeta_{\mathbb{Q}(\zeta_k, a^{1/k})})$ for squarefree k; these fields are generally nonabelian.	The known proof retains this family beyond A. Hooley (1967, translation pp. 40–45, 47).
Least prime in a progression	For $q > 3$ and $\gcd(a,q) = 1$, $P(a,q) \leq (\varphi(q)\log q)^2$. This is a numerical bound, not polynomial enumeration in $\log q$.	Rational Dirichlet characters modulo q.	From A. Lamzouri et al. (2015, Cor. 1.2, corrected author p. 3).
Chebyshev limiting distribution	Normalized progression errors have a limiting logarithmic distribution against bounded continuous test functions.	Rational Dirichlet characters modulo q.	From A. Rubinstein and Sarnak (1994, Thm. 1.1, p. 174).

Problem	Result and algorithmic meaning	Required zero statement	Source and status
Strict prime races	The standard strict ordering densities, symmetry formulas, and numerical biases require additional control of zero ordinates.	Dirichlet RH plus the grand simplicity hypothesis, a rational independence condition.	A does not supply that independence. Rubinstein and Sarnak (1994, pp. 176–177).
Ideal-class generators	Prime ideals of norm at most $12(\log \lvert D_K \rvert)^2$ generate $\mathrm{Cl}(K)$; this alone supplies no relation-search runtime.	Ideal-class Hecke $L_K(s,\chi)$, including ζ_K in Bach's argument. RH for the Hilbert class field's zeta is sufficient.	The quoted sharp bound retains Hecke RH. Accepting CT gives exponent 32 in Proposition 17.1. Belabas et al. (2008, author pp. 2–3) identify Bach's zero sums.
Class-group random walks	A sufficiently long random walk hits any specified subset with at least half its uniform probability. Degree and the exponent $B > 2$ are fixed; the modulus parameter is sufficiently large.	Ideal-class Hecke characters, including the trivial character in the spectral estimate.	A covers quadratic ζ_K; generally it leaves nontrivial class characters. Jao et al. (2009, Cor. 1.3, author p. 4; proof pp. 9–10). Uniform exponent-vector sampling is a different claim; see Section 17.
Schoof point counting	The original theorem gives deterministic $O(\log^9 q)$ point counting for presented fields of characteristic other than $2, 3$.	No number-field RH; the finite-field estimate used is proved.	Already unconditional. Schoof (1985, pp. 483, 485–490).
ECPP	Certificate correctness is unconditional. The cited practical algorithm's runtime analysis is explicitly heuristic.	None for verification; prime and smoothness distributions enter the heuristic analysis.	No general proved runtime follows just by supplying A. Schoof (2008, Thm. 5.2, p. 121; pp. 122, 124–125).
Isogeny expanders	Spectral bounds on ray-class Cayley graphs, with fixed degree for uniform rapid mixing. The elliptic application fixes the finite field, point count and endomorphism level.	Finite-order Hecke characters of the narrow ray class group over K; imaginary quadratic K in the elliptic application.	$K = \mathbb{Q}$ follows from A; the elliptic case generally retains Hecke RH. Jao et al. (2009, Thms. 1.1, 1.5–1.6; proof p. 9).
CM/CRT split-prime selection	Equation (21) and Lemma 2 control rational primes splitting completely in the ring class field. Only this stage is classified.	RH for the ring class field's Dedekind zeta function.	Generally remains conditional. Sutherland (2011, corrected author pp. 21–22).
Hilbert Nullstellensatz	Complex feasibility belongs to AM under the number-field prime-density input; this is not a deterministic-time conclusion.	ζ_K for a primitive-element field containing one algebraic solution; no normal closure is needed for this estimate.	The A-only transfer leaves this input. Accepting CT gives AM; Proposition 17.2. Koiran (1996a, Thm. 9, p. 282), with the erratum (Koiran, 1997).

Theorem 1.6 (The pure-field hypothesis for the classical construction)**.** *Let p and r be distinct primes, with r odd, and put $F = \mathbb{Q}(p^{1/r})$, $K = \mathbb{Q}(\zeta_r)$, and $E = K(p^{1/r})$. Under A,* $\mathrm{RH}(\zeta_F)$ *and* $\mathrm{RH}(\zeta_E)$ *are equivalent. Either assertion implies a prime $\ell \neq p, r$ such that*

$$\ell \equiv 1 \pmod{r}, \qquad p^{(\ell-1)/r} \not\equiv 1 \pmod{\ell},$$

with $\ell \ll r^2 \log^2(rp)$. A completely explicit, less sharp bound is

$$\ell \leq \left(4r(2r-3)\log r + 4(r-1)^2 \log p + \frac{5}{2}r(r-1) + 5\right)^2. \tag{2}$$

Given such a prime, one can construct $\mathbb{F}_{p^r}$ in deterministic time polynomial in ℓ and $\log p$*. Thus the displayed analytic assumption gives time polynomial in r and* $\log p$*.*

The construction is Adleman and Lenstra's Gaussian-period construction (Adleman and Lenstra, 1986, Algorithm A, pp. 350–351); the numerical prime bound uses Bach and Sorenson (Bach and Sorenson, 1996, Theorem 5.1, p. 1729). The sharper order of magnitude is the classical least-nonsplit-prime bound of Murty, as recorded by Zaman (2018, p. 1, equation (1.2)); Theorem 18.5 proves the required specialization. What is made explicit here is the single pure-field zeta function remaining after the abelian factors have been covered by A. This is a sufficient hypothesis for this construction, not a necessary condition for every algorithm constructing the same field. The algorithm of OpenAI (2026a) is a case in point: its auxiliary primes, of size polynomial in the input length, come from a zero-free strip of fixed width for Hecke L-functions (OpenAI, 2026a, Proposition 11.3, p. 46), without $\mathrm{RH}(\zeta_F)$.

Theorem 1.7 (Lenstra's Teichmüller construction, in prime degree)**.** *Let $p \neq r$ be primes, with r odd, and set*

$$s = \mathrm{ord}_r(p), \quad a = v_r(p^s - 1), \quad R = r^a, \quad D = \mathbb{F}_p[T]/\Phi_r(T), \quad \zeta = T \bmod \Phi_r.$$

For $b \in \mathbb{F}_r^\times$, let ρ_b be the automorphism $\rho_b(\zeta) = \zeta^b$, and let $\omega_a(b)$ denote the unique lift of b modulo r^a whose order divides $r-1$. The following two inputs are interconvertible, without an analytic assumption, in deterministic time polynomial in r and $\log p$*:*

1. *a presentation of $\mathbb{F}_{p^r}$;*
2. *an element $\eta \in D^\times$ of order R such that $\rho_b(\eta) = \eta^{\omega_a(b)}$ for every $b \in \mathbb{F}_r^\times$.*

If a presentation of $\mathbb{F}_{p^m}$ is given and $r \mid m$, it also determines such an η in time polynomial in m and $\log p$*.*

This is a prime-degree form of Lenstra's constructions (Lenstra, 1991, Theorems (3.1), (5.1), and (5.2), pp. 334–340). The point of writing the reductions out is to identify the input which is absent when only p and r are known. Reading a nonresidue from an already constructed field does not construct that field from its order.

Theorem 1.8 (Nonresidues in a presented extension)**.** *Let p be prime, let r be a prime divisor of $p-1$, and let a presentation of $\mathbb{F}_{p^m}$ be given. Under A, an rth nonresidue can be found deterministically in time polynomial in m and* $\log p$*. For $r = 2$ and even m no analytic assumption is needed.*

Here an rth nonresidue means a nonzero element outside $(\mathbb{F}_{p^m}^\times)^r$. The prime-field input is the small-character nonresidue theorem. The extension step is linear algebra in Frobenius eigenspaces, as in Bhargava, Ivanyos, Mittal, and Saxena's work. The proof below includes a nonvanishing choice of eigenspace vector.

The sufficient GRH hypothesis in classical construction results must also be read at its stated base field. Evdokimov's uniform construction theorem (Evdokimov, 1989, Theorem 1.1, pp. 107–108) uses the hypothesis explained on p. 105, which permits characters over abelian number fields, rather than only characters over $\mathbb{Q}$. His joint primitive-root and nonresidue interface therefore has a broader stated analytic premise than A. We do not infer that the same premise is necessary for every subroutine of that construction.

Theorem 1.9 (The two polynomial-time interfaces after Tang)**.** *Assume the algebraic prime-field theorem and bit bound of OpenAI (OpenAI, 2026a, Theorem 1.2, p. 3; Proposition 10.1, pp. 41–44), and* CT*, the uniform Dedekind zero-free statement of Cureton and Thorner's 2026 preprint, version 2 (Cureton and Thorner, 2026, Theorem 1.1, p. 3), with the revised Lemma 8.1 posted by Tang (Tang, 2026b). There are deterministic algorithms of polynomial bit complexity for both of the following tasks:*

1. *completely factor a nonzero dense polynomial over a presented finite field, including its multiplicities;*
2. *from a binary integer $q > 1$ promised to be a prime power, construct a power-basis presentation of the field of order q.*

The first interface reduces a degree-n polynomial to at most $\max(0, n-1)$ split squarefree prime-field queries of degree at most n. For $q = p^m$, the second uses $O(m^2(1 + \log_2 p))$ prime-field queries of degree at most m and polynomial additional bit cost.

These algorithms and the polynomial conclusion are Tang's (Tang, 2026a, Theorem 1.1 and Propositions 3.1, 4.1, pp. 2, 4–7). Section 21 gives the interface reductions in full and traces each auxiliary prime to just two Dedekind zeta functions. For $q = p^r$ the input length $\log_2 q = r \log_2 p$, so this construction includes every prime-degree field in Theorem 1.1, regardless of its cofactor. The two 2026 preprints assert unconditional theorems. The conclusion above accepts their named statements, including CT with Tang's revision of Lemma 8.1; it is not a verification of either whole preprint. Their analytic input differs from rational Dirichlet GRH. The uniform question in this paper asks for a derivation with A as its sole new analytic input. Accepting CT also gives the auxiliary-prime bound, so that question concerns the ingredients of a derivation, not logical independence.

Theorem 1.10 (A finite family for quasipolynomial factorization)**.** *Let a field $\mathbb{F}_{p^m}$ be presented, and let $0 \neq f \in \mathbb{F}_{p^m}[X]$ have degree $d \geq 2$. If $p \leq d$, its complete factorization can be computed deterministically in polynomial time in d and $\log(p^m)$. If $p > d$, put*

$$\mathcal{B} = \{r \leq d : r \text{ prime}, \quad r \nmid (p-1)m, \quad v_r(p^{\mathrm{ord}_r(p)} - 1) > 1\}.$$

Under A and RH for $\zeta_{\mathbb{Q}(p^{1/r})}$ for $r \in \mathcal{B}$, complete factorization has deterministic bit complexity $(d^{\log d} \log(p^m))^{O(1)}$. Only primes encountered by the algorithm need their additional zeta hypothesis. If $p > d$ and every prime $r \leq d$ either divides m or satisfies $v_r(p^{\mathrm{ord}_r(p)} - 1) = 1$, the latter bound holds without an analytic assumption.

This identifies a finite list of extra zeta hypotheses sufficient for the earlier quasipolynomial factorization algorithm. Prime-field nonresidues are covered by A, whereas missing extension fields can require the pure-field hypotheses displayed here. It refines the analytic accounting in Lenstra's construction problem, without giving a uniform polynomial bound for that earlier algorithm.

This finite family refines the earlier quasipolynomial algorithm; it is not the analytic input to Tang's polynomial theorem.

The next average statements use character orthogonality over the varying input prime. Their lower-search comparisons specialize Montgomery's construction, presented in Montgomery and

Vaughan (2007, p. 432, Exercise 9(a)–(e)). That source treats large least quadratic nonresidues; here the characters encode the auxiliary-prime and Tate-cubic searches, with their exact input restrictions.

Theorem 1.11 (Short auxiliary primes for almost all inputs)**.** *Assume A and fix an odd prime r. Let*

$$C_r > \frac{r-1}{2\log r}, \qquad B = C_r \log P \log\log P.$$

For all sufficiently large P, depending on r and C_r, all but

$$O_{r,C_r}\Big(\sqrt{P}\log\log P\Big)$$

primes $p \in (P/2, P]$ admit a prime $\ell \le B$, with $\ell \ne p, r$, such that $\ell \equiv 1 \pmod r$ and $p^{(\ell-1)/r} \not\equiv 1 \pmod \ell$. In particular this bounds the exceptions among inputs with $r \nmid p-1$ and $r^2 \mid p^{\mathrm{ord}_r(p)} - 1$. For $r = 3$ one may take $C_3 = 1$; the assertion remains asymptotic, and no explicit numerical threshold for P is asserted here.

On the successful inputs, deterministic enumeration and certification of such an ℓ, followed by the established auxiliary-prime reduction, constructs $\mathbb{F}_{p^r}$ in time polynomial in $\log p$ for the fixed r and C_r. An algorithm stopping at the displayed bound can report failure on the exceptional inputs; the theorem is not an all-input construction.

The averaging is over input primes. It makes the classical construction polynomial on a density-one set of inputs of each fixed degree, while leaving the prescribed exceptional input unresolved from A alone; in outcome that input is covered by OpenAI (2026a, Theorem 1.1, p. 2). The proof uses only rational Dirichlet characters.

Theorem 1.12 (Uniform auxiliary-prime menus)**.** *Assume* A*. There are absolute positive constants C, C_1, c_1, C_2 with the following properties. Let $P \ge 4$, let $r \ge 3$ be prime, and put*

$$B = Cr^2\log^2(Pr), \qquad k = \left\lceil \frac{\log P}{2\log r} \right\rceil, \qquad A_r(P) = 1 + \frac{\log r}{\log P} + \frac{\log\log(Pr)}{\log r}.$$

There are at least k primes $\ell \le B$ with $\ell \equiv 1 \pmod r$. Choose any k of them, forming a set $\mathcal{S}$. The number of primes $p \in (P/2, P]$, $p \ne r$, for which no $\ell \in \mathcal{S}$ satisfies

$$\ell \ne p, \qquad p^{(\ell-1)/r} \not\equiv 1 \pmod \ell$$

is at most $C_2\sqrt{P}\,A_r(P)$. This count includes input primes belonging to $\mathcal{S}$, and the same bound holds when the whole menu of progression primes up to B is used. For every fixed $A_0 > 0$, the exceptional proportion among all primes in the interval tends to zero uniformly for $r \le P^{A_0}$.

Let $H_r(P)$ count the hard primes in that interval, namely those with

$$d = \mathrm{ord}_r(p) > 1, \qquad v_r(p^d - 1) \ge 2.$$

If $P \ge C_1 r^2 \log^2(2r)$, then

$$H_r(P) \ge c_1 \frac{P}{r\log P}.$$

Among these hard primes, the proportion failing the same menu tends to zero uniformly for $r = o(\sqrt{P}/\log P)$. More precisely, for every positive function $\epsilon(P) \to 0$, the supremum of that proportion over odd primes $r \le \epsilon(P)\sqrt{P}/\log P$ tends to zero; an empty supremum is zero. Menu construction, testing, and the field construction after a successful test have bit cost polynomial in $r, \log P$, with an absolute exponent.

Only enough auxiliary primes to suppress the principal-character term are needed in this estimate. This gives a uniform average form of the classical auxiliary-prime construction, including a separately justified density statement within the hard inputs. It does not decide whether a prescribed exceptional characteristic has an auxiliary prime of this size. An auxiliary prime below a larger fixed power of $r \log p$ exists for every prime p above a fixed power of $r \log p$, by a zero-free strip for Hecke L-functions (OpenAI, 2026a, Proposition 11.3, p. 46).

For the cubic family, a second classical ingredient is the quadratic-value product estimate in Cilleruelo (2011, Lemmas 1, 3, and 4 and §2.4, pp. 1131–1134). The specialization below turns its large prime divisors into independent cyclic-cubic conductor primes. It gives the following explicit output bound for all prime congruence classes.

Theorem 1.13 (A logarithmic cubic search for almost all inputs)**.** *Assume A. Fix $C > 1/\log 3$ and, for sufficiently large P, put*

$$K = \lfloor C \log P \rfloor .$$

All but $O_C(\sqrt{P} \log\log P)$ primes $p \in (P/2, P]$ admit an integer $1 \le k \le K$, $k \not\equiv 3 \pmod p$, for which

$$Y^3 - (k+3)Y^2 + kY + 1$$

is irreducible over $\mathbb{F}_p$. In particular, the same exception bound holds among the hard cubic inputs $p \equiv -1 \pmod 9$.

Testing these $O_C(\log p)$ cubics deterministically constructs $\mathbb{F}_{p^3}$ in time polynomial in $\log p$ on the successful inputs. The search can report failure on the exceptional inputs; the assertion is not an all-input parameter bound.

A short explicit list of cubics therefore constructs cubic extensions for all but the stated exceptional primes. This is a direct average answer to Lenstra's construction problem: it avoids an auxiliary field on the successful inputs, while leaving the exceptional characteristic unresolved by this search bound.

Theorem 1.14 (A simultaneous auxiliary table for almost all characteristics)**.** *Assume A, and fix $a_0 > 0$. There is an absolute constant C_0 such that the following holds for all sufficiently large P (the threshold may depend on a_0). Let $1 \le D \le (\log P)^{a_0}$ be an integer. Except for*

$$O\left(\frac{\sqrt{P}\, D^3 \log^2(PD)}{\log P} \right) \tag{3}$$

primes $p \in (P/2, P]$, every prime $r \le D$ simultaneously has a rational prime ℓ satisfying

$$\ell \le C_0 r^2 \log^2(Pr), \qquad \ell \equiv 1 \pmod{12r}, \qquad p^{(\ell-1)/r} \not\equiv 1 \pmod \ell. \tag{4}$$

All such primes automatically satisfy $\ell \notin \{2, 3, p, r\}$. The constants in the O-bound are absolute; only the lower threshold for P depends on a_0.

One table of small auxiliary primes now works simultaneously for all prime degrees in the stated range. It supplies the classical Kummer and Gaussian-period inputs, and also the auxiliary entries in Tang's factoring interface, for almost all characteristics. The exception estimate is uniform in the degree cutoff, with the stated restriction on that cutoff.

Corollary 1.15 (Average polynomial-time interfaces after Tang)**.** *Assume A and the purely algebraic auxiliary-prime algorithm of OpenAI (2026a, Theorem 1.2, p. 3; Proposition 10.1, pp. 41–44). For every prime outside the exceptional set in Theorem 1.14, a single uniform deterministic procedure has the following guarantees.*

1. *Over any explicitly presented finite extension of characteristic p, it completely factors every nonzero dense polynomial of degree $n \le D$, with leading coefficient and multiplicities, in time polynomial in $n+1$ and the bit length of the field order.*

2. *Given an integer promised to equal $q = p^m$, with $1 \le m \le D$, it constructs a presentation of the field of order q in time polynomial in $\log q$.*

The polynomial exponents are absolute and the algorithms do not receive P or D as input. Their running time uses the input degree, not the global degree cutoff.

The average table transfers to both factoring in a known field and constructing a field from its order. The algebraic prime-field theorem of OpenAI (OpenAI, 2026a, Theorem 1.2, p. 3) remains an explicit input, while the analytic part uses only A; the resulting procedures can fail on the exceptional primes. With a zero-free strip for Hecke L-functions as its analytic part, OpenAI (2026a, Theorem 1.1, p. 2) has no exceptional primes.

Why this matters. A finite field is specified abstractly by its order, but computation in it requires an explicit presentation. Lenstra's construction problem asks for that presentation with a deterministic polynomial bit bound. In outcome it is answered by OpenAI (2026a, Theorem 1.1, p. 2) with Shoup's reduction, from a zero-free strip for Hecke L-functions. The results here show what is reached with less: no analytic input, or A alone. The cofactor theorem uses no analytic input and allows an unbounded valuation when the numerical residual cofactor is small, after removing the stated small-prime factors to their full valuations. Under A, the bounded-congruence variant removes a larger prime set at each fixed target degree; the fixed-valuation CM theorem covers every fixed degree and fixed valuation with no exceptional characteristic. These complement one another and separate the uniform problem with analytic input A alone, which is unresolved, from the prime-field consequences of Dirichlet GRH: Miller's cutoff and small nonresidues follow from rational characters, whereas Hooley's classical proof and the usual Kummer auxiliary-prime proof use further fields.

What is new here. The fields are not new in outcome, as the version-2 note above records; what is new is the routes to them, with their inputs and bounds. The cofactor bound comes from extracting roots simultaneously in a formal cyclotomic algebra. The rational-function count and the odd-prime commutator are due to Bach, von zur Gathen and Lenstra. A comparison along a cyclic orbit removes the selected prime powers without auxiliary roots. Alternatively, a norm descends the answer from a cyclotomic factor of degree prime to r. The fixed-valuation construction uses a varying imaginary quadratic field in the CM model. For a varying totally real Galois field of fixed degree, the residue-map bound charges the height of its integral multiplication table and its discriminant. It supplies the quotient for the bounded-valuation extension. The average estimates count the input characteristics where a short search can fail.

Organization and the remaining routes. The proofs of the construction, classification, pure-field identity, and average results form the first part. The later part, *Routes that do not give a uniform construction*, keeps the full proofs of the analytic, elliptic, class-field, Jacobi-sum, fixed-menu, and sparse-compression statements. Those results identify exact additional inputs or limitations of particular methods. They do not supply a uniform selector for the exceptional characteristic. One exists for every characteristic above a fixed power of the input length: the

auxiliary primes of OpenAI (2026a, Proposition 11.3, p. 46) come from a zero-free strip for Hecke L-functions. The routes of the later part seek a selector from A alone; they neither find one nor prove that none exists. The final section states the two uniform questions and the precise obstructions still left by these methods; both questions are to be read with the analytic input restricted to A.

2 Proof overview

This section explains the mechanisms and the role of each analytic or algebraic hypothesis before the detailed proofs.

The cofactor construction. The difficulty is to lift a known primitive r-th root to the full r-primary order without first factoring the cyclotomic algebra or constructing roots of unity for its other prime divisors. We use Bach, von zur Gathen and Lenstra's separating elements in a Kummer algebra and project them to the relevant cyclotomic torsion. If every projected coboundary were scalar, a rational function would take a fixed power value at more points than its numerator degree allows; the distinct values at zero and infinity exclude that possibility. Their odd-prime commutator then produces a zero divisor and hence a root. Every other selected primary part needs only a root of its prime order, because the first nontrivial primary transition and the cyclic product-one identity give a proper factor. The two-primary part uses only the known root -1. For an order-ℓ element, compare its cyclic shift with each of its powers. If this fails to split the algebra, the shift induces a nonidentity order-r automorphism of μ_ℓ. This forces $r \mid \ell - 1$. A second proof takes a norm after adjoining a root in a formal cyclotomic algebra. It also applies when a partial factor of Φ_ℓ of degree prime to r is given. Only one such auxiliary algebra is needed at a time. Arithmetic runs over the squarefree cyclotomic algebra, splitting only when a tested coefficient vanishes on some components. There are at most its dimension many final branches. The resulting roots are combined and projected to Lenstra's Teichmüller character, which gives the extension. No analytic hypothesis enters; the number of candidate shifts explains the numerical cofactor in the bound.

The bounded-congruence refinement. The full cyclotomic algebra for a prime $\ell \equiv 1 \pmod{r}$ has rank divisible by r, so its norm cannot directly descend an r-th root. Put $b = v_r(\ell-1)$. The real cyclic subfield of $\mathbb{Q}(\zeta_\ell)$ of degree r^b has an explicit integral Gaussian-period basis, and p splits completely in it. The CM construction yields a residue map with cost polynomial in the numerical multiplication-table height and discriminant when this degree is fixed. Period-value equations then select a cyclotomic quotient of rank $(\ell-1)/r^b$, which is prime to r; the preceding full norm argument applies unchanged. To make the varying field data legitimate, the proof bounds the sign element, lattice index, CM orbit degree, theta height and recognition precision explicitly, and treats coordinate denominators by projective reduction. Only the auxiliary quadratic-prime choice uses A. Bounding r, b bounds the CM dimension; it does not impose a bound on the valuation of ℓ in the cyclotomic value.

The fixed-valuation construction. The missing input is a residue embedding of a fixed real wild cyclotomic field. Its degree is $g = r^{b-1}$, and the proof keeps r, b fixed. Dirichlet GRH finds a small imaginary quadratic conductor for which the input prime splits. The product of that quadratic field with the fixed real field has explicit principal periods and a primitive CM type. Degree and height bounds make certified algebraic recognition of the theta coordinates polynomial in the numerical auxiliary conductor; the sharper height bound uses only Dirichlet characters because this CM field is abelian. Mumford's theta equations turn the recognized point into a projective abelian variety with a complete addition atlas. Potential good reduction

together with the full prime-to-characteristic marking keeps every residue component in the moduli open. Pila's algorithm then computes Frobenius polynomials over all those components. The prime supports of the conjugate CM Frobenius elements separate the real embeddings by polynomial greatest common divisors. A residue embedding gives the prescribed root of unity, and Lenstra's descent gives the extension when the valuation is maximal. The sextic case is the first instance, with an explicit cubic output. Fixed dimension is essential to Pila's exponent and the number of theta coordinates; the proof gives no uniform exponent as r, b grow.

The classification and pure-field identity. Miller's test and the prime-field nonresidue bound reduce to rational Dirichlet characters. A primitive-root algorithm also needs the factorization data stated in the table, and Hooley's density theorem uses zeta functions of generally nonabelian Kummer fields. For the pure-field identity, the normal closure of $X^r - p$ has Galois group $\mathbb{F}_r \rtimes \mathbb{F}_r^\times$. Its permutation and regular representations isolate one nonabelian constituent. Artin induction identifies its L-function with a nontrivial finite-order Hecke L-function over the cyclotomic field, so it is entire. The zeta identity then shows exactly what remains after A accounts for the abelian factors. The classical least-nonsplit-prime theorem gives the auxiliary-prime bound under that extra pure-field hypothesis. None of these identities turns a general Hecke character into a rational Dirichlet character.

Presented fields and algorithmic transfers. If the field is already presented, Frobenius eigenspaces supply extension-field nonresidues from the small prime-field witness; a basis search ensures the chosen eigenvector is nonzero. The quasipolynomial route calls for only the finite list of missing prime-degree constructions in its statement. Tang's two interfaces instead transfer OpenAI's algebraic prime-field theorem together with the separately named preprint input CT, including Tang's revised Lemma 8.1. Distinguishing these inputs is necessary: factoring in a known field and constructing a field from its order are different tasks.

The averages. The uniform auxiliary-prime question fixes the characteristic; the average proofs vary it. For candidate auxiliary primes, character orthogonality expresses simultaneous failure as a sum of rational Dirichlet characters. With m independent order-r characters, the principal contribution is reduced by r^{-m}, while the remaining errors depend on the conductor product M through $\log M$. Choosing $r^m \geq \sqrt{P}$ balances these two costs. For the logarithmic cubic search, the large prime divisors of $k^2 + 3k + 9$ give private ramified primes for a large subfamily of cyclic cubic fields. Their private primes prove the independence needed for the same orthogonality argument. Dirichlet GRH estimates the resulting character sums; the quadratic-value product estimate supplies the independent family. The bounds still allow exceptional input primes, as their statements explicitly say.

Part I
Constructions and their analytic inputs

3 The maximal Teichmüller unit

This section proves both directions of Theorem 1.7, so that a use of the field as an input cannot be mistaken for its construction.

Proof of Theorem 1.7. Each irreducible factor of Φ_r over $\mathbb{F}_p$ has degree $s = \mathrm{ord}_r(p)$. Thus D is a product of $(r-1)/s$ copies of $\mathbb{F}_{p^s}$. This description proves the properties used below; the algorithm never needs to factor Φ_r.

First let a field $L \cong \mathbb{F}_{p^r}$ be presented. Compute a normal-basis element α for $L/\mathbb{F}_p$ by Lenstra's deterministic algorithm, and put $C = L \otimes_{\mathbb{F}_p} D$. Let σ be p^s-power Frobenius on L and the identity on D. Since $s \mid r-1$, $\gcd(s,r) = 1$, and σ generates $\mathrm{Gal}(L/\mathbb{F}_p)$. In every field component of C, the elements $\sigma^i(\alpha)$, $0 \le i < r$, are a basis over $\mathbb{F}_{p^s}$: L and $\mathbb{F}_{p^s}$ are linearly disjoint. It follows that

$$\beta = \sum_{i=0}^{r-1} \zeta^i \sigma^i(\alpha)$$

is nonzero in every component, hence is a unit. A cyclic shift gives $\sigma(\beta) = \zeta^{-1}\beta$. On all of C, σ agrees with p^s-power Frobenius. Writing $w = (p^s-1)/R$ and $\delta = \beta^w$, we obtain

$$\delta^R = \zeta^{-1}, \qquad \delta^{rR} = 1.$$

Use the Teichmüller lifts modulo rR, and define

$$y = \prod_{b \in \mathbb{F}_r^\times} \rho_b(\delta)^{\omega_{a+1}(b)^{-1}},$$

where ρ_b acts on D and fixes L. The exponents are read modulo rR. Reindexing this product by $b \mapsto cb$ shows $\rho_c(y) = y^{\omega_{a+1}(c)}$. Also

$$y^R = \prod_b \zeta^{-b\omega_{a+1}(b)^{-1}} = \zeta^{-(r-1)} = \zeta.$$

Thus y has order rR, and $\eta = y^r$ has order R. Since $R \mid p^s - 1$, $\eta^{p^s} = \eta$. The fixed algebra of σ in C is D, so $\eta \in D$. Reduction of the equivariance congruence modulo R proves the required property of η.

Conversely, let such an $\eta \in D^\times$ be given. The automorphisms ρ_b are transitive on the field components of D, and exponentiation by $\omega_a(b)$ preserves order. Hence the image of η in every component has the same order, namely R. It is an rth nonresidue there, because R is the full r-part of $p^s - 1$. Over a field containing μ_r, a polynomial $Y^r - c$ for prime $r \neq p$ is irreducible when c is an rth nonresidue: adjoining one root splits the polynomial, and its Galois group embeds into the prime-order group μ_r, so a nontrivial extension has degree r. Consequently

$$C' = D[Y]/(Y^r - \eta)$$

is a product of fields of degree sr over $\mathbb{F}_p$.

Extend the automorphisms to C' by $\rho_b(Y) = Y^{\omega_{a+1}(b)}$. They preserve $Y^r = \eta$ and form a group action: Y has order rR, and the lifts are multiplicative modulo rR. The action is transitive on the field components, so its fixed algebra is a field. To compute its degree, extend scalars to an algebraic closure. There are $r(r-1)$ geometric points of C', and the action on them is free because its action on the primitive rth roots ζ is free. There are therefore r orbits.

Invariants commute with this scalar extension, being kernels of finite systems of linear equations. It follows that $(C')^{\mathbb{F}_r^\times}$ is a degree-r field over $\mathbb{F}_p$.

Both constructions involve algebras of dimension at most $r(r-1)$. The integers R and rR have $O(r \log p)$ bits. The lifts are computed by Hensel lifting the simple roots of $X^{r-1}-1$ modulo r; all exponentiations and linear systems have polynomial bit cost. Finally, if a field of degree m is given with $r \mid m$, its unique degree-r subfield is the kernel of $z \mapsto z^{p^r} - z$. This is an $\mathbb{F}_p$-linear computation of polynomial size. □

Corollary 3.1 (A presented finite ground field)**.** *Theorem 1.7 remains valid with the ground field $\mathbb{F}_p$ replaced by a presented finite field $k = \mathbb{F}_q$, where r is odd and prime to its characteristic, $s = \mathrm{ord}_r(q)$, and $R = r^{v_r(q^s-1)}$. The reductions have bit cost polynomial in r and $\log q$.*

Proof. Use qth-power Frobenius and linear algebra over k throughout that proof. The factors of Φ_r have degree s over k; the degree-r extension of k and its degree-s extension are linearly disjoint. The same normal-basis resolvent, multiplicative projection, and free action on geometric points therefore give both reductions. Each algebra has dimension at most $r(r-1)$ over k, and the exponents have $O(r \log q)$ bits. Arithmetic and linear algebra over the given presentation of k have polynomial bit cost. □

Corollary 3.2 (The easy valuation case)**.** *If $v_r(p^{\mathrm{ord}_r(p)} - 1) = 1$, the degree-$r$ field can be constructed without an analytic assumption.*

Proof. In this case $R = r$, and $\eta = \zeta \in D$ has order R and satisfies the equivariance condition modulo r. Apply Theorem 1.7. □

When $r \mid p-1$, the small prime-field nonresidue theorem under rational Dirichlet GRH (Theorem 1.8) supplies the input. Together with the unconditional case just proved, this leaves $r \nmid p-1$ and $r^2 \mid p^{\mathrm{ord}_r(p)} - 1$ for this approach. Theorem 1.7 states an equivalence of two output problems in this case; it supplies neither input from p, r alone.

4 Arithmetic algorithms over products of finite fields

Arithmetic algorithms over fields can be executed simultaneously over a finite product by splitting only the supports of their zero tests. This supplies the coefficient interface for the root-extraction and CM constructions below.

Proposition 4.1. *Let $h \in \mathbb{F}_p[T]$ be monic and squarefree of degree N, and put $A = \mathbb{F}_p[T]/(h)$. Consider a deterministic algorithm over finite fields whose coefficient-dependent operations are field arithmetic, division by tested nonzero elements, and equality tests. Suppose it uses at most U such operations and at most U bit operations on non-field data on every component field of A, when run on the reductions of a given input over A. Integer operations and loop bounds must be determined by the field size and previous test outcomes; they must not inspect a representation-specific encoding of a field element. Assume also that U is at least the number of input field elements and the bit length of the non-field input.*

There is a deterministic simulation, polynomial in $N, U, \log p$, that partitions h into coprime factors and computes the correct output over every resulting quotient. It does not require an irreducible factor of h or a selected root of h. When the outputs are integers or integer polynomials, the output on every factor is a single such value, valid for all of that factor's component fields.

Proof. First use distinct-degree factorization to partition h into products of irreducibles of a common degree e. This requires only modular powering and greatest common divisors, and has a deterministic polynomial cost. Every component in such a block then has the same known cardinality p^e.

Run the algorithm modulo a current factor v of h. Arithmetic operations are polynomial arithmetic modulo v. For a zero test on an element represented by $c(T)$, compute $w = \gcd(v, c)$. If $w = v$, the element vanishes on every component. If $w = 1$, it vanishes on none, and its inverse, when requested, is obtained by the extended Euclidean algorithm. If $0 < \deg w < \deg v$, split the state into its reductions modulo w and v/w, and continue the two appropriate branches. Squarefreeness ensures that these factors are coprime and that these cases exactly describe the componentwise zero test.

Every irreducible factor follows exactly the path of the original field algorithm, by induction on its operations. There are at most N final branches and at most $N - 1$ genuine splits. Charging an operation to every final branch below it bounds the number of simulated operations by NU. All polynomial arithmetic, greatest common divisors, inversions, and state reductions have polynomial cost in N and $\log p$. The stored state has polynomial size because the original run has at most U operations. This proves the claimed complexity. If an output is an integer, its computation depends on the common field size and the common test outcomes along that branch, so the output is identical in all its components. □

Proposition 4.2. *The preceding conclusion also holds for any finite $\mathbb{F}_p$-algebra A that is a product of finite fields and is given by structure constants in a basis of size N. The output partition is then a direct-product decomposition into explicitly represented ideals with their own identities. A monogenic presentation is unnecessary.*

Proof. All the required decompositions can be computed by linear algebra over $\mathbb{F}_p$. To separate the component degrees, let $F : x \mapsto x^p$. For $e = 1, \dots, N$, in the as yet unassigned part of the algebra, form the ideal J_e generated by $F^e(b) - b$, where b runs through a basis. A component of degree d survives in A/J_e exactly when $d \mid e$: Frobenius to the e-th power is the identity precisely in that case, and otherwise some generator is nonzero in that field. All proper divisors of e have already been assigned at earlier iterations. Thus the new nonzero quotient consists exactly of the degree-e components. To compute J_e, span the products of its generators with the algebra basis. Compute its annihilator by solving linear systems. Since the algebra is a product of fields, $A = J_e \oplus \mathrm{Ann}(J_e)$; decomposing 1 gives the two identities and realizes the quotient as a subalgebra. All Frobenius computations use repeated powering with $e \le N$.

On any current block B, a test on $c \in B$ is handled by the multiplication matrix M_c. Its kernel and image are complementary ideals: the former consists of the components on which $c = 0$, and the latter of those on which $c \neq 0$. Decomposing 1_B again computes their identities. On the latter ideal, solve $cd = 1$ by linear algebra. This gives exactly the same branch simulation and the same NU operation charge as before. The structure constants, subspace bases, state restrictions, and all linear systems have polynomial size in $N, U, \log p$. □

Corollary 4.3. *Given an explicitly represented order $\mathcal{R}$ in a number field and a prime p, one can deterministically compute, in polynomial time in the input length, a p-maximal overorder $\mathcal{R}'$ and the finite $\mathbb{F}_p$-algebra $(\mathcal{R}'/p\mathcal{R}')_{\mathrm{red}}$, by structure constants. The latter is a product of the residue fields above p, and the simulation just proved applies without selecting any of those fields.*

Proof. Algorithm 6.1 and Proposition 6.2 of Buchmann and Lenstra (1994, pp. 246–247) compute the p-maximal overorder. Each iteration computes a nilradical by Frobenius linear algebra and then a multiplier order; the proof bounds the number of strict iterations by $\log |\operatorname{disc}(\mathcal{R})|/(2 \log p)$. For any finite algebra of dimension n, its nilradical is the kernel of $x \mapsto x^{p^t}$ as soon as $p^t \ge n$. This follows either from its decomposition into Artinian local rings, or from the descending chain of images of multiplication by an element. Compute this kernel and quotient by it. The resulting finite reduced algebra is a product of finite fields and has the asserted interpretation because the order is maximal at p. Neither operation uses finite-field factorization. □

Corollary 4.4. *Let A be a finite $\mathbb{F}_p$-algebra as in either simulation proposition. Suppose an abelian scheme over A, including its identity section, is explicitly given by projective equations whose affine-chart ideals are radical, together with a complete addition law of fixed shape: the ambient dimension, number of equations, and all their degrees are bounded by constants. Then one can compute a direct-product decomposition of A and the Frobenius characteristic polynomial of every component fiber in time polynomial in N and $\log p$, without selecting a component field. The Frobenius on a component of degree e is the p^e-power Frobenius.*

Proof. Apply Propositions 4.1 and 4.2 to the algorithm of Pila (1990, Theorem A; Sections 3–4, author/PDF pp. 8–14). Its input format, including radical affine-chart ideals, addition atlas, and identity, is specified on pp. 8–9. The monomial-dependence tests on pp. 10–11 are explicitly reduced to inhomogeneous linear systems over the coefficient field, using Theorem 2.7 on p. 8. The Frobenius powers are computed by repeated squaring and reduction on p. 11. Section 4, p. 12, expressly lists the four activities: polynomial arithmetic, linear-system solution, enumeration over the small torsion prime fields, and integer Chinese remaindering. Its chosen irreducible-polynomial presentation is used to implement coefficient-field arithmetic; none of these four activities branches on a component's representation-specific encoding. The enumeration on p. 9 is over $\mathbb{F}_\ell$, with $\ell = O(\log(p^e))$, not over the coefficient field. Thus the entire algorithm has the arithmetic form used in the proposition. The operation bounds on pp. 12–14 are polynomial in $\log(p^e) = e \log p$, with constants depending on the fixed shape. Since $e \le N$, the simulation has the stated cost. □

This interface assumes that the equations and addition law already define smooth abelian varieties in every component. It neither proves a polynomial cost for constructing a varying CM model nor provides the characteristic-zero identification of the resulting Frobenius with a specified CM element. Those requirements are established in Sections 7–13.

5 A uniform construction when the cyclotomic cofactor is small

We adapt the rational-function and primary-coboundary arguments of Bach et al. (2001, Proposition (5.2), pp. 20–21; Sections 5.7 and 5.9–5.11, pp. 22–23). Their Theorem 3, printed p. 7, takes an explicit field of cardinality p^d and primitive prime-order roots for *every* prime dividing $\Phi_d(p)$. Theorem 2 on the same page constructs such an ℓ-th root from an explicit $\mathbb{F}_{p^h}$ and irreducible prime-degree polynomials for primes dividing $\ell - 1$ but not dividing h; see its proof, pp. 16–20. Theorem 1, p. 6, retains those construction inputs in its smooth-cyclotomic factorization reduction. Here only selected prime-order roots are supplied; certain others are replaced by one formal cyclotomic algebra and a full-algebra norm.

The restriction $r \nmid \ell - 1$ has an earlier counterpart in Ivanyos et al. (2012, Theorem 1.5, author p. 6, journal p. 498): for a degree-n polynomial splitting over $\mathbb{F}_p$, their smooth-$p-1$ algorithm always finds a proper factor when $n \nmid \operatorname{lcm}\{\ell - 1 : \ell \mid p - 1\}$. For prime $n = r$, this excludes primes $\ell \equiv 1 \pmod r$. Their Theorem 6.3, author p. 32 (journal p. 524), and its preceding Lemmas 6.1–6.2 concern that special-field setting. The introduction on author p. 6 explicitly distinguishes the general $\Phi_d(p)$-smooth case and the problem of computing roots. The projection and residual-cofactor bound below therefore require their own proof.

Sze permits a partial factorization of $q - 1$, charging the numerical residual and construction of selected roots of unity (Sze, 2011, Definition 2.1 and Theorem 2.2, p. 4; Lemma 3.21, pp. 17–18). His prime-two input is a fourth root of unity; the odd-degree argument below uses only -1. Replacing $q - 1$ by $\Phi_d(p)$, or deleting his root inputs, is not an immediate substitution in that theorem. Formal cyclotomic ring extensions and tame fixed algebras are classical tools of Lenstra (1991, Section 4, pp. 336–338; Theorems (5.1)–(5.2), pp. 338–340). His Theorem (9.1), pp. 344–346, constructs extensions from the specified prime-degree irreducibles,

and Theorem (1.2), p. 330, finds isomorphisms between supplied fields. Neither supplies an arbitrary missing presentation.

Theorem 5.1. *Let $p \neq r$ be primes, with r odd, and put*

$$d = \operatorname{ord}_r(p), \qquad q = p^d, \qquad a = v_r(q-1),$$
$$N = \Phi_d(p), \qquad s = v_2(N), \qquad \mathcal{D} = \mathbb{F}_p[Z]/\Phi_r(Z).$$

Let S, T be disjoint finite sets of odd primes other than r dividing N. For $\ell \in S$, supply $\eta_\ell \in \mathcal{D}$ whose image has exact order ℓ in every field component. For $\ell \in T$, assume $r \nmid \ell - 1$; no algebraic data attached to ℓ is supplied. Define

$$C = \frac{N}{2^s r^a \prod_{\ell \in S \cup T} \ell^{v_\ell(N)}}.$$

An unconditional deterministic algorithm constructs a monic irreducible polynomial of degree r over $\mathbb{F}_p$ and an element $U \in \mathcal{D}$ such that

$$U^{r^{a-1}} = Z, \qquad \sigma_b(U) = U^{\omega_a(b)} \quad (b \in \mathbb{F}_r^\times),$$

with order exactly r^a in every component. Its bit cost is

$$\left(r + C + \sum_{\ell \in S \cup T} \ell + \log p \right)^c$$

for an absolute c, independent of the valuations and the number of listed primes. Numerical prime values and the numerical cofactor occur in this bound. No factorization of C, component of $\mathcal{D}$, or additional field presentation is required as input.

The same bound gives simultaneous root extraction: for any given nonconstant monic $g \mid \Phi_r$, put $A = \mathbb{F}_p[Z]/g$ and restrict the supplied roots to A. For every given $z \in A^\times$ satisfying $z^{(q-1)/r} = 1$, the algorithm returns $y \in A$ with $y^r = z$ on every field component, without a supplied decomposition. It never constructs an algebra whose dimension is a product over the primes in T.

Proof. Every irreducible factor of Φ_r over $\mathbb{F}_p$ has degree d, and all factors are distinct. Thus the components of A have size q, and $\zeta = Z$ is a strict primitive r-th root. The condition on z says exactly that it is an r-th power in each component. Also $d \mid r-1$, and r divides none of $p^e - 1$ for proper divisors $e \mid d$. Hence r^a is the full r-primary part of N.

Arithmetic, splitting, and a proper-factor formula. All algebras are represented by explicit $\mathbb{F}_p$-bases. For an element v of a finite étale algebra E, its multiplication matrix detects units and computes inverses. If v is a nonzero nonunit, then

$$E = vE \oplus \operatorname{ann}(v)$$

as ideals: this holds on every field component. Compute the image and kernel of the multiplication matrix, and decompose 1 in these complementary subspaces. The resulting identities are orthogonal idempotents, giving explicit projection and inclusion maps. Process both nonzero factors and patch their outputs by addition. For $A = \mathbb{F}_p[Z]/g$, this is also the usual CRT split from $\gcd(g, v)$ and its complementary factor.

Use this rule whenever a field calculation encounters a nonzero nonunit leading coefficient or a component-dependent equality. Polynomial Euclid then operates with unit leading coefficients on each branch. A split partitions a positive integer dimension; a tree starting with dimension

m has at most m leaves and $m-1$ splits. Restarting a bounded calculation after each split adds polynomial overhead. No nonzero component is discarded.

If a monic divisor $h \mid X^r - z$ has constant degree j with $0 < j < r$ on one such branch, let $b = (-1)^j h(0)$. The roots are $\alpha\zeta^i$, so $b^r = z^j$ componentwise. For integers u, v with $uj + vr = 1$,

$$(b^u z^v)^r = z. \tag{*}$$

The constant coefficient is a unit, so negative powers are valid. Degrees zero and r are not accepted as proper factors. An equality for an element of a binomial algebra is tested on all its coefficients. The common zero support is split before declaring that equality false on a branch. Thus every nontriviality and exact order used below holds on each coefficient-field component of its branch, rather than merely somewhere in the product algebra.

The cyclotomic projection and bounded search. First describe the field calculation in a component K of size q, with primitive root ζ. Form

$$R = K[X]/(X^r - z), \qquad \alpha = X \bmod (X^r - z), \qquad \tau(\alpha) = \zeta\alpha,$$

where τ fixes K. The promise makes $R \cong K^r$, and τ has order r. Absolute Frobenius $F(x) = x^p$ satisfies $F^d = 1$ and commutes with τ. Put $h = 2^{\omega(d)} \le d$, with $h = 1$ if $d = 1$. For squarefree $e \mid d$, define

$$F_e = \prod_{\lambda \mid e} F^{d/\lambda}, \qquad E_d = \prod_{\lambda \mid d} (1 - p^{d/\lambda}), \qquad P(v) = \prod_{e \mid \mathrm{rad}(d)} F_e(v)^{\mu(e)} = v^{E_d}.$$

Products indexed by λ are over primes; empty products mean the identity, or 1 for E_d. Each proper cyclotomic factor of $T^d - 1$ divides some $T^{d/\lambda} - 1$, whence

$$P(v)^N = 1, \qquad r \nmid E_d. \tag{5}$$

The second assertion follows from the definition of d.

Let $B = rhC$, and suppose $p \ge B$. For $i = 1, \dots, B-1$, first test $\alpha + i$ for being a unit. A failure gives a proper factor of $X^r - z$, and hence a root by (*). Otherwise set

$$b_i = P(\alpha + i), \qquad \delta_i = \tau(b_i)/b_i.$$

There is an i with $\delta_i^{rC} \ne 1$. To prove this, choose a root β of $X^r - z$ solely for the argument, and specialize ζ to its given value. Consider

$$H(Y) = \prod_{e \mid \mathrm{rad}(d)} \left(\frac{F_e(\zeta\beta) + Y}{F_e(\beta) + Y} \right)^{\mu(e)}.$$

Its numerator and denominator are monic of degree at most h, with the same degree after cancellation. If all searches failed, $H(i)^{rC} = 1$ at all $B-1$ tested integers. These values are defined by the unit tests. At zero the equality also holds, but

$$H(0) = \zeta^{E_d} \ne 1, \qquad H(\infty) = 1.$$

The numerator of $H^{rC} - 1$ has degree at most $B-1$, because its leading terms cancel. Its B distinct zeros $0, 1, \dots, B-1$ force it to vanish identically. A rational function with a positive power equal to 1 is constant, contradicting the two displayed values. Thus the search succeeds. This uses no factorization of C and no numerical prime-power enumeration.

For a successful shift set

$$v = \delta_i^C, \qquad Q = N/C.$$

Then $v^Q = 1$ and $v^r \neq 1$. It remains a coboundary: $v = \tau(c)/c$ for the known unit $c = b_i^C$. Every power of a coboundary satisfies

$$\prod_{j=0}^{r-1} \tau^j(v) = 1. \qquad (**)$$

In particular, a scalar coboundary has r-th power 1.

The non-r primary branches. The primes dividing Q are known: r, possibly 2, and the members of $S \cup T$. For each $\ell \neq r$ among them test $v^{Q/\ell^{v_\ell(Q)}}$. If a test is nontrivial, successive ℓ-th powers produce an element ϵ of order exactly ℓ, still a coboundary. In a field coefficient algebra it is not scalar, by (**) and $\ell \neq r$. For a known primitive ℓ-th root $\eta \in K$, the gcds with $\epsilon - \eta^j$, $0 \leq j < \ell$, therefore yield a proper positive factor of $X^r - z$. Only order ℓ, not maximal ℓ-primary order, is used. For $\ell = 2$, the root is -1; this needs no fourth root because r is odd. In characteristic 2 the factor 2^s is absent.

Over the original coefficient algebra A, a nontrivial ϵ might equal 1 on the whole rank-r algebra above some components. Compute that support from the simultaneous zero set of all coefficients of $\epsilon - 1$, using the splitting rule. Restart on that support, if nonzero, and continue on its complement. On every coefficient-field component of the complement, ϵ is nonconstant: a scalar value would be both an ℓ-th and an r-th root, hence 1. The complement is nonzero, so this is a genuine decrease of base dimension. All remaining equality tests are refined in the same way.

The r-primary branch, including the order-r-squared case. If every non-r projection is trivial, v has r-power order r^t, with $t \geq 2$ since $v^r \neq 1$. Find t by repeated powering, and write

$$\nu = v^{r^{t-2}} = \tau(c_0)/c_0, \qquad c_0 = c^{r^{t-2}}.$$

Thus ν has order r^2. If ν^r is not a scalar, comparison with $1, \zeta, \ldots, \zeta^{r-1}$ gives a proper factor, since it takes different r-th-root values on the split algebra. Otherwise ν^r is a nontrivial scalar r-th root. Put $\epsilon = \tau(\nu)/\nu$; then $\epsilon^r = 1$. A nonscalar ϵ again gives a proper factor by the same comparisons. If it too were scalar, induction would give

$$\tau^j(c_0) = \epsilon^{j(j-1)/2} \nu^j c_0.$$

At $j = r$, oddness of r and $\epsilon^r = 1$ imply $\nu^r = 1$, a contradiction. This proves termination in the order-r^2 case. It works for an arbitrary promised z; no scalar primitive r^2-th root, recursive root oracle, or special presentation of z is assumed. Over an étale coefficient algebra, scalar equalities and their failures are refined as above. If a comparison reveals only a proper base support, process both sides and restart where needed. This adds at most the bounded base-splitting overhead.

Cyclic comparison inside the original algebra. The following comparison specializes the discrete-logarithm or zero-divisor lemma of Ivanyos et al. (2012, Lemma 2.3) to a prime-order element and its shift. The order of the shift determines when the discrete-logarithm outcome can occur. Let ϵ be the nonscalar element of order $\ell \neq r$ found above. Work first over one coefficient field K, so the algebra $R = K[X]/(X^r - z)$ has r geometric points, cyclically permuted by τ. A nonempty support of $\epsilon = 1$ is proper and gives the factor $\gcd(X^r - z, \epsilon - 1)$. Otherwise all its values are primitive ℓ-th roots. For $1 \leq j < \ell$ compute

$$\gcd(X^r - z,\ \tau(\epsilon) - \epsilon^j).$$

At each point exactly one exponent j gives equality. Unless some gcd has degree strictly between 0 and r, one exponent works at all the points. Then $\tau(\epsilon) = \epsilon^j$ and applying $\tau^r = 1$ gives $j^r \equiv 1 \pmod{\ell}$. The exponent $j = 1$ would make ϵ constant on the orbit, contrary to its construction. Thus j has order r in $(\mathbb{Z}/\ell)^\times$. For $r \nmid \ell - 1$ this is impossible. Formula (*) supplies the root from the proper factor. The same comparisons over a product of coefficient fields use the unit tests and the support splits above, with at most ℓ gcds on each branch. This proves the assertion for every prime in T within the original binomial algebra. The congruence condition in the split-polynomial algorithm of Ivanyos et al. (2012, Theorem 1.5) was recalled at the start of this section.

A second proof by descent of a norm. If the non-r branch chooses $\ell \in T$, work on the base branch A_0 where ϵ is nonconstant above every component, and form

$$E = A_0[Y]/\Phi_\ell(Y), \qquad \theta = Y \bmod \Phi_\ell(Y).$$

This is free of rank $n = \ell - 1$ over A_0. As $\ell \mid q - 1$, it is étale, all its field components still have size q, and θ is a strict primitive ℓ-th root.

After base change, compute the gcds with $\epsilon - \theta^j$, $0 \le j < \ell$, using unit leading coefficients. A nonunit coefficient splits E; recursively process both factors with the same comparisons. Nonconstancy over each coefficient field is preserved. A gcd of degree r would mean $\epsilon = \theta^j$ on that entire branch and is impossible. Degree zero means try the next j; degree strictly between zero and r gives a root by (*). Every component has such a proper gcd, because at least two ℓ-th-root values occur on its r geometric points. Different gcd degrees on different components force an explicit split before a monic constant-degree result is returned. There are at most $\dim_{\mathbb{F}_p} E$ leaves. Patching all their answers yields

$$y \in E^\times, \qquad y^r = z\, 1_E.$$

This subroutine performs no new rational search inside E and adjoins no other prime-order root. It is also the simultaneous comparison subroutine for the supplied roots in S.

The norm uses the full original rank. Return y to the original free presentation of E, and compute $b = \mathrm{Norm}_{E/A_0}(y)$ as the determinant of multiplication in the basis $1, Y, \ldots, Y^{n-1}$. A division-free determinant, or Gaussian elimination with the explicit base splitting and CRT rule, has polynomial cost. Then

$$b^r = z^n.$$

Since $r \nmid n$, take integers u, v with $un + vr = 1$. The element $b^u z^v \in A_0$ is the required simultaneous root. A root on only one component of E would not suffice: the norm is taken only after all components are patched, and uses the full rank $\ell - 1$, not the possibly unsuitable rank of a selected quotient. This branch is now finished; discard E. This argument also gives the following version with partial data. For a selected ℓ suppose a monic divisor $h_\ell \mid \Phi_\ell$ in $\mathcal{D}[Y]$ is given, of positive degree n prime to r. Restrict its coefficients to the current branch and use $E = A_0[Y]/h_\ell$ in place of the full cyclotomic algebra. It is free of rank n. Each fiber splits into distinct roots of order ℓ, so the same comparisons find and patch roots on every component. The determinant in this rank again has $b^r = z^n$. This recovers the stated bound since $n \le \ell - 1$. A supplied root gives a linear h_ℓ; the full Φ_ℓ works without further data if $r \nmid \ell - 1$. For example, take $p = 251$, $r = 3$ and $\ell = 7$. Here $\Phi_2(251) = 252 = 2^2 \cdot 3^2 \cdot 7$. Although $3 \mid 6$, the supplied quadratic $Y^2 - 106Y + 1$ suffices: $106^3 + 106^2 - 2 \cdot 106 - 1 \equiv 0 \pmod{251}$ and $\Phi_7(Y) = Y^3((Y + Y^{-1})^3 + (Y + Y^{-1})^2 - 2(Y + Y^{-1}) - 1)$. Its norm uses rank two. In general the existence of a divisor over the coefficient field leaves the computational problem of finding it.

Other base branches may choose other primes, but no tensor product over all primes in T is formed. When $r \mid \ell - 1$, this Bezout descent is unavailable, which explains the stated restriction rather than removing it.

Small characteristic and the case a equals one. If $p < rhC \le r^2 C$, an elementary absolute Frobenius algorithm suffices. In any finite étale $\mathbb{F}_p$-algebra J, compute a basis of $\ker(F - 1)$. This is $\mathbb{F}_p^k$ when J has k field factors. For each fixed-space basis element and every scalar in $\mathbb{F}_p$, split by their difference. The joint coordinates of this basis separate all factors, so the process terminates in fields, in time polynomial in $p, \dim J, \log p$. Apply it first to A, then, for each resulting K, to $K[X]/(X^r - z)$ as an absolute algebra. Each final factor has dimension d, and the inclusion of K is a vector-space isomorphism. Pull the image of X back by linear algebra, and patch the roots across all components of A. The enumeration is over p scalars, not p^d elements. This covers $p = 2$ without an analytic assumption.

If $a = 1$, root extraction can instead be done directly. Put $m = (q - 1)/r$; the promise gives $z^m = 1$ and $\gcd(m, r) = 1$. For $m > 1$, choose e with $re \equiv 1 \pmod{m}$ and return z^e. For $m = 1$, the promised input is 1 and its root is 1.

Lifting and tame normalization. In $\mathcal{D}$, start with $W_1 = Z$ and repeatedly use the simultaneous root interface to obtain $W_{j+1}^r = W_j$, for $j < a$. An element of order r^j, $j < a$, is an r-th power in every component of cardinality q. Any returned root has order exactly r^{j+1}. Consequently $W = W_a$ satisfies $W^{r^{a-1}} = Z$, with strict order r^a; for $a = 1$, simply take $W = Z$. Choose $e(r - 1) \equiv 1 \pmod{r^a}$, and set

$$U = \left(\prod_{b \in \mathbb{F}_r^\times} \sigma_b(W)^{\omega_a(b)^{-1}} \right)^e .$$

Reindexing proves $\sigma_c(U) = U^{\omega_a(c)}$. Raising the inner product to r^{a-1} gives Z^{r-1}, so $U^{r^{a-1}} = Z$, proving strict order. The lifts can be computed by $\omega_a(b) = b^{r^{a-1}} \pmod{r^a}$. The inverse of $r - 1$ is taken modulo r^a, not in $\mathbb{F}_p$; no restriction $p \nmid r - 1$ is needed.

Descent to a field of degree r. For an explicit form of Lenstra's tame descent, form

$$B_0 = \mathcal{D}[X]/(X^r - U), \qquad \sigma_b(X) = X^{\omega_{a+1}(b)}.$$

The relation is respected by the equivariance of U, and the multiplicativity of the tame lifts gives the group law. This is an étale algebra of dimension $r(r - 1)$, and X has strict order r^{a+1}. The action of $\mathbb{F}_r^\times$ on its geometric points is free, and is transitive on the $r - 1$ base images of Z. Each orbit meets any chosen fiber over Z once, so there are r orbits. Frobenius to the power d fixes that base image and cycles through all r points in its fiber, since $v_r(p^d - 1) = a$. It therefore acts transitively on the orbits. Thus the fixed algebra is one field of degree r over $\mathbb{F}_p$. Compute it by intersecting the kernels of $\sigma_b - 1$; this uses no division by $r - 1$. Any vector outside the scalar line generates this prime-degree field. Its first linear dependence among $1, x, \ldots, x^r$ gives the required monic irreducible polynomial.

Absolute bit bound. Put $P = 2 + r + C + \sum_{\ell \in S \cup T} \ell + \lceil \log_2 p \rceil$. We have $d, h \le r$, $\log q \le r \log p$, and all valuations are at most $\log_2 q$. The binary lengths of N, Q, E_d and every powering exponent are polynomial in P; for example $\log(1 + |E_d|)$ is bounded by a constant times $1 + d\omega(d) \log(p + 1)$. Trial division of $d < r$, cyclotomic evaluation, and repeated division by the listed primes have absolute polynomial cost; C is not factored. There are fewer than $r^2 C$ shifts. Primary chains have at most $\log_2 N$ steps and root comparisons use at most ℓ values. Each auxiliary binomial algebra has dimension at most $r(r - 1)(\ell - 1) \le P^3$, and every split

tree is bounded by its starting dimension. Arithmetic, Frobenius, gcds, kernels, inverses and determinants therefore have absolute polynomial bit cost, even with dense tables and elementary arithmetic. There are at most $a-1 \le \log_2 q$ lifting calls. In the small-characteristic branch, $p < r^2 C \le P^3$, so its scalar enumeration has the same type of bound. No loop is indexed by a numerical primary power and no dimension is a product over the listed primes. This proves all the claims with one absolute exponent. □

Proof of Theorem 1.1. Discover the primes in $\mathcal{T}_L$ by a sieve or trial division up to L, and compute their full valuations by division into N. Apply Theorem 5.1 with $S = \varnothing$ and $T = \mathcal{T}_L$. Since $\sum_{\ell \le L} \ell \le L^2$, enlarging the absolute exponent gives the stated bound. No algebraic input was introduced by this preliminary integer computation. □

5.1 Numerical cofactor special cases

Theorem 5.2. *Let $p \neq r$ be primes, with r odd, and put*

$$d = \mathrm{ord}_r(p), \qquad a = v_r(p^d - 1), \qquad M = \frac{\Phi_d(p)}{r^a}.$$

There is an unconditional deterministic algorithm constructing a monic irreducible polynomial of degree r over $\mathbb{F}_p$ in bit time polynomial in $r, M, \log p$, with one absolute exponent. It also constructs, in the formal algebra $\mathcal{D} = \mathbb{F}_p[Z]/\Phi_r(Z)$, an element U satisfying

$$U^{r^{a-1}} = Z, \qquad \sigma_c(U) = U^{\omega_a(c)} \quad (c \in (\mathbb{Z}/r\mathbb{Z})^\times),$$

where $\sigma_c(Z) = Z^c$ and $\omega_a(c)$ is the tame lift modulo r^a. Every component of U has exact order r^a. No factorization of M, presentation of $\mathbb{F}_{p^d}$, or roots of unity for prime divisors of M are inputs.

Proof. Apply Theorem 5.1 with $S = T = \varnothing$. Its residual is $M_2 = N/(2^{v_2(N)} r^a) \le M$, so it gives the stated bound and all the displayed output properties. □

Theorem 5.3. *Let $p \neq r$ be primes, with r odd, and put*

$$d = \mathrm{ord}_r(p), \qquad a = v_r(p^d - 1), \qquad M_2 = \frac{\Phi_d(p)}{2^{v_2(\Phi_d(p))} r^a}.$$

There is an unconditional deterministic algorithm constructing a monic irreducible polynomial of degree r over $\mathbb{F}_p$ in bit time polynomial in $r, M_2, \log p$, with one absolute exponent. It also constructs, in the formal algebra $\mathcal{D} = \mathbb{F}_p[Z]/\Phi_r(Z)$, an element U satisfying

$$U^{r^{a-1}} = Z, \qquad \sigma_c(U) = U^{\omega_a(c)} \quad (c \in (\mathbb{Z}/r\mathbb{Z})^\times),$$

where $\sigma_c(Z) = Z^c$ and $\omega_a(c)$ is the tame lift modulo r^a: it is congruent to c modulo r and has order dividing $r-1$. Every component of U has exact order r^a. No factorization of M_2, presentation of $\mathbb{F}_{p^d}$, or roots of unity for prime divisors of M_2 are inputs.

Proof. This is precisely Theorem 5.1 with $S = T = \varnothing$. The proof above includes the two-primary comparison with -1, the small-characteristic branch, simultaneous lifting, and tame descent. □

Lemma 5.4. *Use p, r, d, a, M as above, let K be a presented field of cardinality p^d, and suppose a primitive r-th root $\eta \in K$ is supplied. Put $t = 2^{\omega(d)}$, where $\omega(d)$ counts distinct prime divisors and $t = 1$ if $d = 1$. If $p \geq trM$, an r-th root in K of any promised nonzero r-th power $c \in K$ can be computed using a number of field operations and equality tests polynomial in $r, d, M, \log p$. The same bound holds for the non-field bit operations. The algorithm uses no other operations on the representation of field elements.*

Proof. Use the field calculation in the proof of Theorem 5.1, with $S = T = \varnothing$ and residual $M_2 \leq M$. The hypothesis $p \geq trM$ puts it in the large-characteristic branch. That calculation uses only arithmetic and equality in K, polynomial arithmetic of degree at most r, and integer computations with polynomial-length exponents. Its bound with M_2 implies the asserted bound with M. No decomposition or operation on the representation of elements of K is needed in this field case. □

Corollary 5.5. *Let $p \equiv -1 \pmod 3$ be prime. Suppose every odd prime $\ell \mid p+1$ satisfies that $\ell - 1$ is a power of two, and let S be the largest prime divisor of $p+1$. There is an unconditional deterministic algorithm producing an element of maximal three-power order in $\mathbb{F}_{p^2}$ and a monic irreducible cubic over $\mathbb{F}_p$, with bit cost polynomial in $S, \log p$, with an absolute exponent. In particular, this cost is polynomial in $\log p$, uniformly in both exponents, when $p + 1 = 2^u 3^a$.*

Proof. Trial division until the remaining factor is 1 discovers the full support of $p + 1$ in time polynomial in $S, \log p$. Every prime in that support other than 2 and 3 satisfies $3 \nmid \ell - 1$. Apply Theorem 5.1 with those primes in T, with no supplied roots, and with $C = 1$. Here $\mathcal{D} = \mathbb{F}_p[Z]/(Z^2 + Z + 1)$ is the required quadratic field. For the explicit classical cubic, put $u = U + U^{-1} \in \mathbb{F}_p$, using $U^p = U^{-1}$. Then $Y^3 - 3Y - u$ is irreducible. A root $b \in \mathbb{F}_p$ would give $w \in \mathbb{F}_{p^2}^\times$ with $w + w^{-1} = b$, and hence $w^3 \in \{U, U^{-1}\}$, contradicting maximal three-primary order. For $p = 2$, this is the explicit polynomial $Y^3 + Y + 1$.

This corollary is also a direct classical transfer, not a new smooth-support assertion: BGL Theorem 2, printed p. 7 (proof, pp. 16–20), with $h = 2$, constructs the primitive ℓ-th roots without any extra prime-degree inputs, because the only prime dividing $\ell - 1$ is 2 and it already divides h. Their Theorem 3, p. 7 (proof, pp. 20–24), then supplies a noncube in that presented quadratic field. Raising it to $(p^2 - 1)/3^{v_3(p+1)}$ and using the same norm-one cubic proves the stated bound (Bach et al., 2001). □

An explicit cubic example. Take $p = 197$, $r = 3$, and $L = 11$. Then $d = 2$, $N = 198 = 2 \cdot 3^2 \cdot 11$, and the numerical cofactor after removing only two and three is 11. Since $3 \nmid 11 - 1$, the threshold theorem has $C_{11} = 1$. The prime eleven does not satisfy the power-of-two support condition of the preceding classical corollary. A cubic can be checked directly in this example: $u = -6$ satisfies $u^3 - 3u + 1 = -197 = 0$ in $\mathbb{F}_{197}$. A root U of $X^2 - uX + 1$ therefore has order nine. Indeed $U^3 + U^{-3} = -1$, so U^3 is a nonidentity third root. The discriminant is 32, a nonsquare because $197 \equiv 5 \pmod 8$; hence U lies in the quadratic field and has norm one. Nine is its full three-primary order. The same trace argument proves that $Y^3 - 3Y + 6$ is irreducible over $\mathbb{F}_{197}$. This finite example illustrates the changed cofactor; it is not an infinitude assertion about primes with a prescribed factorization.

5.2 Fixed selected primes for cubic construction

Corollary 5.6. *Fix a finite set S of odd primes different from 3, and assume A, the Riemann hypothesis for all rational Dirichlet L-functions. For a prime $p \equiv 2 \pmod 3$, put*

$$C_S(p) = \frac{p+1}{2^{v_2(p+1)}3^{v_3(p+1)}\prod_{\ell\in S}\ell^{v_\ell(p+1)}}.$$

There is a deterministic algorithm constructing a monic irreducible cubic over $\mathbb{F}_p$ in at most $K_S(C_S(p)+\log p)^{c_S}$ bit operations. In $\mathcal{D} = \mathbb{F}_p[Z]/(Z^2+Z+1)$, it also returns U of exact order 3^a, $a = v_3(p+1)$, with

$$U^{3^{a-1}} = Z, \qquad \sigma_2(U) = U^{-1}.$$

The constants depend only on the fixed set S. No roots of unity, field presentation, residual factorization, or number-field residue map are inputs. The assertion includes $p = 2$. For primes $p \equiv 1 \pmod 3$, a separate deterministic polynomial-time cubic construction holds under A, and $p = 3$ is explicit. No bound uniform in a growing S is asserted.

Proof. We use Proposition 16.2 below only for fixed totally real Galois fields, with their fixed integral basis and multiplication table. That proposition is a proved dependency, not a supplied residue-map oracle. Handle $p = 2$ directly by $U = Z$ and $X^3 + X + 1$. For odd $p \equiv 2 \pmod 3$, the displayed quadratic algebra is a field, and $\beta = 2Z + 1$ satisfies $\beta^2 = -3$.

Only primes $\ell \in S$ dividing $p + 1$ need roots. For such $\ell \geq 7$, fix $K_\ell = \mathbb{Q}(\zeta_\ell + \zeta_\ell^{-1})$. It is totally real Galois of fixed degree $(\ell - 1)/2 \geq 3$. Since $p \equiv -1 \pmod \ell$, p splits completely in K_ℓ. The proposition constructs $\lambda : \mathcal{O}_{K_\ell} \to \mathbb{F}_p$ under A. Compute $t = \lambda(\zeta_\ell + \zeta_\ell^{-1})$ from its fixed integral coordinates. If $\ell = 5$, do not apply that degree-at-least-three proposition: instead compute $b^2 = 5$ in $\mathbb{F}_p$ and put $t = (-1 + b)/2$. Existence follows from $p \equiv -1 \pmod 5$.

If h_ℓ is the minimal polynomial of $\zeta_\ell + \zeta_\ell^{-1}$, then

$$\Phi_\ell(X) = X^{(\ell-1)/2}h_\ell(X + X^{-1}).$$

Thus the roots of $X^2 - tX + 1$ are primitive ℓ-th roots in $\mathbb{F}_{p^2}$. They cannot lie in $\mathbb{F}_p$, since the odd prime ℓ divides $p + 1$ but not $p - 1$. Consequently $\Delta = t^2 - 4$ is a nonzero nonsquare. So is -3, and $\Delta/(-3)$ is a square. Compute $w^2 = \Delta/(-3)$ in $\mathbb{F}_p$ and set

$$\eta_\ell = (t + w\beta)/2 \in \mathcal{D}.$$

This gives the root in the required displayed algebra, not merely in an unspecified quadratic extension.

For these prime-field square roots, rational Dirichlet GRH supplies a quadratic nonresidue in a polynomial-length initial interval (Lamzouri et al., 2015, Theorem 1.1(2), corrected author p. 2); primes below 3000 are handled by finite search. Here is the square-root computation explicitly. For a nonzero square b, write $p - 1 = 2^s m$, with m odd, and choose the quadratic nonresidue n. Start with

$$c = n^m, \quad x = b^{(m+1)/2}, \quad t = b^m, \quad M = s.$$

The invariants are $x^2 = bt$, $\operatorname{ord}(c) = 2^M$, and $t^{2^{M-1}} = 1$. If $t \neq 1$, find the least i with $t^{2^i} = 1$; then $1 \leq i < M$. Put $v = c^{2^{M-i-1}}$ and replace

$$(x, t, c, M) \longmapsto (xv, tv^2, v^2, i).$$

Both t and v^2 have the same nontrivial 2^{i-1}-st power, namely -1, so the new t has order dividing 2^{i-1}. The invariants persist and M strictly decreases. Return x when $t = 1$; return zero for a zero input. At most s stages and polynomial-length exponentiations are needed. The

nonresidue is reused for every ℓ. No extension-field nonresidue or stronger zero hypothesis is introduced.

Apply Theorem 5.1 with $r = 3$, $T = \varnothing$, and the constructed roots for the subset of S dividing $p + 1$. Its residual is exactly $C_S(p)$. The tame lift of 2 is -1, giving the stated equivariance. All fields K_ℓ are fixed, so their precomputation and polynomial costs are absorbed into K_S, c_S. Integral coordinates introduce no additional denominator exceptions. The fixed-field proposition covers all completely split primes; any explicit finite exceptional set in its implementation can also be handled by enumerating the p^2 elements of the displayed quadratic field, testing the required orders, and enumerating monic cubics until one has no prime-field root. These finite terminating computations are absorbed in the fixed constant.

For $p \equiv 1 \pmod 3$, apply the same rational Dirichlet subgroup bound to obtain a cubic nonresidue $b \in \mathbb{F}_p$; then $X^3 - b$ has no root and is irreducible. For $p = 3$, use $X^3 - X - 1$. The cofactor formula was used only for $p \equiv 2 \pmod 3$, as stated. □

For the concrete fixed choice $S = \{7\}$, the parameter is

$$C_7(p) = \frac{p+1}{2^{v_2(p+1)}3^{v_3(p+1)}7^{v_7(p+1)}}.$$

Only the fixed real cubic field $\mathbb{Q}(\zeta_7 + \zeta_7^{-1})$ is used for the residue map, so all constants are absolute. No density or infinitude assertion about primes with small $C_7(p)$ follows or is asserted.

6 A small quadratic splitting prime

The CM construction uses only a quadratic splitting condition for its varying auxiliary field; Dirichlet GRH supplies it within a polynomial search.

Lemma 6.1. *Assume rational Dirichlet GRH. Let p, r be distinct odd primes, let $d = \operatorname{ord}_r(p)$ be even, and put $t = v_2(d)$. There is a prime $q \notin \{2, p, r\}$ such that*

$$q \equiv 1 \pmod{2^{t+1}}, \qquad \left(\frac{p}{q}\right) = -1, \qquad q \ll d^2 \log^2(2pd),$$

with an absolute effective implied constant. Such a prime can be found deterministically in time polynomial in r and $\log p$. If $a = v_r(p^d - 1)$, then for every $1 \le b \le a$ there is no integer j satisfying $p^j \equiv -1 \pmod{r^b q}$.

Proof. Put $M_2 = 2^{t+1}$, and let κ be the primitive quadratic Dirichlet character of $\mathbb{Q}(\sqrt{p})$. On odd primes different from p, it takes the value (p/q). Expand the indicator of 1 mod M_2 by the characters modulo M_2, and multiply it by $(\kappa_0 - \kappa)/2$, where κ_0 is the principal character of the same modulus as κ. All terms have modulus dividing $4pM_2$. Exactly one is principal, with coefficient $1/(2\varphi(M_2))$; the sum of absolute coefficients is at most 1. The nontrivial character at the odd prime p ensures that no term involving κ is principal.

The complex smoothed Dirichlet estimate of Lemma 25.5 gives

$$\sum_{m \le x} \Lambda(m) \log(x/m)\, 1_{m \equiv 1 \bmod M_2} \frac{\kappa_0(m) - \kappa(m)}{2} = \frac{x}{2\varphi(M_2)} + O(\sqrt{x}\log(2pM_2)).$$

Even prime powers contribute zero. Higher odd prime powers contribute at most $O(x^{1/3}(\log x)^2) = O(\sqrt{x})$. Taking x a sufficiently large absolute multiple of $\varphi(M_2)^2 \log^2(2pM_2)$ therefore gives a prime of the required kind. Since $M_2 \le 2d$, this is the displayed bound. Testing all integers up to the bound, using primality testing and the Jacobi symbol, has the stated polynomial cost.

The quadratic condition gives

$$v_2(\operatorname{ord}_q p) = v_2(q-1) \geq t+1.$$

For $b \leq a$, the order of $p \bmod r^b$ is still d. An exponent giving -1 modulo r^b must have 2-adic valuation $t-1$, whereas one giving -1 modulo q has valuation at least t. They cannot coincide. The same order comparison excludes $q = r$, the quadratic condition excludes $q = p$, and the progression excludes $q = 2$. □

The assumption that p is odd is necessary for this construction: when $p = 2$, a sufficiently strong 2-power progression forces $(2/q) = 1$. That characteristic already belongs to the separate small-characteristic algorithmic case.

7 Explicit primitive CM periods for the sextic construction

The particular sextic fields arising above admit explicit principally polarized periods. Their computation does not need a general period-reduction algorithm.

Proposition 7.1. *Let $q \neq 3$ be an odd prime with $q \equiv 1 \pmod 4$, put $D = 3q$, and let*

$$f(X) = X^3 - 3X + 1, \quad R = \mathbb{Q}(\theta), \quad f(\theta) = 0, \quad s = \sqrt{-D}, \quad L = R(s).$$

Let $\theta_1, \theta_2, \theta_3$ be the real roots of f, $S = (\theta_i^{j-1})_{1 \leq i,j \leq 3}$, and $\epsilon_i = \operatorname{sign} f'(\theta_i)$. The matrix

$$B = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 3 \end{pmatrix}, \qquad \tau = \frac{1}{2}B^{-1} + \frac{i\sqrt{D}}{2} S^{-1} \operatorname{diag}(|f'(\theta_i)|) S^{-T}$$

belongs to $\mathcal{H}_3$. The principally polarized torus $\mathbb{C}^3/(\mathbb{Z}^3 + \tau\mathbb{Z}^3)$ has primitive CM by the order

$$\mathcal{O} = \mathbb{Z}[\theta, (1+s)/2] \subset L.$$

There are absolute positive constants c, C such that

$$c\sqrt{D}\, I_3 \preceq \operatorname{Im}\tau \preceq C\sqrt{D}\, I_3.$$

The periods and all genus-three theta constants at τ with characteristic denominator 2 or 4 can be approximated with certified absolute error 2^{-N} in deterministic time polynomial in D and N. Moreover,

$$|\operatorname{disc}(L)| = 3^9 q^3, \qquad [\mathcal{O}_L : \mathcal{O}] = 3.$$

Proof. The polynomial f has discriminant 81 and three distinct real roots; it defines the real cubic subfield of $\mathbb{Q}(\zeta_9)$. Its power-basis trace dual is $f'(\theta)^{-1}\mathbb{Z}[\theta]$. The quadratic ring of integers is $\mathbb{Z}[\omega]$, where $\omega = (1+s)/2$, with trace dual $s^{-1}\mathbb{Z}[\omega]$. The two fields are linearly disjoint. The product basis is an integral basis of the order $\mathcal{O}$, and its trace dual is

$$\mathcal{O}^\vee = \xi\mathcal{O}, \qquad \xi = (f'(\theta)s)^{-1}.$$

This assertion concerns the product order and does not require that it be the maximal order of L. In fact, the six characters of the cyclic extension $L/\mathbb{Q}$ are $1, \nu, \nu^2, \chi, \chi\nu, \chi\nu^2$, where ν is a cubic character of conductor 9 and χ is the quadratic character of conductor $3q$. Their conductors are respectively $1, 9, 9, 3q, 9q, 9q$. At 3, multiplication by the tame quadratic character cannot cancel the nontrivial wild cubic character; at q the quadratic character remains ramified. The conductor-discriminant identity therefore gives $|\operatorname{disc}(L)| = 3^9 q^3$. The product basis has discriminant $81^2(3q)^3 = 3^{11}q^3$, proving the index assertion.

Choose the CM embeddings ϕ_i by $\phi_i(\theta) = \theta_i$ and $\phi_i(s) = \epsilon_i i\sqrt{D}$. On $\mathcal{O}$, define

$$E(x, y) = \mathrm{Tr}_{L/\mathbb{Q}}(\xi \overline{x} y).$$

The trace-dual identity shows that E is an integral unimodular alternating form. With this sign convention, $\phi_i(\xi) = -i/(|f'(\theta_i)|\sqrt{D})$, so $E(z, iz) > 0$ for every nonzero complex tangent vector. It is therefore a principal Riemann form.

For $0 \le j, k \le 2$,

$$E(\theta^j, \omega\theta^k) = \mathrm{Tr}_{R/\mathbb{Q}}\left(\frac{\theta^{j+k}}{f'(\theta)}\right) = B_{j+1,k+1}.$$

Here the traces for exponents $0, 1, 2, 3, 4$ are $0, 0, 1, 0, 3$, by the power-basis dual identity and the relation $\theta^3 = 3\theta - 1$. Both remaining diagonal blocks of E are zero. Since $\det B = -1$, the basis with first block $(1, \theta, \theta^2)$ and second block $\omega(1, \theta, \theta^2)B^{-1}$ is symplectic. Its complex period blocks are

$$\Omega_1 = S, \qquad \Omega_2 = \left(\frac{1}{2}I_3 + \frac{i\sqrt{D}}{2}\,\mathrm{diag}(\epsilon_i)\right) SB^{-1}.$$

The identity $B = S^T \,\mathrm{diag}(1/f'(\theta_i))S$ now gives exactly $\tau = \Omega_1^{-1}\Omega_2$ as displayed. Its imaginary part is positive definite. Its eigenvalues are $\sqrt{D}$ times those of one fixed positive definite matrix, which proves the two uniform inequalities.

The field $L/\mathbb{Q}$ is cyclic of degree 6. Its only proper CM subfield is $\mathbb{Q}(s)$. An induced CM type would choose the same sign of s at all three real embeddings. The signs of f' at the three ordered real roots are $+, -, +$, so the chosen type is primitive. The Riemann form construction gives a principally polarized abelian variety; primitivity implies simplicity, by the standard CM criterion recalled in Balakrishnan et al. (2016, author/PDF p. 3).

For completeness, certified theta evaluation here can use direct summation. For $r \in \{2, 4\}$ and a characteristic $a/r, b/r$, where $a, b \in \{0, \dots, r-1\}^3$, the absolute value of its term indexed by $n \in \mathbb{Z}^3$ is at most

$$\exp(-\pi c\sqrt{D}\,\|n + a/r\|^2).$$

The sum outside $\|n\|_\infty \le K$, for $K \ge 2$, is bounded by $C_1 \exp(-c_1\sqrt{D}\,K^2)$, with absolute effective positive constants. This follows by separating the three coordinates, bounding the one-dimensional Gaussian tails by an integral, and taking a union over the three coordinates. It suffices to take $K = O(1 + \sqrt{(N+1)/\sqrt{D}})$, and the number of terms is polynomial in N. All entries of the period matrix are fixed-degree algebraic numbers given explicitly above. Certified algebraic arithmetic and interval evaluation of exponentials therefore calculate the finite sum and its tail bound with a polynomial cost in D and N; guard precision polynomial in these quantities suffices. No assumption about a nonzero theta constant is used. □

The next proposition gives effective degree and height bounds for these theta coordinates and recovers them exactly. Subsequent sections construct the finite-field model and its addition atlas.

8 Effective algebraization of the explicit sextic theta point

The exact period description above gives more than a numerical candidate: under rational Dirichlet GRH its theta-null point can be recovered as an exact algebraic point with a polynomial cost in the numerical auxiliary prime. The statement here concerns that point, rather than a smooth integral model and an addition atlas.

Proposition 8.1. *Assume rational Dirichlet GRH, and use the notation of Proposition 7.1. Fix $r \in \{2, 4\}$. Let $T_r \in \mathbb{P}^{r^6-1}(\overline{\mathbb{Q}})$ be the theta-null point with coordinates $\vartheta[a/r, b/r](\tau, 0)$, for $a, b \in \{0, \dots, r-1\}^3$, and write $T = T_r$. There are effectively computable absolute constants C_0, C_1 such that*

$$[\mathbb{Q}(T) : \mathbb{Q}] \le B := C_0 q^{3/2}(1 + \log q)^5, \qquad h(T) \le C_1 \log(2q).$$

One can choose a nonzero coordinate, recover all ratios to that coordinate exactly as algebraic numbers, and construct a common number-field presentation of those ratios in deterministic time polynomial in the numerical value of q. The polynomial's exponent is absolute; it does not depend on q.

Proof. Write A for the principally polarized complex abelian variety defined by the periods. It is algebraic over $\overline{\mathbb{Q}}$, by CM algebraicity; see Milne (2020, Proposition 7.9, p. 54). The algebraic theta embedding and its identification with these analytic coordinates are given in Pazuki (2012, Sections 2.3.1–2.3.4, pp. 27–31).

First we bound the height. Saturating its lattice from $\mathcal{O}$ to $\mathcal{O}_L$ gives a degree-3 isogeny $A \to A'$, where A' has full CM by $\mathcal{O}_L$. Multiplication by 3 gives an isogeny $A' \to A$ of degree 3^5, because $3\mathcal{O}_L \subset \mathcal{O}$ and $[\mathcal{O}_L : \mathcal{O}] = 3$. The isogeny inequality therefore gives

$$h_F(A) \le h_F(A') + \tfrac{5}{2} \log 3;$$

see Pazuki (2012, Remark 2.3, p. 27). The abelian Colmez formula applies to A'. In the original source, Theorem 0.3(ii), p. 631, combines with Theorem III.2.9(i), p. 680; the conductor here is odd, so the relevant function is invariant under inertia at 2 and belongs to the subspace defined on p. 676 (Colmez, 1993). Thus no undetermined multiple of $\log 2$ occurs.

For an explicit upper bound one may use Wei (2023, Corollary 4.3, author/PDF pp. 14–15; Theorem 4.8, p. 18). The reflex field is a subfield of the abelian field L, so its Galois closure is itself. Every nontrivial odd Artin character in that corollary is consequently a one-dimensional Dirichlet character over $\mathbb{Q}$, of conductor at most $9q$. Rational Dirichlet GRH excludes its zeros from the short region in Corollary 4.3 with $c = 1/4$. That corollary and the Colmez formula imply

$$h_F(A') \le C \log |\operatorname{disc}(L)| + C \le C' \log(2q),$$

with effective absolute constants; the dimension is always 3. Pazuki's Corollary 1.3(2), pp. 21–22, now gives $h(T) = h_\Theta(A) \le C_1 \log(2q)$. The Euclidean-norm height used there is at least the usual maximum-norm projective height, so it also gives the latter bound claimed here.

Next we count conjugates. An $\mathcal{O}$-stable lattice $I \subset L$ satisfies

$$3\mathcal{O}_L I \subset I \subset \mathcal{O}_L I.$$

Up to multiplication by $L^\times$, the ideals $\mathcal{O}_L I$ have h_L possible classes. For each class the intermediate lattices correspond to subgroups of a group of order 3^6, hence their number is bounded by an absolute constant. There are at most 2^3 CM types. On a fixed simple CM variety, two principal polarizations differ by a totally positive unit of the real subfield; equivalence under automorphisms identifies units differing by $u\overline{u}$. This quotient is a quotient of $\mathcal{O}_R^{\times,+}/(\mathcal{O}_R^\times)^2$, because $\mathcal{O}_R \subset \mathcal{O}$, and has at most 2^3 elements. These lattice and isomorphism classifications are the analytic CM classification of Milne (2020, Section 3, pp. 29–31). Choosing a symmetric representative of the polarization and a fixed theta structure introduces only an absolute finite number of choices. In particular, Pazuki's construction on pp. 28–29 changes its coordinates only by the fixed finite group of permutations and root-of-unity diagonal multipliers described there. Since the construction commutes with field automorphisms, the Galois orbit of T has size at most Ch_L.

Minkowski's ideal-class bound gives an integral ideal of norm at most $C\sqrt{|\operatorname{disc}(L)|}$ in every class. The number of ideals of a given norm n is at most the six-fold divisor function $d_6(n)$: this follows coefficientwise from the local Euler factors. The elementary bound $\sum_{n\le X} d_6(n) \le X(1+\log X)^5$ follows by summing over five of the six factors. Since $|\operatorname{disc}(L)| = 3^9 q^3$, the asserted degree bound follows.

There is also an effective analytic normalization. Since $\operatorname{Re}\tau = B_0^{-1}/2$, where B_0 denotes the fixed trace matrix in the period formula, and B_0^{-1} is integral, summing over the r^3 characteristics with $a = 0$ gives

$$\sum_{b\in\{0,\dots,r-1\}^3} \vartheta[0,b/r](\tau,0) = r^3 \sum_{n\in(r\mathbb{Z})^3} \exp(-\pi n^T \operatorname{Im}\tau\, n) \ge r^3.$$

The Fourier sum over b kills every other lattice vector, and the real phases on the remaining vectors equal 1, since r is even. Therefore some coordinate has absolute value at least 1. Evaluating these r^3 coordinates to a fixed certified accuracy chooses one with absolute value at least $1/2$. All theta constants have an absolute upper bound independent of q, from the Gaussian estimate in the period proposition. Their ratios can thus be approximated to 2^{-N} with polynomial cost in q, N.

For any normalized coordinate α, its degree is at most B and $h(\alpha) \le h(T)$. Its primitive integer minimal polynomial therefore has coefficient height at most

$$H = 2^B \exp\bigl(BC_1 \log(2q)\bigr),$$

by the relation between Mahler measure and absolute logarithmic height. The reconstruction theorem of Kannan et al. (1988, Theorem 1.19, p. 241) recovers that polynomial from a certified approximation using polynomially many bit operations in $B, \log H$. For example, $N = O(B^2 + B\log H + B\log B)$ suffices. Root separation then gives an isolating complex description of the intended root. There are only $r^6 \le 4096$ coordinates. Standard characteristic-zero factorization and primitive-element constructions combine their presentations in a common number field; Lenstra (1983, Theorem 4.5, pp. 253–254) provides deterministic polynomial-time number-field factorization. Even the crude compositum degree bound B^{4096} is polynomial in q. Intermediate heights, separation precision, and primitive-element coefficient sizes are polynomial in these degree and height bounds. This proves the cost assertion. □

This proposition constructs the characteristic-zero marked theta point. Simultaneous projective reduction and the marked good-reduction theorem below identify its reduction at every prime over the target characteristic.

Proposition 8.2. *For the sextic family of Proposition 7.1, exact algebraic recognition of its fixed-level theta point has deterministic cost polynomial in the numerical auxiliary prime q without any Riemann-hypothesis assumption. A sufficient unconditional height bound is $h(T_4) \le Cq^3$. The sharper logarithmic height bound in Proposition 8.1 uses A.*

Proof. For a maximal-order CM variety, apply von Känel and Kret (2023, Lemma 7.10, printed p. 63; equation (7.7), p. 62) with dimension three and $\epsilon = 1/2$. Its real field is fixed, and the conductor quantity in that lemma is $c = |D_L|/D_R = 3^5 q^3$. The real class number is fixed, the CM class number is at least one, the Hasse unit index is at most two, and the number of roots of unity is bounded in degree six. The explicit bound is therefore $h_F \le Cc + C\log c + C \le C'q^3$. That lemma uses the proved averaged Colmez formula and effective analytic estimates; it assumes no zero hypothesis. The degree-three lattice isogeny and its reverse change the height by a constant. Pazuki's fixed-level comparison gives $h(T_4) \le Cq^3$. The orbit degree and certified theta approximation in Proposition 8.1 are unconditional. Replacing its logarithmic height bound by Cq^3 leaves the minimal-polynomial coefficient length, recognition precision, and all algebraic-number operations polynomial in q. This proves the assertion. □

9 Simultaneous projective reduction at a given prime

Reduction of projective coordinates can require division by different common factors at different primes. The coordinate ideal retains these factors without choosing any residue-field component.

Lemma 9.1. *Let B be an explicitly presented finite étale algebra over $\mathbb{F}_p$, and let M be an explicitly presented B-module of rank one, with specified generators $x_0, \dots, x_m$. A basis w of M over B, and the coordinates $c_i \in B$ defined by $x_i = c_i w$, can be computed deterministically using polynomially many operations in $\mathbb{F}_p$. The c_i generate the unit ideal of B. No decomposition of B into fields is required.*

Proof. For $w \in M$, compute the kernel $A_w = \mathrm{Ann}_B(w)$ of the linear map $B \longrightarrow M$, $b \longmapsto bw$. Every ideal of a product of fields has an identity idempotent. Find the identity f_w of A_w by solving

$$f_w \in A_w, \qquad f_w a_j = a_j$$

for a vector-space basis a_j of A_w. These are linear equations, with a unique solution; if $A_w = 0$, take $f_w = 0$. In a field component, f_w is 1 where $w = 0$ and 0 where $w \neq 0$.

Starting with $w = 0$, successively replace

$$w \longleftarrow w + f_w x_i \qquad (0 \le i \le m).$$

An already nonzero component is preserved, and a zero component acquires the component of x_i. Since the x_i generate M, the final w is nonzero in every component and hence is a basis. Solve $c_i w = x_i$ by linear algebra. The isomorphism $B \longrightarrow M$, $b \mapsto bw$, takes the ideal generated by the c_i onto the submodule generated by the x_i, so that ideal is B. □

Proposition 9.2. *Let a number field K be given by an irreducible rational polynomial or an explicit rational multiplication table. Given a rational prime p and explicit coordinates of a point*

$$[\alpha_0 : \cdots : \alpha_m] \in \mathbb{P}^m(K),$$

there is an unconditional deterministic algorithm, polynomial in the input bit length and $\log p$, that outputs structure constants for a finite étale algebra

$$B \simeq \prod_{v \mid p} \kappa(v)$$

and elements $c_0, \dots, c_m \in B$ generating its unit ideal. Under the natural residue-field identification, the component $[c_0(v) : \cdots : c_m(v)]$ is the projective reduction of the given point at each $v \mid p$. The algorithm neither factors a polynomial over $\mathbb{F}_p$ nor selects a residue field.

Proof. Construct an order $O_0 \subset K$ of polynomial bit size. For a defining polynomial $a_n X^n + \cdots + a_0 \in \mathbb{Z}[X]$, one may use $\mathbb{Z}[a_n\theta]$; the scaled generator is integral. For a rational multiplication table on a basis e_i, choose D clearing its multiplication constants and use the lattice $\mathbb{Z}1 + \sum_i \mathbb{Z}De_i$, which is a subring.

Compute the overorder O of O_0 that is maximal at p and has p-power index over O_0. This is the unconditional deterministic polynomial-time algorithm of Buchmann and Lenstra (1994, Algorithm 6.1 and Proposition 6.2, pp. 246–247). Its operations are Frobenius kernels and ideal arithmetic; it requires no finite-field factorization. Thus

$$O \otimes \mathbb{Z}_p = O_K \otimes \mathbb{Z}_p.$$

Set $A = O/pO$. If $p^h \ge n = [K : \mathbb{Q}]$, the nilradical N of A is the kernel of the $\mathbb{F}_p$-linear map $x \mapsto x^{p^h}$. Indeed a nilpotent element of an n-dimensional algebra has nilpotence exponent at most n. Let $J \subset O$ be the inverse image of N, and compute

$$B = A/N = O/J.$$

Since O is maximal at p, its localizations above p are the corresponding discrete valuation rings of O_K. Hence B is naturally the product of their residue fields, including when p ramifies in K.

Clear the rational denominators of all coordinates by a common positive integer C, so that $\beta_i = C\alpha_i \in O$, and form the nonzero ideal

$$I = \sum_{i=0}^{m} \beta_i O.$$

Compute a lattice basis by Hermite normal form. The quotient

$$M = I/JI$$

is a B-module computable by linear algebra over $\mathbb{F}_p$: since $pI \subseteq JI$, first form I/pI, then quotient by the span of the products of generators of J and I modulo pI. The action of B is computed from multiplication in K.

At each $v \mid p$, the ideal I_v is principal in the discrete valuation ring O_v. Therefore

$$M \otimes_B \kappa(v) = I_v/\mathfrak{m}_v I_v$$

has dimension one over $\kappa(v)$. Thus M has rank one over B, and the images x_i of β_i generate it. Apply Lemma 9.1 to obtain w and coordinates $x_i = c_i w$.

To verify the reduction, choose any lift $\widetilde{w} \in I$ of w. Write

$$t_v = \min_i v(\beta_i), \qquad I_v = \mathfrak{m}_v^{t_v}.$$

The component of w in $I_v/\mathfrak{m}_v I_v$ is nonzero, so $v(\widetilde{w}) = t_v$. The coordinate equation then gives

$$c_i(v) = \overline{\beta_i/\widetilde{w}} \in \kappa(v).$$

At least one coordinate is nonzero in each component, and this is exactly the projective reduction of the original point. Different module bases multiply all coordinates by a unit in B, which preserves the projective point.

For complexity, the initial order and its discriminant have polynomial bit size. The overorder's index squared divides the initial discriminant, giving a polynomial bound for its lattice encoding; the source also supplies this bound in Section 5.9, p. 246. Ideal sums, products and bases are polynomial-time integer lattice operations. The coordinate ideal and all its product generators have polynomial bit length. Both B and M have dimension at most n. Frobenius powering uses an exponent with polynomially many bits, and support patching takes $m+1$ rounds of finite-field linear algebra. There is no loop of length p, and no loop proportional to the numerical size of a common local factor. The output has at most n^3 multiplication constants and $(m+1)n$ coordinate entries. □

The role of common factors. Reducing the original coordinates directly in O/J can erase every coordinate at a component. The quotient I/JI instead removes the minimum valuation at each prime simultaneously. The proposition outputs the full product of residue fields; it provides no rule for choosing an individual field factor.

10 Integral reduction of a marked CM theta point

Fix the theta type $\delta = (16, 16, 16)$. Let U_δ denote Mumford's moduli scheme of abelian threefolds with a normalized, relatively ample, totally symmetric invertible sheaf of type δ and a symmetric δ-theta structure. Theorem and Further Theorem of Mumford (1967, Section 6, pp. 83–84) realize U_δ as an open subscheme of the explicit projective theta-null scheme over $\mathbb{Z}[1/2]$, with its universal abelian scheme. The hypothesis that every entry of δ is divisible by 8 is satisfied. The fixed type here avoids any ambiguity between different conventions for a "level-eight" marking.

Proposition 10.1. *Let H be a number field and let $u \in U_\delta(H)$. Assume that its associated abelian variety has geometric complex multiplication. At every finite place v of H of odd residue characteristic, the normalized projective reduction of the theta-null point u belongs to U_δ. In particular, if the coordinates of u are given explicitly, Proposition 9.2 computes, deterministically in polynomial bit time, a point of U_δ over*

$$B = \prod_{v|p} \kappa(v)$$

for every odd rational prime p. Pulling back the fixed universal family gives an abelian scheme over B, without selecting a residue-field factor.

Proof. Write (A, L, α) for the marked variety. The theta structure identifies the polarization kernel with

$$H(L) \simeq (\mathbb{Z}/16\mathbb{Z})^3 \times \mu_{16}^3.$$

This group is killed by 16 and has order 16^6, so $H(L) = A[16]$. The marking therefore makes $A[16]$ unramified at every odd place: both the constant factor and μ_{16} are unramified there. Adjoining ζ_{16}, if desired, makes the full torsion rational by a fixed-degree extension unramified at these places.

The CM action is defined after a finite extension of H. Theorem 6(a) of Serre and Tate (1968, p. 503) implies potential good reduction, which consequently also holds over H_v. Corollary 2(b) to their Theorem 2, printed p. 497, now applies with $m = 16$: potential good reduction and unramified 16-torsion imply good reduction. The potential-good hypothesis is part of this implication. Equivalently, inertia has finite image on $T_2(A)$, while the principal congruence kernel modulo 4 in $\mathrm{GL}_6(\mathbb{Z}_2)$ has no nonidentity finite-order element. Thus inertia acts trivially, and the Néron–Ogg–Shafarevich criterion applies.

It remains to extend the specified marking, rather than only the underlying abelian variety. Work over the strict henselization R of the local valuation ring, with fraction field F. Let $\mathcal{A}/R$ be the good model. Since $\ker \lambda_L = A[16]$, the polarization factors uniquely as $\lambda_L = 16\lambda_0$, with λ_0 principal. The homomorphism λ_0, its inverse, and its polarization property extend over R; see Milne (1992, Propositions 2.11 and 2.14, author pp. 17–18).

Let $\mathcal{P}$ be the rigidified Poincaré bundle on $\mathcal{A} \times \mathcal{A}^\vee$. Set

$$\mathcal{M} = (\mathrm{id}, \lambda_0)^*\mathcal{P}, \qquad \mathcal{L} = \mathcal{M}^{\otimes 8}.$$

The normalized symmetric bundle $\mathcal{M}$ induces $2\lambda_0$. Hence $\mathcal{L}$ is normalized, totally symmetric, relatively ample, and induces $16\lambda_0$. Its generic fibre is the specified L. Indeed two totally symmetric bundles with the same polarization differ by an element of $\mathrm{Pic}^0[2]$ whose character on $A[2]$ is trivial. Perfect 2-Weil duality makes that element zero, and rigidification makes the isomorphism unique. This is the uniqueness observation of Mumford (1966, Section 2, pp. 304–305 and 307).

The theta group $G(\mathcal{L})$ is an extension of $\mathcal{A}[16]$ by $\mathbb{G}_m$. Over the strictly henselian ring R, its finite quotient and μ_{16} are constant. Choose a symplectic basis of the quotient and lift its basis elements to the theta group. Such lifts exist because their fibres are $\mathbb{G}_m$-torsors over the local ring R. Their sixteenth powers are units. Hensel's lemma permits rescaling the lifts so that these powers equal one. Their commutators are the prescribed pairing, and the resulting generators define an integral theta structure.

Compare this structure with the specified generic structure α. A central-fixing automorphism of the standard Heisenberg group induces a symplectic automorphism of its finite quotient. For each chosen quotient generator, its central coefficient satisfies a sixteenth-power equation with a unit right side. A solution in F is an integral unit. Thus this generic automorphism extends over R, and so does α. Its symmetry condition is an equality of morphisms true on the generic fibre and hence over R. We have therefore obtained an R-point of U_δ.

Its image in projective space is the unique extension of the generic theta-null point, since projective space is separated. Its special point consequently equals normalized projective reduction. Membership in an open subscheme can be checked after strict henselization, so the assertion holds at the original place v. Applying the simultaneous projective-reduction algorithm gives coordinates over B whose component at every $v \mid p$ lies in U_δ. The resulting B-point factors through this open, proving the final claim. □

Remark 10.2. All geometric extension arguments above certify the location of the calculated point; the algorithm does not compute local uniformizers or local theta structures. Evaluating a given fixed universal presentation and a fixed complete morphism atlas has polynomial cost. One may also spread additional fixed characteristic-zero coordinate formulas over $\mathbb{Z}[1/M]$, for an effectively computable fixed integer M independent of the input. Its finitely many prime divisors are constant exceptional cases. This does not establish the degree or height bounds needed to recognize the initial algebraic theta point, nor any Frobenius-separation or field-construction claim.

11 A complete addition atlas from one fixed theta level

The theta constants used for recognition can be converted to an explicit projective abelian variety by fixed-size linear algebra. Using level sixteen supplies both the equations and the addition law for a level-eight embedding. In particular, this step does not require selecting a root of a CM class polynomial or searching for an addition law over the finite field. The complete biquadratic addition-law mechanism is classical; compare Arene et al. (2012, Proposition 2.2 and Lemma 2.4, pp. 310–312). The construction below makes its fixed-size input requirements explicit.

Lemma 11.1. *Fix an even positive integer r, put $n = r^2$, and let τ be a point of the Siegel upper half space of degree g. For $A, B \in (\mathbb{Z}/r\mathbb{Z})^g$, represented by vectors with entries in $\{0, \dots, r-1\}$, put*

$$T_{A,B}(z) = \theta[A/r, B/r](\tau, rz), \qquad S_c(z) = \theta[c/n, 0](n\tau, nz), \quad c \in (\mathbb{Z}/n\mathbb{Z})^g.$$

Then, with $\zeta_n = e^{2\pi i/n}$,

$$T_{A,B}(z) = \sum_{j \bmod r} \zeta_n^{(A+rj)\cdot B} S_{A+rj}(z), \tag{6}$$

$$S_{A+rj}(z) = r^{-g} \sum_{B \bmod r} \zeta_n^{-(A+rj)\cdot B} T_{A,B}(z). \tag{7}$$

Thus the r^{2g} recognized theta constants at zero determine the standard level-n theta-null point by a fixed invertible linear transformation over $\mathbb{Q}(\zeta_n)$. For $g = 3, r = 4$, this has fixed degree and height cost and introduces only powers of two in its denominators.

Proof. In the defining, absolutely convergent theta series write its summation index as $k = rt + j$, with $j \bmod r$. Then $k + A/r = r(t + (A + rj)/r^2)$. The factor involving B is $\exp(2\pi i(k + A/r) \cdot B/r) = \zeta_n^{(A+rj)\cdot B}$. The remaining terms are precisely the series defining S_{A+rj}. The inverse follows from orthogonality of the characters of $(\mathbb{Z}/r\mathbb{Z})^g$. This identity also identifies the line-bundle sections: all the displayed functions of z have the common automorphy factor of the n-th power of the principal theta line bundle. In the S_c basis its standard theta-group action is translation and multiplication by characters. Consequently this is Mumford's type-$(n, \dots, n)$ marking, rather than merely a projective coordinate change with an unspecified moduli interpretation. □

Proposition 11.2. *Fix $g = 3$ and $n = 8$. Let k be a field of characteristic different from two. Suppose a projective vector*

$$v = (v_a)_{a \in (\mathbb{Z}/2n\mathbb{Z})^g}$$

is a valid point of Mumford's theta-null moduli problem of type $(2n, \dots, 2n)$, including a totally symmetric line bundle and a symmetric marking. There is a deterministic construction, using a fixed number of field arithmetic operations and equality tests, of the following data for the same underlying abelian variety in its compatible level-n embedding in $\mathbb{P}^{n^g-1}$: its identity, its inverse, a set of quadratic equations defining it scheme-theoretically in projective space, and a complete set of n^g addition laws of bidegree $(2, 2)$.

The construction also works over a finite étale $\mathbb{F}_p$-algebra of dimension N, given by structure constants, with a projective vector v whose coordinates generate the unit ideal and with the same validity promise on every geometric component. It returns the data on a partition into direct factors in deterministic time polynomial in $N, \log p$ and the input length, without finding an irreducible factor or selecting a component field.

Proof. Write $G = (\mathbb{Z}/n\mathbb{Z})^g$, $H = (\mathbb{Z}/2n\mathbb{Z})^g$, and $E = \{0, 1\}^g$. Use integer representatives for $i, j \in G$ in formulas with indices in H; changing a representative only permutes the summation over E. In compatible theta bases let s_i be the level-n sections and t_a the level-$2n$ sections. Mumford's fundamental addition formula and multiplication formula give, up to a common nonzero scalar arising from choices of trivialization,

$$s_i(x + y)s_j(x - y) = \sum_{\eta \in E} t_{i+j+n\eta}(x)t_{i-j+n\eta}(y), \tag{8}$$

$$s_i(z)s_j(z) = \sum_{\eta \in E} v_{i-j+n\eta}t_{i+j+n\eta}(z). \tag{9}$$

The scalars do not affect any of the kernels or projective maps below. For the analytic standard bases these formulas follow directly by replacing the two theta summation variables by their sum and difference and sorting their common parity. Their algebraic validity is Mumford (1966, §3, fundamental addition formula, p. 324; multiplication formula, p. 330); the proof extends to abelian schemes and to residue characteristics different from two, as checked in Mumford (1967, §6, Step III, pp. 92–94).

First form the symmetric matrix

$$P_{ij} = \sum_{\eta \in E} v_{i+j+n\eta}v_{i-j+n\eta}.$$

It is a nonzero scalar multiple of $(e_ie_j)_{i,j \in G}$, where (e_i) are the compatible level-n coordinates of zero. Thus it has rank one and is nonzero. Choose j with $P_{jj} \neq 0$; its j-th column gives the identity $[e_i]$. This is also the explicit level-forgetting morphism of Mumford (1967, §6, Step II, p. 85). No square root is needed. The inverse is $[X_i] \mapsto [X_{-i}]$.

Next use (9) to form the matrix

$$M : \mathrm{Sym}^2(k^G) \longrightarrow k^H, \qquad X_iX_j \longmapsto \sum_{\eta \in E} v_{i-j+n\eta}[i + j + n\eta].$$

Here $[a]$ denotes the standard basis vector indexed by $a \in H$, and one uses one monomial for each unordered pair $\{i, j\}$. Symmetry is consistent because $v_a = v_{-a}$. This matrix is surjective: the polarization kernel contains the full four-torsion, so Mumford (1966, §4, Theorem 1, pp. 338–340) applies to the multiplication of sections. Gaussian elimination gives a basis of $\ker M$, hence all the quadratic relations on the s_i. These relations cut out the abelian variety scheme-theoretically in projective space. This is the projective consequence of Mumford (1966, §4, the graded-ring lemmas on pp. 336–338 and Theorem 2, pp. 341–348), used explicitly in Mumford (1967, §6,

Step III, p. 94); equivalently they span the quadratic equations of the universal family in the Further Theorem on p. 84 of the latter paper. Only the saturated projective ideal is asserted here: Theorem 2 itself states a relation in degree four, whereas the preceding graded-ring argument supplies equality after saturation. On each standard affine chart the resulting ideal is radical, since this scheme is an abelian variety.

Choose by the same elimination a linear right inverse of M. Its columns define quadratic polynomials $Q_a(X)$, for $a \in H$, representing the sections t_a, up to the same common nonzero normalization. The map $[X_i] \mapsto [Q_a(X)]$ identifies this level-eight model with the given level-sixteen model; hence the construction preserves its form over the ground field, and not only its geometric isomorphism class. Form the n^g by n^g matrix of biquadratic polynomials

$$F_{ij}(X,Y) = \sum_{\eta \in E} Q_{i+j+n\eta}(X) Q_{i-j+n\eta}(Y).$$

At every geometric pair (x, y) on the variety this matrix is a common nonzero scalar multiple of the outer product $(s_i(x+y))_i (s_j(x-y))_j^{\mathsf{T}}$. Both factors are nonzero, because the complete level-n linear system gives an embedding. Consequently at least one column is nonzero, and every nonzero column represents $x+y$. The n^g columns therefore give a complete addition atlas, including over the algebraic closure; this is stronger than completeness merely on k-rational pairs.

All matrices have sizes depending only on g, n: here $|G| = 512$, $|H| = 4096$. The kernel, right inverse, and a nonzero diagonal entry of P require only arithmetic and zero tests. Over a finite étale algebra, perform these operations with the structure-constant simulation of Proposition 4.2. Split when a pivot is zero on some components and nonzero on others, and invert it only on the latter factors. Surjectivity and the rank-one identity hold on every geometric component, so all branches terminate with the stated data. The total number of leaves is at most N, and all fixed-size arithmetic has polynomial bit cost. No search through the elements of a component field occurs. □

Remark 11.3. The input promise in Proposition 11.2 is essential: satisfying a subset of theta relations does not establish that an arbitrary projective vector parametrizes an abelian variety. For a CM theta-null point, one must separately prove that its projective reduction is the point of a good-reduction abelian scheme with its marking. Once that statement has been established, the proposition supplies all the geometric input required by Corollary 4.4, with ambient dimension, equation degrees, and addition-law degrees fixed independently of the auxiliary prime and the characteristic. There is no growing-degree CM endomorphism among these inputs.

12 Extracting a ninth root from sextic CM Frobenius

Proposition 12.1. *Let $p > 3$ be prime with $p \equiv -1 \pmod 9$, and let $q \notin \{2, 3, p\}$ be prime with $q \equiv 1 \pmod 4$ and $(p/q) = -1$. Put $R = \mathbb{Q}(\theta)$, $\theta^3 - 3\theta + 1 = 0$, and $L = R(\sqrt{-3q})$. Suppose a principally polarized abelian threefold of the primitive CM type in Proposition 7.1 has been reduced with good reduction to a field $\mathbb{F}_{p^e}$, and that its Frobenius characteristic polynomial is given. Then one can deterministically find a root of $X^3 - 3X + 1$ in $\mathbb{F}_p$, and an element of order 9 in any presented quadratic extension of $\mathbb{F}_p$, in time polynomial in $e, \log p, \log q$ and the given polynomial's length. This implication uses no unproved analytic hypothesis. If $v_3(p+1) = 2$, that element generates the entire 3-primary subgroup of $\mathbb{F}_{p^2}^{\times}$.*

Proof. The prime p splits completely in R, since its class modulo 9 is -1. It also splits in $\mathbb{Q}(\sqrt{-3q})$: $(-3/p) = -1$, while quadratic reciprocity and $q \equiv 1 \pmod 4$ give $(q/p) = (p/q) = -1$. Thus it splits completely in L.

Let F denote the p^e-power Frobenius of the reduction. The characteristic-zero endomorphism algebra is L, since the CM type is primitive. Galois acts on it through $\mathrm{Aut}(L/\mathbb{Q})$, of order 6. Compatibility of good reduction with endomorphisms implies that F^6 centralizes the embedded L. For an auxiliary prime $\ell \neq p$ the injection from geometric endomorphisms tensored with $\mathbb{Q}_\ell$ into $\mathrm{End}(V_\ell)$ embeds the full algebra $L \otimes \mathbb{Q}_\ell$. Each of its field factors occurs with positive multiplicity. Their degrees sum to $\dim V_\ell = 6$, so each multiplicity is one. Its centralizer is the regular algebra itself. The rational centralizer thus has dimension at most six and contains L. Consequently F^6 is an algebraic integer $\pi \in \mathcal{O}_L$.

The Shimura–Taniyama formula, in Milne (2020, Theorem 8.1 and Corollaries 8.2–8.3, pp. 66–68), gives

$$(\pi) = \prod_{\mathfrak{p} \in \Phi_p} \mathfrak{p}^{6e},$$

where Φ_p is a conjugate of the primitive CM type, viewed as three of the six primes over p, with one prime from each complex-conjugate pair. To apply the stated maximal-order version, one may first make the degree-3 lattice isogeny to full CM and extend the local field until the isogeny and its endomorphisms are defined. Good reduction is preserved, and the formula for a power of π gives the displayed valuations after division by that power.

This support has trivial stabilizer in the cyclic group of order 6. A subgroup of order 2 would be complex conjugation and cannot preserve a CM type. Invariance under the subgroup of order 3 would make the type induced from the quadratic subfield, contrary to primitivity. Hence $\mathbb{Q}(\pi) = L$. The characteristic polynomial of F^6, computed by a resultant from that of F, is therefore the degree-six minimal polynomial of π. Factoring it over the explicit degree-six field L produces its six roots in deterministic polynomial time (Lenstra, 1983, Theorem 4.5, pp. 253–254). Any one of those roots has the same support property.

Because $[\mathcal{O}_L : \mathbb{Z}[\theta, (1 + \sqrt{-3q})/2]] = 3$, express the chosen root in this product basis with denominators dividing 3. These denominators and 2 are invertible modulo p. The known nonsquare -3 permits deterministic square-root extraction in $\mathbb{F}_p$; apply it to the square $-3q$ to obtain t. Reduce the chosen expression for π after substituting $\sqrt{-3q} = t$, obtaining a polynomial $g(\theta)$ of degree at most 2 modulo p.

Among the three embeddings restricting to either chosen sign of the quadratic subfield, a primitive CM type contains exactly one or two. Thus

$$1 \leq \deg \gcd(g(X), X^3 - 3X + 1) \leq 2.$$

A linear factor, either this greatest common divisor or its complement, yields a root $u \in \mathbb{F}_p$.

Finally solve $Z^2 - uZ + 1 = 0$ in the presented quadratic extension. Its discriminant is a nonsquare in $\mathbb{F}_p$, since otherwise that field would contain a ninth root of unity. More explicitly, any root z satisfies $z^6 + z^3 + 1 = 0$, hence has order 9 when $p \neq 3$; and $9 \nmid p - 1$. Quadratic extraction here is deterministic: choose a nonzero trace-zero element v in the presented extension by linear algebra. Then v^2 is a nonsquare in $\mathbb{F}_p$, so $(u^2 - 4)/v^2$ is a square there. Extract its square root using the same known nonsquare -3, multiply by v, and use the quadratic formula. The polynomial and field arithmetic and all square-root computations have the claimed polynomial cost. If $v_3(p + 1) = 2$, then $v_3(p^2 - 1) = 2$, proving the last claim. □

The input promise of good CM reduction is essential to this proposition. In the proposed construction it must follow from an integral model with a complete addition law. For $v_3(p+1) > 2$, this procedure gives order 9, not the maximal 3-power order; that lifting problem is also a separate step.

13 A fixed-dimensional CM construction of ninth roots

The following construction combines classical CM height theorems, algebraic-number recognition, Mumford's theta equations, and Pila's Frobenius algorithm. The use of an auxiliary quadratic conductor allows a primitive sextic CM type, while computation over the full finite residue algebra avoids selecting a root of a CM invariant polynomial. The constituent geometric and computational tools are classical; Theorem 1.5 is the specific composition proved here. Pila's own root-of-unity Theorem D treats a fixed odd prime order and primes congruent to one modulo that order (Pila, 1990, Theorem D, author p. 2; Section 5, pp. 14–18); it is not the ninth-root statement below.

Proof of Theorem 1.5

Proof. Put $\lambda = 2 + \lceil \log_2 p \rceil$. The congruence on p implies $p > 3$, $\operatorname{ord}_3(p) = 2$, and $(-3/p) = -1$.

Auxiliary prime and exact theta point. Apply Lemma 6.1 with $r = 3$ and $d = 2$. By testing primes in increasing order, find

$$q \equiv 1 \pmod 4, \qquad (p/q) = -1, \qquad q \notin \{2, 3, p\}, \qquad q \le C\lambda^2.$$

For completeness, this specialization needs only the four Dirichlet characters obtained by multiplying the indicator of 1 (mod 4) by the quadratic nonresidue indicator for p. Their smoothed prime count has main term $x/4$, and its error under A is $O(\sqrt{x}\log(2p))$. Even powers have zero nonresidue indicator; the remaining higher prime powers contribute $O(x^{1/3}(\log x)^2)$. Thus a sufficiently large absolute multiple of λ^2 contains a suitable prime. This search can be implemented in $O(\lambda^6)$ bit operations using trial division for its small candidates and integer arithmetic; the bound is intentionally coarse.

Let

$$R = \mathbb{Q}(\theta), \quad \theta^3 - 3\theta + 1 = 0, \qquad L = R(\sqrt{-3q}).$$

Proposition 7.1 constructs the period matrix of a principally polarized simple abelian threefold with primitive CM type and multiplication by $\mathbb{Z}[\theta, (1+\sqrt{-3q})/2]$. Its lattice has index 3 in $\mathcal{O}_L$, and $|\operatorname{disc}(L)| = 3^9 q^3$. Apply Proposition 8.1 with $r = 4$, obtaining all 4096 theta ratios exactly in a number field. This step has deterministic polynomial cost in the numerical q, rather than in an unspecified CM preprocessing oracle.

Apply Lemma 11.1, adjoining ζ_{16} by characteristic-zero algebraic computation. It transforms the recovered point into the standard level-sixteen theta-null point v, defined over an explicit number field K. It is a point of Mumford's moduli problem of type $(16, 16, 16)$, including its totally symmetric line bundle and symmetric marking. All number-field operations here are in characteristic zero.

Reduction without choosing a residue field. Apply Proposition 9.2 to the homogeneous coordinates of v at p. It computes structure constants for

$$\mathcal{B} = \prod_{\mathfrak{p}\mid p} \kappa(\mathfrak{p})$$

and a coordinate vector generating its unit ideal. This normalization removes the separate common coordinate valuations at all primes above p; reducing one arbitrary affine chart would not suffice. Proposition 10.1 shows that every component is still a valid marked abelian threefold. Its proof uses potential good reduction of CM varieties and unramified full sixteen-torsion, followed by extension of the specified marking.

Proposition 11.2 now constructs, over a partition of $\mathcal{B}$, the same abelian threefold in a level-eight embedding in $\mathbb{P}^{511}$. It gives quadratic defining equations with radical affine-chart ideals,

the identity, the inverse, and 512 complete biquadratic addition laws. The construction uses multiplication-matrix kernels and right inverses; its matrix dimensions are absolute constants. In particular, the equation degrees and the addition atlas do not grow with q.

Frobenius and prime separation. Separate $\mathcal{B}$ into blocks whose field components have a common residue degree e. This uses the Frobenius-linear construction in Proposition 4.2, not irreducible factor selection. Run the deterministic Pila algorithm with field size p^e, using Corollary 4.4. At every equality test split a block only as required by the tested element. At termination there is a nonempty list of blocks, each with an integer Frobenius characteristic polynomial valid on all its components. Choose any polynomial on this list. Choosing this integer output does not choose a residue-field embedding.

All such components are reductions of conjugates of the same primitive CM construction. The field L is Galois of degree six, and p splits completely in it: it splits in R because $p \equiv -1 \pmod 9$, and

$$(-3q/p) = (-3/p)(q/p) = (-1)(-1) = 1$$

by quadratic reciprocity and $q \equiv 1 \pmod 4$. Consequently Proposition 12.1 applies to any selected integer Frobenius polynomial. Briefly, the sixth power of Frobenius is a primitive element $\pi \in L$, whose ideal is supported at the three primes in a primitive CM type. Compute its degree-six polynomial and factor it over the explicit characteristic-zero field L. Express a root in the product order with denominators dividing 3, extract $\sqrt{-3q}$ modulo p using the known nonsquare -3, and substitute either quadratic sign. The resulting greatest common divisor with f has degree one or two. A linear factor, or the complementary linear factor, yields u.

The same proposition solves $Z^2 - uZ + 1$ in the presented quadratic field and yields an element z of exact order nine. The canonical quadratic presentation is valid because $(-3/p) = -1$. For another presentation, a nonzero trace-zero vector is found by linear algebra, and its square supplies the quadratic coordinate conversion using only square roots in $\mathbb{F}_p$ with a known nonsquare. This completes the first two outputs.

The cubic extension when the valuation is two. Now suppose $v_3(p+1) = 2$. Then $v_3(p^2-1) = 2$, so an element of exact order nine in $\mathbb{F}_{p^2}$ is not a cube. We have $z + z^{-1} = u$. If H had a root $\beta \in \mathbb{F}_p$, choose a root w of $W^2 - \beta W + 1$ in $\mathbb{F}_{p^2}$. Then

$$w^3 + w^{-3} = \beta^3 - 3\beta = u,$$

so w^3 is one of the roots z, z^{-1} of $Z^2 - uZ + 1$. This contradicts their being noncubes. Hence H has no root in $\mathbb{F}_p$, and a cubic with no root is irreducible. Equivalently, a cube root of z has order twenty-seven and cannot lie in $\mathbb{F}_{p^2}$. The displayed polynomial is the promised explicit presentation.

Uniform bit bound. Here are parameter bounds sufficient to make the composition's bit complexity explicit. Put

$$B_0 = C_0 q^{3/2}(1 + \log q)^5, \qquad N = 8\lceil B_0 \rceil^{4096}.$$

The factor eight allows $\mathbb{Q}(\zeta_{16})$. The number field degree and the dimension of $\mathcal{B}$ are at most N. Let S be the total bit length of its characteristic-zero presentation and theta coordinates. The recognition proof, including its fixed number of composita, gives

$$S \le C(B_0 \log(2q) + 2)^\kappa$$

for an absolute effectively computable $\kappa \geq 4096$. The degree and coefficient bounds used by recognition are specified in Proposition 8.1; in particular its requested precision is polynomial in B_0 and $B_0 \log(2q)$.

Let $c_0 \geq 8$ be a fixed exponent bounding the deterministic recognition, characteristic-zero, lattice, projective-reduction, and final fixed-degree operations by $C(S+\lambda+2)^{c_0}$. Let $\Delta_0 \geq 2$ be a fixed exponent for Pila's algorithm on the model just constructed, counting its coefficient-field operations and its non-field bit operations. These are effective absolute constants supplied by the cited algorithms; neither is an exponent depending on the input. The Pila parameters may be taken as ambient dimension 511, at most $512 \cdot 513/2$ quadratic equations, at most 512 addition charts, total addition degree 4, and dimension 3. The explicit parameter-dependent bound for its exponent is in Pila (1990, Section 4, author pp. 12–14).

Arithmetic, inversion, support splitting, changes of basis, and reconstruction of structure constants in an algebra of dimension at most N have the conservative bound $O(N^6\lambda^3)$ per simulated step with school arithmetic and Gaussian elimination. There are at most N branches, and each field computation has $e \leq N$. After increasing c_0 for elementary preprocessing, a bit bound for the entire construction is

$$C\left\{\lambda^6 + (S+\lambda+2)^{c_0} + N^7\lambda^3(N\lambda+2)^{\Delta_0}\right\}. \tag{$*$}$$

Since $q \leq C\lambda^2$, we may bound $B_0 \leq C\lambda^8$, $N \leq C\lambda^{32768}$, and $S \leq C\lambda^{9\kappa}$. For example the absolute exponent

$$C_* = 10 + \max\{6,\, 9\kappa c_0,\, 229379 + 32769\Delta_0\}$$

suffices in the theorem. Formula $(*)$ records the source of the bound; it is not a claim that these deliberately large exponents are practical.

Finally, the two conditional uses of A were the quadratic auxiliary-prime estimate and the effective height bound for an abelian CM field. In the latter, the odd Artin characters in the Colmez expression are Dirichlet characters over $\mathbb{Q}$, because $L/\mathbb{Q}$ is abelian. The algebraic recognition, marked reduction, complete atlas, Frobenius computation, and extraction from π are unconditional once their stated inputs have been obtained. Proposition 8.2 allows the height step to use an unconditional polynomial bound instead; then only the auxiliary-prime step needs A. □

Scope of the output. For $v_3(p+1) > 2$, the construction still produces an element of exact order nine and a root of $X^3 - 3X + 1$. It does not produce an element of maximal three-power order in $\mathbb{F}_{p^2}^{\times}$, and the displayed cubic is then not an irreducible extension polynomial. The construction has fixed dimension three; it gives no polynomial-in-degree extension of the method to unbounded cyclotomic conductor.

Example 13.1. At $p = 17$, the auxiliary quadratic prime can be $q = 5$: $5 \equiv 1 \pmod 4$ and $(17/5) = -1$. The roots of $X^3 - 3X + 1$ modulo 17 are $7, 13, 14$. To display the final conversion, take $u = 7$ and write $s^2 = -3$ in $\mathbb{F}_{17^2}$. Since $u^2 - 4 = 11 = 2s^2$ and $6^2 = 2 \pmod{17}$,

$$z = \frac{7+6s}{2} = 12 + 3s, \qquad z^3 = 8 + 8s, \quad z^6 = 8 + 9s, \quad z^9 = 1.$$

Thus z has exact order nine. The irreducible cubic in Theorem 1.5 is $Y^3 - 3Y - 7$. Direct substitution of the seventeen field elements also verifies that it has no root. These calculations illustrate the exact final conversion; they are not an execution of the theta recognition and Frobenius algorithm.

14 The CM construction for fixed wild cyclotomic parameters

The same construction extends to every fixed odd prime and every fixed wild layer. The real field needed here is the wild subfield of the cyclotomic field, rather than its full maximal real subfield. This distinction allows every tame residue degree. The constants in this section depend on the fixed parameters; their dependence is displayed sufficiently to distinguish this result from a polynomial bound uniform in the cyclotomic degree.

Proof of Theorem 1.3

Proof. All characteristic-zero data depending only on r, b may be computed once by finite exact algorithms. In particular, realize R_b as the fixed field of the order-$r-1$ subgroup of $\mathrm{Gal}(K_b/\mathbb{Q})$, compute an integral basis, and choose an integral primitive element by enumeration. Put $I = [\mathcal{O}_{R_b} : \mathbb{Z}[\theta]]$. Both I and the discriminant of f are now fixed integers. The finitely many input primes dividing $2rI$ can be handled by exhaustive finite-ring arithmetic. This has a constant cost depending on r, b. The existence of the required outputs in these cases follows from the cyclotomic reduction at $p \neq r$ used at the end of the proof. We henceforth assume $p \nmid 2rI$.

The order of p modulo r^b divides d, by the input valuation condition. Its Frobenius therefore lies in the tame subgroup, so p splits completely in R_b. The field is totally real, cyclic, and of odd degree g. Its discriminant is

$$D_R = r^{E_{r,b}}, \qquad E_{r,b} = \sum_{j=1}^{b-1} (j+1)\varphi(r^j) = bg - \frac{g-1}{r-1} - 1. \tag{10}$$

Indeed its characters of exact order r^j have conductor r^{j+1}, and there are $\varphi(r^j)$ of them; the conductor–discriminant identity gives (10).

An auxiliary imaginary quadratic field. Let $\lambda = 2 + \lceil \log_2 p \rceil$. Under A, a prime q can be found with

$$q \equiv 3 \pmod 4, \qquad (p/q) = 1, \qquad q \notin \{p, r\}, \qquad q \le C\lambda^2, \tag{11}$$

where C is absolute. Here is the precise character content. On integers coprime to $4p$, the desired indicator is

$$\tfrac{1}{4}(1 - \chi_{-4})(1 + \chi_{D_p}), \qquad D_p = \begin{cases} p, & p \equiv 1 \pmod 4, \\ 4p, & p \equiv 3 \pmod 4, \end{cases}$$

and it is zero on other integers. The four characters are Dirichlet characters of modulus $4p$, with only the first principal. The log-tent explicit formula in Lemma 25.5 under A gives main term $x/4$ and error $O(\sqrt{x}\log(2p))$. Even prime powers have zero indicator; higher odd powers contribute $O(x^{1/3}(\log x)^2)$. Removing the one possible prime r, including its powers, costs $O((\log x)^2)$. Taking a sufficiently large absolute multiple of λ^2 proves (11). Trial division and quadratic-symbol tests suffice for the small-prime search, at bit cost $O(\lambda^6)$. Quadratic reciprocity gives $(-q/p) = (p/q) = 1$, so p splits in $E = \mathbb{Q}(s)$, where $s^2 = -q$.

Set $L = R_b E$ and $\mathcal{O} = \mathbb{Z}[\theta, (1+s)/2]$. The discriminants of R_b and E are coprime. Hence

$$\mathcal{O}_L = \mathcal{O}_{R_b}\mathcal{O}_E, \qquad |D_L| = D_R^2 q^g, \qquad [\mathcal{O}_L : \mathcal{O}] = I^2. \tag{12}$$

The usual coprime-discriminant assertion can also be checked locally: at every prime at least one of the two integer algebras is unramified, and their tensor product is the maximal order of the linearly disjoint compositum. The extension $L/\mathbb{Q}$ is abelian of degree $2g$, has odd conductor dividing $r^b q$, and p splits completely in it.

Explicit principal periods and a primitive type. Order the real roots $\theta_1 < \cdots < \theta_g$ of f, put $S = (\theta_i^j)_{1\le i\le g,\,0\le j<g}$, and set

$$B_{jk} = \operatorname{Tr}_{R_b/\mathbb{Q}} \frac{\theta^{j+k}}{f'(\theta)} \quad (0 \le j,k < g).$$

The trace-dual formula for a monogenic order shows that B is integral. Its entries with $j+k < g-1$ vanish, and its anti-diagonal entries are one. Thus $\det B = \pm 1$. For $\epsilon_i = \operatorname{sign} f'(\theta_i)$, use the embeddings $\phi_i(s) = \epsilon_i i\sqrt{q}$ and the form

$$E(x,y) = \operatorname{Tr}_{L/\mathbb{Q}}\left(\frac{\overline{x}y}{f'(\theta)s}\right) \quad (x,y \in \mathcal{O}).$$

The product-order trace dual is $(f'(\theta)s)^{-1}\mathcal{O}$. The form is consequently integral, alternating, and unimodular. At the chosen embeddings its Riemann positivity follows from $\phi_i((f'(\theta)s)^{-1}) = -i/(|f'(\theta_i)|\sqrt{q})$. The symplectic period computation therefore gives

$$\tau = \tfrac{1}{2}B^{-1} + \tfrac{i\sqrt{q}}{2}S^{-1}\operatorname{diag}(|f'(\theta_i)|)S^{-T} \in \mathcal{H}_g. \tag{13}$$

In particular $\operatorname{Im}\tau$ is $\sqrt{q}$ times one fixed positive definite matrix.

The CM type is primitive. To see this without assuming any particular ordering of the Galois action on the real roots, write $\operatorname{Gal}(L/\mathbb{Q}) = G \times \{1,\iota\}$, where $|G| = g$ is odd and ι is complex conjugation. A stabilizer element with negative second component has an odd power equal to ι, which cannot preserve a CM type. The stabilizer is therefore a subgroup of G. Derivative signs at ordered simple roots alternate, and g is odd; the number of positive signs is $(g+1)/2$. An invariant subset of the regular G-set is a union of orbits of the stabilizer. Its order must thus divide both g and $(g+1)/2$, which are coprime. The stabilizer is trivial. This is the primitive CM criterion, so the principally polarized variety defined by (13) is simple and has rational endomorphism algebra L.

Algebraization with fixed dimension. We give the parameter bounds needed to apply the preceding theta construction in genus g. Its ingredients are valid in every fixed genus. The analytic theta sums at denominator four have 16^g coordinates. They satisfy the tail bound

$$\operatorname{Tail}(K) \le C_{r,b}\exp(-c_{r,b}\sqrt{q}\,K^2) \qquad (K \ge 2).$$

Thus certified precision 2^{-M} costs a polynomial in q, M with exponent depending only on r, b. Moreover,

$$\sum_{a\in(\mathbb{Z}/4\mathbb{Z})^g} \vartheta[0,a/4](\tau,0) = 4^g \sum_{n\in(4\mathbb{Z})^g} \exp(-\pi n^T \operatorname{Im}\tau\, n) \ge 4^g. \tag{14}$$

The phases in (14) equal one because B^{-1} is integral. Thus a nonzero normalization coordinate can be selected by certified approximation with a fixed positive lower bound.

Saturating the lattice from $\mathcal{O}$ to $\mathcal{O}_L$ changes the Faltings height by at most a constant depending on I, g. The abelian Colmez formula applies without a conductor-at-two ambiguity (Colmez, 1993, Theorem 0.3(ii), p. 631; Theorem III.2.9(i), p. 680). The bound in Wei (2023, Corollary 4.3, pp. 14–15; Theorem 4.8, p. 18), followed by Pazuki (2012, Corollary 1.3(2), pp. 21–22), gives theta height at most $C_{r,b}\log(2q)$. Every conditional zero estimate here concerns a Dirichlet character over $\mathbb{Q}$ of conductor at most $r^b q$, because the CM field and its reflex field are abelian.

For clarity, the orbit-degree bound also remains polynomial. Any $\mathcal{O}$-stable lattice J lies between $I^2\mathcal{O}_L J$ and $\mathcal{O}_L J$, giving a fixed number of choices for each ideal class. The real units congruent to one modulo I^2 lie in the order. Their squares are norms of automorphism units, so

the number of inequivalent principal polarizations is bounded by a constant depending on the fixed field and order. Fixed-level markings and CM types likewise introduce only fixed factors. Minkowski's class bound, together with $\sum_{n\le Y} d_{2g}(n) \le Y(1+\log Y)^{2g-1}$, therefore gives the theta-point degree bound

$$B_0 = C_{r,b} q^{g/2} (1+\log q)^{2g-1}. \tag{15}$$

This is the CM lattice classification of Milne (2020, Section 3, pp. 29–31), with the fixed order index accounted for explicitly.

The degree bound (15), the logarithmic height bound, and (14) permit exact algebraic recognition of every coordinate ratio by Kannan et al. (1988, Theorem 1.19, p. 241). For example coefficient height $2^{B_0}\exp(C_{r,b}B_0\log(2q))$ and precision polynomial in its logarithm and B_0 suffice. Characteristic-zero factorization and compositum construction are deterministic polynomial-time operations (Lenstra, 1983, Theorem 4.5, pp. 253–254). Adjoin ζ_{16} and apply the fixed Fourier identity of Lemma 11.1. The resulting exact level-sixteen theta point has coefficient-field degree at most

$$N = 8\lceil B_0 \rceil^{16^g}. \tag{16}$$

The simultaneous projective reduction and integral marked-CM reduction arguments apply in any fixed genus: full sixteen-torsion, potential good reduction, and the torsion-free principal congruence kernel give good reduction at odd p; the chosen symmetric marking extends. Mumford's moduli theorem and universal family apply to type $(16,\dots,16)$ (Mumford, 1967, pp. 83–84,92–94). The multiplication-matrix construction in Proposition 11.2, with g in place of three, gives a level-eight embedding with 8^g homogeneous coordinates, quadratic equations with radical affine-chart ideals, and 8^g complete biquadratic addition charts. Surjectivity of the multiplication map and projective equality are the same dimension-independent theorems of Mumford (1966, pp. 324,330,338–341) and Mumford (1967, p. 94). All matrix sizes are fixed here.

Compute over the full finite residue algebra, separate equal residue-degree blocks by Frobenius linear algebra, and use the arithmetic/equality simulation of Pila's algorithm. This gives at least one integer Frobenius characteristic polynomial of a good CM reduction over some $\mathbb{F}_{p^e}$, with $e \le N$. It does not choose a field factor of the residue algebra. Pila's exponent depends on the genus and model sizes, which are fixed, rather than on q (Pila, 1990, Theorem A, p. 1; Section 4, pp. 8–14).

Separating all real embeddings by conjugate supports. Let F be the selected Frobenius. Its $2g$-th power centralizes the CM field, since the Galois action on L has order dividing $2g$. The centralizer is L: on a Tate module at $\ell \ne p$, every field factor of $L\otimes\mathbb{Q}_\ell$ occurs by injectivity of the scalar-extended endomorphism realization. Their degrees sum to $2g = \dim V_\ell$, so each occurs once, as in the sextic case. Thus $F^{2g} = \pi \in \mathcal{O}_L$. The Shimura–Taniyama formula gives

$$(\pi) = \prod_{\mathfrak{p}\in\Phi_p} \mathfrak{p}^{2ge}, \tag{17}$$

where Φ_p is a translate of the primitive CM type (Milne, 2020, Theorem 8.1 and Corollaries 8.2–8.3, pp. 66–68). One can pass to the fixed-degree isogeny with full CM and a finite extension to apply the maximal-order formulation, then divide valuations by the extension degree. The type's trivial stabilizer implies $\mathbb{Q}(\pi) = L$.

Compute the characteristic polynomial of F^{2g} from that of F, and factor it over the explicit number field L in characteristic zero. It has exactly the $2g$ conjugate roots $\pi_1,\dots,\pi_{2g}$ in L. Since $I^2\mathcal{O}_L \subset \mathcal{O}$, their product-basis expressions have denominators dividing I^2. Use a quadratic nonresidue of $\mathbb{F}_p$, obtainable under A by testing $O(\lambda^2)$ small integers, to extract a square root t of $-q$ in $\mathbb{F}_p$. Substituting $s = t$ in every π_j yields $G_j(X) \in \mathbb{F}_p[X]/(f)$; all denominators are units because $p \nmid 2I$.

The zero sets of the G_j separate the roots of f. Indeed, identify the primes above p with the regular Galois set of L. For distinct primes P, Q restricting to the chosen quadratic sign, equality of membership in every translate of Φ_p would say that translation by $QP^{-1} \neq 1$ preserves Φ_p, a contradiction. By (17), membership is exactly the vanishing of the appropriate conjugate of π. Repeatedly replace each current factor h of f by the nonconstant members of

$$\gcd(h, G_j), \qquad h/\gcd(h, G_j).$$

After all $2g$ conjugates, every factor is linear. Thus a root $u \in \mathbb{F}_p$ and the corresponding residue homomorphism on $\mathcal{O}_{R_b}/p\mathcal{O}_{R_b}$ are computed using only greatest common divisors. Integrality of the basis conversion follows from $p \nmid I$.

The formal Teichmüller unit. The fields R_b and T_r have coprime degrees, so $K_b = R_b T_r$. Their integer rings generate the full integer ring locally away from r: both extensions are unramified there, and their tensor product is finite étale with fraction field K_b. Consequently a fixed product-basis expression

$$\zeta_{r^b} = A(\theta, \zeta_r) \tag{18}$$

has denominators with prime factors only among those of rI. It can be precomputed in characteristic zero. Evaluate (18) at (u, Z) in $\mathcal{D}_p$ to obtain U. The compatible choice $\zeta_r = \zeta_{r^b}^{r^{b-1}}$ also gives $U^{r^{b-1}} = Z$. Every component is the reduction of a primitive r^b-th root, and remains primitive because $p \neq r$. The automorphism of K_b fixing R_b and sending ζ_r to ζ_r^c sends ζ_{r^b} to $\zeta_{r^b}^{\omega_b(c)}$. Reducing this identity proves (1). All components of $\mathcal{D}_p$ have degree d, which proves the final maximal-order assertion when the valuation equals b.

Bit bound and analytic scope. Let S_0 denote the characteristic-zero presentation length. Recognition and the fixed number 16^g of composita give

$$S_0 \leq C_{r,b}(B_0 \log(2q) + 2)^{\kappa_{r,b}}$$

for an effective fixed exponent $\kappa_{r,b}$. Let c_0 bound the polynomial-time characteristic-zero, lattice, and reduction algorithms, and let Δ_g bound Pila's arithmetic and non-field operations for the model above. Including support splitting and changes of structure constants, a conservative total bit bound is

$$C_{r,b}\{\lambda^6 + (S_0 + \lambda + 2)^{c_0} + N^7\lambda^3(N\lambda + 2)^{\Delta_g}\}. \tag{19}$$

For $\nu = (3g-1)16^g$, (15)–(16) imply $N \leq C_{r,b}\lambda^\nu$ and $S_0 \leq C_{r,b}\lambda^{3g\kappa_{r,b}}$. For example

$$10 + \max\{6, 3g\kappa_{r,b}c_0, 7\nu + 3 + (\nu+1)\Delta_g\}$$

is an effective sufficient exponent in the theorem. The fixed precomputation and finitely many excluded primes are absorbed in $C_{r,b}$. The only conditional analytic steps used Dirichlet characters over $\mathbb{Q}$: conductors dividing $4p$ for (11), conductor p for quadratic nonresidue selection, and conductors at most $r^b q$ for the abelian CM height bound. □

Proof of Corollary 1.4

Proof. The unit U of Theorem 1.3 has the maximal r-power order and the equivariance required by Theorem 1.7. The unconditional reduction in that theorem therefore supplies the degree-r field. Explicitly, form $\mathcal{D}_p[Y]/(Y^r - U)$, let c act by $Z \mapsto Z^c$ and $Y \mapsto Y^{\omega_{b+1}(c)}$, and compute the fixed algebra by linear algebra. The theorem proves that this algebra is a field of dimension r. Any basis element outside its scalar subspace generates it, since r is prime. Its minimal polynomial, computed from linear dependences among its powers, is the desired monic irreducible polynomial. These operations have bit cost polynomial in $r, \log p$ and preserve the displayed bound. □

The theorem outputs an element of the formal algebra $\mathcal{D}_p$; it does not implicitly select an embedding of a component into an independently presented extension field. It applies to all tame residue degrees d, including odd ones. Its dimension is r^{b-1}, and (19) does not supply a polynomial bound uniform in r, b.

Example 14.1. The fixed-parameter construction also has a noncubic endpoint that can be checked entirely by polynomial arithmetic. Let $r = 5$, $b = 2$, $p = 43$, and $\zeta = \zeta_{25}$. The period $\theta = \zeta + \zeta^7 + \zeta^{18} + \zeta^{24}$ satisfies

$$f(T) = T^5 - 10T^3 + 5T^2 + 10T + 1.$$

Its roots modulo 43 are $10, 19, 24, 36, 40$. Taking $\theta \mapsto 10$, the resulting element of $\mathbb{F}_{43}[Z]/\Phi_5(Z)$ is

$$U = 12 + 31Z + 18Z^2 + 32Z^3, \qquad U^5 = Z, \qquad U^{25} = 1.$$

For $c = 1, 2, 3, 4$, the lifts $\omega_2(c)$ are $1, 7, 18, 24$, respectively, and substitution verifies $U(Z^c) = U(Z)^{\omega_2(c)}$.

Here is an exact characteristic-zero certificate for the transfer. Write $\zeta_5 = \zeta^5$ and define

$$\begin{aligned} A_0(T) &= 25 + 25T - 23T^2 - T^3 + 2T^4, \\ A_1(T) &= 19 + 12T + T^2 - 3T^3 - T^4, \\ A_2(T) &= 34 - T - 31T^2 + 2T^3 + 3T^4, \\ A_3(T) &= 47 + 54T - 62T^2 - 3T^3 + 6T^4. \end{aligned}$$

Reduction modulo $\Phi_{25}(X)$ proves both $f(\theta) = 0$ and

$$35\zeta = \sum_{j=0}^{3} A_j(\theta)\zeta_5^j.$$

The period is nonrational, so its minimal polynomial has degree five, the prime degree of the tame fixed field. The displayed polynomial is therefore its minimal polynomial. Its discriminant is $49 \cdot 5^8$, whereas the field discriminant is 5^8 by the conductor–discriminant formula; the period order has index seven. The identity can consequently be reduced at 43, where 35 is invertible. Substituting $T = 10$ and reducing its four coefficients modulo 43 gives the stated U. Multiplication modulo Φ_5 verifies all its power and equivariance identities. This checks the exact algebraic conversion from a period root; it does not execute the CM search.

15 Fixed higher three-power strata

The order-nine construction extends to every fixed three-adic valuation. Its dimension dependence must remain explicit. Write $C_m \in \mathbb{Z}[X]$ for the monic Chebyshev polynomial defined by

$$C_m(Z + Z^{-1}) = Z^m + Z^{-m}; \qquad C_3(X) = X^3 - 3X.$$

Proposition 15.1. *Let $p > 3$, $p \equiv -1 \pmod 9$, and $a = v_3(p+1)$. Suppose $z_b \in \mathbb{F}_{p^2}$ has exact order 3^b, with $2 \le b \le a$, and put $u_b = z_b + z_b^{-1} \in \mathbb{F}_p$. If $b < a$, the polynomial $X^3 - 3X - u_b$ splits into three distinct linear factors over $\mathbb{F}_p$. Choosing any root, followed by quadratic extraction, yields a cube root of z_b and hence an element of order 3^{b+1}. Conversely a cube root yields a root of the displayed polynomial. If $b = a$, that polynomial is irreducible over $\mathbb{F}_p$.*

Proof. A cube root w of z_b has order 3^{b+1}. When $b < a$, all three such roots lie in the norm-one subgroup of $\mathbb{F}_{p^2}^{\times}$, so their traces $w + w^{-1}$ belong to $\mathbb{F}_p$ and satisfy the displayed cubic. Equal

traces make the corresponding roots equal or inverse. The inverse case would imply $z_b = z_b^{-1}$, which is impossible. Conversely, from a trace root v, solve $W^2 - vW + 1 = 0$ and choose the root whose cube is z_b by direct testing. The known nonsquare -3 supplies deterministic quadratic extraction. If $b = a$, a root of the cubic would produce an element of order 3^{a+1} in $\mathbb{F}_{p^2}$, a contradiction. □

Theorem 15.2 (Every fixed three-power stratum)**.** *Fix $b \geq 2$. Under rational Dirichlet GRH, there is an effective deterministic algorithm which, given a prime $p \equiv -1 \pmod{3^b}$, constructs an element of exact order 3^b in any presented quadratic extension of $\mathbb{F}_p$, in bit time*

$$O_b((1 + \log p)^{C_b}).$$

Here C_b is an effective constant depending on b. If $v_3(p+1) = b$, the algorithm also constructs an irreducible cubic over $\mathbb{F}_p$. No uniform bound on C_b as b varies is asserted.

Proof of Theorem 15.2. Apply Theorem 1.3 with $r = 3$. The condition $p \equiv -1 \pmod{3^b}$ gives $v_3(p^2 - 1) \geq b$, so it constructs an element U of exact order 3^b in the quadratic field $\mathbb{F}_p[Z]/(Z^2 + Z + 1)$. Its tame equivariance says $U^p = U^{-1}$, because the nontrivial tame lift is minus one. Thus $u_b = U + U^{-1}$ belongs to $\mathbb{F}_p$.

To transfer to another presented quadratic field, choose a nonzero trace-zero vector v by linear algebra. Then v^2 is a nonsquare in $\mathbb{F}_p$, and $-3/v^2$ is a square. The known nonsquare -3 permits deterministic square-root extraction in $\mathbb{F}_p$. If $a^2 = -3/v^2$, the map $Z \mapsto (-1 + av)/2$ is an isomorphism of the two quadratic presentations. Evaluating U under it gives the required root of unity. All conversion operations are polynomial in $\log p$. When $v_3(p+1) = b$, Proposition 15.1 proves that $Y^3 - 3Y - u_b$ is irreducible. The constants of Theorem 1.3 depend only on the fixed pair $(3, b)$, proving the asserted bit bound. □

16 Completely split residue maps in fixed degree

The CM construction can be made quantitative when the field varies in a fixed degree. The numerical multiplication-table height and discriminant are charged explicitly; the exponent is allowed to depend on the degree.

For comparison, Rónyai proved a deterministic bound under GRH for factoring reductions of integral polynomials in terms of their splitting fields (Rónyai, 1992). For a fixed defining polynomial, factorization at a completely split prime outside its index gives a residue map by selection of a linear factor. Here the CM construction makes explicit that rational Dirichlet GRH suffices with an explicit numerical bound when the totally real Galois field varies in fixed degree.

Proposition 16.1 (Quantitative completely split residue maps)**.** *Fix $m \geq 3$. Given an integral basis $b_1, \ldots, b_m$ and its integer multiplication table $b_i b_j = \sum_k c_{ijk} b_k$ for a totally real Galois field $K/\mathbb{Q}$ of degree m, put*

$$H = \max\{2, |c_{ijk}| : 1 \leq i, j, k \leq m\}, \qquad D_K = \operatorname{disc}(K) > 0.$$

Under rational Dirichlet GRH, a unital ring homomorphism $\mathcal{O}_K \longrightarrow \mathbb{F}_p$ can be constructed deterministically for every input prime p splitting completely in K, in bit time

$$c_m(H + D_K + 1 + \log p)^{C_m},$$

with effective constants c_m, C_m depending only on m. The field and its integral basis may vary. No residue map, prime above p, finite-field factor, CM model, or field-dependent advice is an input. Fixed geometric precomputation may depend on m alone. The bound is polynomial in the numerical H, D_K, not their logarithms, and no degree-independent exponent is asserted.

Proof. All constants in this proof depend on m alone. We charge every field-dependent quantity to the numerical parameters H, D_K. Fixed-degree characteristic-zero arithmetic and real root isolation recover the embeddings and the Galois automorphisms from the multiplication table. For example, a primitive element is found among a degree-dependent finite grid of small integral combinations of the basis elements, and its polynomial has coefficients $H^{O_m(1)}$; characteristic-zero factorization over K then recovers all its automorphisms. The coordinates of 1 are recovered by linear algebra if necessary.

A controlled sign element and lattice index. Let $F = (\rho_i(b_j))$ be the real embedding matrix. The eigenvalues of the multiplication matrix of b_j are its conjugates, so

$$|\rho_i(b_j)| \le mH, \qquad \|F\|_2 \le m^2 H, \qquad |\det F| = \sqrt{D_K} \ge 1.$$

The adjugate formula gives $\|F^{-1}\|_\infty \le c_m H^{m-1}$. Set $s_1 = -1$, $s_i = 1$ for $i > 1$, and $v_i = -s_i$. Approximate $8m^2 H F^{-1} v$ within $1/4$ in each coordinate and round the rational approximations to integers n_j. With $\alpha = \sum_j n_j b_j$,

$$|n_j| \le c_m H^m, \qquad |\rho_i(\alpha) - 8m^2 H v_i| \le \tfrac{3}{4} m^2 H.$$

Thus $\operatorname{sign}(\rho_i(\alpha)) = -s_i$, and every $|\rho_i(\alpha)|$ lies between positive degree-dependent multiples of H. No unit-signature assertion or unbounded weak-approximation search is used. All the required certified approximations have polynomial bit cost. Define

$$M_{ij} = -\operatorname{Tr}_{K/\mathbb{Q}}(\alpha b_i b_j), \qquad h = |\det M|, \qquad S = \operatorname{diag}(s_i),$$
$$V = \operatorname{diag}(-s_i \rho_i(\alpha)), \qquad C = F^{-1} V^{-1} F^{-t}.$$

Then $M = F^t S V F$, $V > 0$, and

$$|M_{ij}| \le c_m H^3, \qquad 1 \le h \le c_m H^{3m}, \qquad \lambda_{\min}(C) = \lambda_{\max}(F^t V F)^{-1} \ge c'_m H^{-3}.$$

The entries of M^{-1}, C and $\lambda_{\max}(C)$ are bounded by $H^{O_m(1)}$.

Precompute the fixed geometric atlas described below, recording a positive integer J_m divisible by its finitely many excluded characteristics and cleared denominators. The explicit theta atlas permits just the exclusion of characteristic two; allowing J_m makes the fixed-data convention immaterial. If $p \le \max(2, J_m, h)$, enumerate all assignments of the m basis elements to $\mathbb{F}_p$, and test the multiplication and unit relations. Complete splitting guarantees a solution. Since $h \le c_m H^{3m}$, this costs $H^{O_m(1)}$ with a constant depending only on m. Henceforth p is odd, $p \nmid h$, and no geometric exception occurs. The splitting promise already implies $p \nmid D_K$.

A polynomially bounded imaginary quadratic field. Choose $Q_0 = \lceil c_m H^6 \rceil$ with c_m large enough that, for $q > Q_0$ and $\beta = \sqrt{q}/2$,

$$\sum_{n \in \mathbb{Z}^m \setminus \{0\}} e^{-16\pi\beta\lambda_{\min}(C)\|n\|^2} < \tfrac{1}{2}.$$

Let P_0 be the product of all primes at most Q_0, computed by a sieve, and put $Q = 4pD_K P_0$. Factoring D_K is unnecessary. On $(\mathbb{Z}/Q\mathbb{Z})^*$ the characters $(-1/\cdot)$ and $(p/\cdot)$ are independent: the second has a nontrivial p-part and the first detects the modulo-four factor. Their values $-1, +1$ specify a coset of index four. Theorem 1.4 of Lamzouri et al. (2015, corrected author p. 3), under the stated Dirichlet hypothesis, gives a prime q in this coset with

$$q \le \max\{10^9, c(1 + \log Q)^2\} = (H + D_K + 1 + \log p)^{O_m(1)}.$$

The fixed 10^9 alternative in that theorem is included; the finitely many small moduli only change the effective degree-dependent constant. The bound follows from $\log P_0 \le Q_0 \log Q_0$. Every prime in the coset is coprime to Q, and therefore

$$q > Q_0, \quad q \equiv 3 \pmod 4, \quad (p/q) = 1, \quad q \nmid pD_K.$$

Trial search to the displayed numerical bound, using primality, gcd and quadratic-character tests, has the required deterministic polynomial cost. This is the sole conditional step of the proof.

Set $L = \mathbb{Q}(\sqrt{-q})$, $E = KL$, $G = \mathrm{Gal}(K/\mathbb{Q})$, and $\mathcal{G} = G \times \langle c \rangle = \mathrm{Gal}(E/\mathbb{Q})$, where c is complex conjugation. Coprime discriminants and reciprocity give

$$\mathcal{O}_E = \mathcal{O}_K \mathcal{O}_L, \qquad D_E = D_K^2 q^m, \qquad (-q/p) = (p/q) = 1.$$

Thus p splits completely in E. The product integral basis and its multiplication table are explicit of polynomial encoding length. This step invokes no estimate for a nonrational L-function.

A primitive type with two trivial stabilizers. Identify the embeddings of E with $\mathcal{G}$ and choose

$$\Phi = \{(g, 1) : g \ne 1\} \cup \{(1, c)\}.$$

This type has one minus sign over the m real embeddings of K. Both its left and its right stabilizer in $\mathcal{G}$ are trivial. A translation without c must fix the unique location of the minus sign, hence is the identity. A translation with c exchanges the numbers 1 and $m-1$ of minus and plus signs, so cannot preserve the type when $m \ge 3$. An induced type from a proper CM subfield would be invariant under right multiplication by the nontrivial subgroup fixing that subfield. Thus Φ is primitive. The CM classification (Milne, 1999, Proposition 2.1, p. 54) implies that the associated abelian variety is simple with rational endomorphism algebra E. This classification does not assume that $\mathcal{G}$ is abelian.

Principalization with a controlled lattice index. Use the controlled α, F, S, M, h, V constructed above, with the embedding order matched to the displayed type. In particular,

$$M_{ij} = -\operatorname{Tr}_{K/\mathbb{Q}}(\alpha b_i b_j), \quad h = |\det M|, \quad V = \operatorname{diag}(-s_i \rho_i(\alpha)).$$

The matrix M is symmetric, integral, nonsingular, and polynomially bounded in H; V is positive and $M = F^t S V F$. Write $\omega = (1 + \sqrt{-q})/2$ and $\beta = \sqrt{q}/2$. The positive integral alternating form

$$\mathcal{E}(x, y) = \operatorname{Tr}_{E/\mathbb{Q}}\left(\frac{\alpha}{\sqrt{-q}} x \overline{y}\right)$$

on $\Phi(\mathcal{O}_E)$ has block matrix $\left(\begin{smallmatrix} 0 & M \\ -M^t & 0 \end{smallmatrix}\right)$. Enlarge its second half to the lattice

$$\Lambda' = F\mathbb{Z}^m + \left(\tfrac{1}{2} I + i\beta S\right) F M^{-1} \mathbb{Z}^m.$$

The index $[\Lambda' : \Phi(\mathcal{O}_E)]$ is h, and the form becomes unimodular on the displayed basis. Consequently $A' = \mathbb{C}^m/\Lambda'$ is principally polarized, and the natural isogeny from $A = \mathbb{C}^m/\Phi(\mathcal{O}_E)$ has degree $h \le c_m H^{3m}$. Since

$$\Phi(\mathcal{O}_E) \subseteq \Lambda' \subseteq h^{-1}\Phi(\mathcal{O}_E),$$

the endomorphism order contains $\mathbb{Z} + h\mathcal{O}_E$, and $\mathrm{End}^0(A') = E$. A normalized period matrix is

$$\tau = \tfrac{1}{2} M^{-1} + i\beta C, \qquad C = F^{-1} V^{-1} F^{-t} > 0.$$

Indeed $F^{-1} S F M^{-1} = C$. Its real part and positive imaginary direction have the numerical bounds proved above; only β depends on the auxiliary prime.

Polynomial degree and effective unconditional height. If a lattice Λ is stable under $\mathbb{Z} + h\mathcal{O}_E$ and $J = \mathcal{O}_E\Lambda$, then $hJ \subseteq \Lambda \subseteq J$. Up to scaling there are h_E choices for the maximal-order ideal class of J. Each intermediate lattice corresponds to a subgroup of $J/hJ \simeq (\mathbb{Z}/h\mathbb{Z})^{2m}$. Such a subgroup has at most $2m$ generators, so their number is at most h^{4m^2}. For a fixed simple CM torus, principal polarizations differ by totally positive units of its real endomorphism order modulo norms of endomorphism units. The norm subgroup contains squares of real units, giving at most 2^m classes. CM types, symmetric line bundles and level-sixteen markings contribute only degree-dependent factors. The degree of the marked theta point is consequently at most

$$\Delta_0 \leq c_m h^{4m^2} h_E.$$

Minkowski's bound and elementary ideal counting in degree $2m$ give

$$h_E \leq c_m \sqrt{D_E}(1 + \log D_E)^{2m-1}.$$

Indeed, the number of ideals of norm n is at most the $2m$-fold divisor function, and every class has an integral ideal of norm at most $c_m\sqrt{D_E}$. Thus $\Delta_0 = (H + D_K + q)^{O_m(1)}$. This is a degree bound; the algorithm does not compute a class group.

The unconditional height estimate of von Känel and Kret (2023, Section 7.2.1, equation (7.7), Lemma 7.10, pp. 62–63) applies to every CM field and type. Take their $\epsilon = 1/2$ and $c = D_E/D_K = D_K q^m$. The class-number ratio h_K/h_E is at most one, the Hasse unit index is at most two, and the number of roots of unity is bounded in degree $2m$. Their estimate therefore gives

$$h_F(A) \leq c_m(1 + D_K q^m).$$

Only this coarse degree-uniform bound is needed. In the elementary absolute-convergence estimate one may use the weaker upper bound $\zeta(s)^m$ for a degree-m Euler product; this merely changes c_m. No precise source constant or individual nonabelian Colmez conjecture is assumed. The isogeny inequality (von Känel and Kret, 2023, (7.11), p. 67), also recorded in (Pazuki, 2012, Remark 2.3, p. 27), gives $|h_F(A') - h_F(A)| \leq \frac{1}{2}\log h$. Pazuki's effective comparison (Pazuki, 2012, Corollary 1.3, pp. 21–22) then bounds the fixed-level logarithmic theta height by $(H + D_K + q)^{O_m(1)}$. All constants depend on the degree and level alone.

Exact theta coordinates and fixed geometric data. Use the fixed level $N = 16$. Mumford's theta moduli scheme of type $(N, \ldots, N)$ and its embedded universal abelian scheme are fixed finite-type objects over the ring obtained by inverting two and adjoining the fixed necessary roots of unity (Mumford, 1967, Section 6, pp. 76–84). Include a finite affine cover, projective scheme equations, and the identity section in the fixed data. Include a complete finite addition atlas as well. To see this last assertion, on each fixed affine base chart let $\mu : \mathcal{A} \times \mathcal{A} \to \mathcal{A}$ be addition. For a sufficiently large fixed t, the line bundle $\mathcal{O}(t, t) \otimes \mu^*\mathcal{O}(1)^{-1}$ is relatively generated by finitely many sections. Multiplying each by the pulled-back output coordinates gives a tuple representing μ wherever that section is nonzero. Increasing t if necessary makes these sections restrictions of ambient bihomogeneous polynomials. Their nonvanishing domains cover every geometric point. Fixed symbolic projective algebra computes such data; any finite characteristics lost on spreading out are fixed exceptions. The equations of each smooth fiber have radical affine dehomogenizations. This is the bounded presentation, including complete group-operation data, required by Pila (Pila, 1990, Section 3, pp. 752–753). The formulas in the proof of Proposition 11.2 provide an alternative explicit atlas in every fixed dimension $g = m$: the multiplication matrix, its quadratic kernel and right inverse, and the biquadratic columns are indexed by $(\mathbb{Z}/8\mathbb{Z})^m$ and $(\mathbb{Z}/16\mathbb{Z})^m$. Mumford's multiplication and graded-ring results used there hold in arbitrary dimension. The resulting 8^m columns cover every geometric pair because the addition matrix is the outer product of the coordinates of $x + y$ and $x - y$.

The quadratic equations define the saturated projective scheme, so their standard affine-chart ideals are radical. Every matrix size, equation degree, and operation count depends only on m.

For clarity, the analytic coordinate convention is matched exactly. For $a, b, j \in (\mathbb{Z}/4\mathbb{Z})^m$, put

$$P_{a,b} = \theta[\,a/4\,;\,b/4\,](\tau, 0), \qquad Q_k = \theta[\,k/16\,;\,0\,](16\tau, 0).$$

Splitting the theta-series index into classes modulo four gives

$$P_{a,b} = \zeta_{16}^{a\cdot b} \sum_j \zeta_4^{j\cdot b} Q_{4j+a}, \qquad Q_{4j+a} = 4^{-m} \sum_b \zeta_4^{-j\cdot b} \zeta_{16}^{-a\cdot b} P_{a,b}.$$

This fixed invertible Fourier transform relates the theta-null conventions in (Pazuki, 2012, Section 2.3.4, p. 30) and is integral away from two. It changes heights by a fixed additive constant.

The lower eigenvalue bound and the choice $q > Q_0$ give

$$|\theta[0;0](16\tau, 0) - 1| < \tfrac{1}{2}.$$

Thus this coordinate supplies a normalization bounded away from zero, uniformly in K. Certified evaluation to P bits uses Gaussian tails and a summation box of side polynomial in P, H, q; the number of terms is polynomial because m is fixed. The upper matrix-entry bounds control phase reduction and precision propagation as well. The degree bound Δ_0 and logarithmic-height bound give polynomial minimal-polynomial coefficient-bit bounds. Algebraic-number recognition (Kannan et al., 1988, Theorem 1.19, p. 241; Proposition 1.6, p. 238), followed by deterministic characteristic-zero factorization (Lenstra, 1983, Theorem 4.5, pp. 253–254), recovers the exact normalized theta coordinates with their selected complex embeddings. There are only 16^m coordinates, so even the product bound for their compositum degrees remains polynomial with exponent depending only on m. Adjoin E and ζ_{16} to the coefficient field B; then

$$[B : \mathbb{Q}] + \operatorname{length}(B, \text{theta point}, E \hookrightarrow B) \le (H + D_K + q)^{O_m(1)}.$$

No CM class polynomial or its finite-field root is an input.

Defined endomorphisms and extension of the specified marking. The marking makes $A'[16]$ rational over B. The finite Galois action on the lattice $\operatorname{End}(A'_{\overline{B}})$ is congruent to the identity modulo 16: an endomorphism vanishing on $A'[16]$ factors through $[16]$. A finite subgroup of a principal integral congruence group of level at least three is trivial. Thus every geometric endomorphism is defined over B, with no assumption about the commutativity of $\mathcal{G}$. CM gives potential good reduction (Serre and Tate, 1968, Theorem 6(a), p. 503). At an odd place, the resulting finite inertia image on the two-adic Tate module lies in the torsion-free group $1 + 4M_{2m}(\mathbb{Z}_2)$. It is trivial, so the good-reduction criterion gives an abelian scheme over the valuation ring.

We also need the *specified* theta-marked object to extend. Let R be this valuation ring, with fraction field B_v, and let $\mathcal{A}/R$ be the good abelian scheme. Regularity and the integral special fiber make $\operatorname{Pic}(\mathcal{A}) \to \operatorname{Pic}(\mathcal{A}_{B_v})$ an isomorphism: divisors extend by closure, while a vertical divisor is a multiple of the principal special fiber. Extend the marked line bundle and normalize its rigidification at zero by rescaling the generic identification. If its polarization is $N\lambda_0$ generically, the principal homomorphism λ_0 and its inverse extend, so the special polarization homomorphism has finite kernel. Semicontinuity gives a nonzero special-fiber section. An effective line bundle with finite translation kernel on an abelian variety is ample (Mumford, 1970, Chapter II, Section 6, Application 1, pp. 60–61). Thus the extended line bundle is relatively ample.

The full torsion marking extends since N is a unit in R. Compare a specified generic theta lift of a standard order-N torsion generator with an arbitrary integral lift. Their ratio is a scalar

$u \in B_v^*$. The integral lift's Nth iterate is a unit and the specified lift's iterate is one; therefore $Nv(u) = 0$, so u is a unit. The exact specified lift extends. Relations extend by equality on the generic fiber, giving the specified theta-group isomorphism. The symmetry involution extends by the same argument with exponent two; its restriction to the two-torsion, and compatibility with the theta structure, extend as well. These are precisely Mumford's normalized, totally symmetric, theta-marked moduli data (Mumford, 1967, Section 6, pp. 76–79, 83). The marked moduli point therefore extends into the moduli open. Its projective extension is the unique projective specialization, so the coordinate reduction used below cannot land on an unwanted boundary point.

Reduction without choosing a residue field. Compute a p-maximal order $\mathcal{O}$ in B by (Buchmann and Lenstra, 1994, Algorithm 6.1 and Proposition 6.2, pp. 246–247); this algorithm does not factor p into individual prime ideals. Let $\mathfrak{r} = \sqrt{p\mathcal{O}}$ and $R_0 = \mathcal{O}/\mathfrak{r}$. Frobenius kernels compute the radical, so R_0 is an explicit product of finite fields. Let I be the fractional ideal generated by the homogeneous theta coordinates. The module $I/\mathfrak{r}I$ is rank-one projective over R_0, since the localizations of $\mathcal{O}$ above p are DVRs. Choose a coordinate with nonzero image and quotient R_0 by its annihilator. On the nonempty support it generates the module; linear algebra computes every affine coordinate ratio. The same annihilator operation selects a nonempty chart and makes its denominators invertible. Validity of the marked specializations guarantees that such a chart exists.

For $i = 1, \ldots, \dim_{\mathbb{F}_p} R_0$, quotient by the ideal generated by $x^{p^i} - x$ for an $\mathbb{F}_p$-basis of R_0. A field factor survives exactly when its degree divides i. The first nonzero quotient consists of copies of a single field $\mathbb{F}_{p^i}$, and $i \le [B : \mathbb{Q}]$. No least common multiple of residue degrees is used. Run Pila's algorithm over that product. Field arithmetic and zero tests are simulated by multiplication matrices. If a test has both zero and nonzero support, retain either nonempty support and restart; each restart strictly decreases dimension. After at most $\dim R_0 - 1$ such decreases, the computation returns the Frobenius polynomial of at least one surviving fiber over $\mathbb{F}_{p^i}$. The complete fixed geometric atlas above supplies exactly the required input. The running time is $(H + D_K + q + \log p)^{O_m(1)}$.

Frobenius in a possibly nonabelian Galois CM field. Write π for Frobenius on a surviving fiber. All CM endomorphisms are defined over its residue field, hence commute with π. For a prime $\ell \neq p$, the Tate module has rank one over $E \otimes \mathbb{Q}_\ell$. The centralizer of that action on the Tate module is $E \otimes \mathbb{Q}_\ell$. Faithfulness implies that the rational centralizer of E inside the endomorphism algebra has dimension at most $2m$; it already contains E, so equals E. Thus $\pi \in E$. It is integral because its Frobenius polynomial is monic integral. This uses the commutativity of the field E, not of its Galois group.

The Shimura–Taniyama formula (Milne, 1999, Section 5, p. 64, before Proposition 5.3) applies with $E = \mathrm{End}^0(A')$ and does not require a maximal integral endomorphism order. As p splits completely in E, the valuations of π at its $2m$ primes over p are 0 or i. Choose a base prime P_1 and label primes by $P_x = x(P_1)$, $x \in \mathcal{G}$. In a compatible embedding convention the positive support is $T = \Phi^{-1}$; a different convention translates this set. Both its stabilizers are trivial: the left stabilizer of Φ^{-1} is the inverse of the right stabilizer of Φ, and conversely. If $\tau\pi = \pi$, then $\tau T = T$, hence $\tau = 1$. Therefore π generates E and its Frobenius polynomial is its degree-$2m$ minimal polynomial. Factor that polynomial over the explicit characteristic-zero field E and choose any root. Its conjugates have the same trivial-stabilizer properties.

Separating supports and reading the residue map. In the explicit split algebra $R_E = \mathcal{O}_E/p\mathcal{O}_E \simeq \mathbb{F}_p^{2m}$ compute

$$e_\tau = (\tau\pi)^{p-1} \pmod{p} \qquad (\tau \in \mathcal{G}).$$

The values of this idempotent are zero on τT and one elsewhere. These masks separate all prime components, including when $\mathcal{G}$ is nonabelian. Indeed, equality of all masks at P_x, P_y would imply

$$\tau^{-1}x \in T \iff \tau^{-1}y \in T \quad \text{for every } \tau \in \mathcal{G}.$$

Putting $u = \tau^{-1}x$ shows $u \in T \iff u(x^{-1}y) \in T$ for all u. The trivial *right* stabilizer of T then forces $x = y$. Enumerate the fixed number 2^{2m} of products choosing e_τ or $1 - e_\tau$ at each τ. Every nonzero such product is a primitive idempotent. Choose one, say e. For each integral basis element b_j of K, linear algebra computes the unique $t_j \in \mathbb{F}_p$ with $b_j e = t_j e$. The assignment $b_j \mapsto t_j$ is the desired ring map.

The CM dimension, geometric degrees, number of coordinates and mask count depend only on m. The varying lattice index, class-count degree bound, height, coefficient-field degree and encoding, and residue degree are all polynomially bounded in H, D_K, q. The auxiliary-prime bound makes the complete bit cost $(H + D_K + 1 + \log p)^{O_m(1)}$. Every field-dependent exception was covered by the explicitly bounded elementary branch at the start, while projective ideal reduction handled coordinate denominators. This proves the asserted quantitative statement. □

The fixed-field interface is the following immediate special case.

Proposition 16.2. *Fix a totally real Galois number field $K/\mathbb{Q}$ of degree $m \geq 3$, with an integral basis and its multiplication table. Assume the Riemann hypothesis for all Dirichlet L-functions over $\mathbb{Q}$. There is a deterministic algorithm which, on a prime p splitting completely in K, constructs a ring homomorphism*

$$\lambda : \mathcal{O}_K \longrightarrow \mathbb{F}_p$$

in time polynomial in $\log p$. The polynomial and all fixed precomputation may depend on K. The Galois group of $K/\mathbb{Q}$ may be nonabelian and may have even order. No hypothesis about Hecke, Dedekind, or nonabelian Artin L-functions is required.

Proof. Apply Proposition 16.1. With K and its integral basis fixed, m, H, D_K are constants, giving the claimed polynomial bound in $\log p$ and the same sole analytic hypothesis. □

Scope. Complete splitting remains essential. The quantitative proposition allows varying numerical field data in fixed degree; it does not assert a polynomial bound with an absolute exponent when the field degree grows.

16.1 Removing primes with bounded target valuation

The numerical cofactor can be reduced further when the part of $\ell - 1$ supported on the target prime has bounded valuation. A residue map from the corresponding small cyclotomic subfield selects a quotient of rank prime to the target; its full norm then descends the root.

Lemma 16.3 (A cyclotomic quotient of coprime rank)**.** *Fix an odd prime r and an integer $b \geq 1$. Let $p \neq r$ and $\ell \neq r$ be primes, with ℓ odd,*

$$d = \operatorname{ord}_r(p), \qquad \ell \mid \Phi_d(p), \qquad v_r(\ell - 1) = b.$$

Under rational Dirichlet GRH, a monic divisor $G_\ell(T)$ of $\Phi_\ell(T)$ over $\mathbb{F}_p$ can be constructed deterministically with degree

$$h_\ell = \frac{\ell - 1}{r^b}, \qquad \gcd(h_\ell, r) = 1,$$

in bit time polynomial in $\ell + \log p$, with constants depending only on (r, b). Every root of G_ℓ has exact order ℓ. For every nonconstant monic $g \mid \Phi_r$ and $A = \mathbb{F}_p[Z]/g$, the algebra

$$B_\ell = A[T]/G_\ell(T)$$

is finite free and étale over A of rank h_ℓ, and all its field components are copies of $\mathbb{F}_{p^d}$. No number-field residue map, cyclotomic root, finite-field factor, or extension presentation is an input.

Proof. Put $m = r^b$ and let H_ℓ be the subgroup of $(\mathbb{Z}/\ell\mathbb{Z})^*$ of order h_ℓ. Its fixed field

$$K_\ell = \mathbb{Q}(\zeta_\ell)^{H_\ell}$$

is cyclic Galois of degree m. It is totally real because m is odd, so $-1 \in H_\ell$. Factoring the numerical integer $\ell - 1$ by trial division and searching for a primitive residue class construct H_ℓ and its cosets in time polynomial in ℓ. Write these cosets as $C_0, \dots, C_{m-1}$ and put

$$\eta_j = \sum_{a \in C_j} \zeta_\ell^a.$$

The elements ζ_ℓ^a for $1 \le a \le \ell - 1$ form an integral basis of $\mathbb{Q}(\zeta_\ell)$, permuted by H_ℓ. Their orbit sums therefore form an integral basis of the invariant ring, which is $\mathcal{O}_{K_\ell}$. This proves maximality of the period order; no separate integral-basis or integer-factorization oracle is used.

Count pairs $a \in C_i$, $a' \in C_j$ according to $a + a' \pmod{\ell}$. For nonzero residues the count is constant on each coset; the zero-residue contribution is eliminated using $1 = -\sum_j \eta_j$. This computes the multiplication table with coefficients of absolute value at most ℓ in polynomial time. Its discriminant is computed by the trace pairing and is bounded by $\ell^{O_m(1)}$. The sharper value ℓ^{m-1} from the conductor–discriminant formula is not needed.

The order of p modulo ℓ divides $d < r$, whereas $\mathrm{Gal}(K_\ell/\mathbb{Q})$ is an r-group. Thus the Frobenius of p in this quotient is trivial, and p splits completely in K_ℓ. Proposition 16.1 constructs a residue map $\lambda_\ell : \mathcal{O}_{K_\ell} \to \mathbb{F}_p$ in time polynomial in $\ell + \log p$, because $m = r^b$ is fixed. Let $t_j = \lambda_\ell(\eta_j)$ and write $\eta_j(T) = \sum_{a \in C_j} T^a$. Compute

$$G_\ell(T) = \gcd\bigl(\Phi_\ell(T), \eta_0(T) - t_0, \dots, \eta_{m-1}(T) - t_{m-1}\bigr),$$

normalized to be monic. This is ordinary prime-field Euclidean gcd arithmetic, not finite-field factorization.

Since $p \neq \ell$, the prime-conductor cyclotomic algebra is étale at p. The period-value conditions select precisely the embeddings extending the chosen residue map of K_ℓ, of which there are h_ℓ. Equivalently, the reduction of the integral period basis distinguishes the m residue maps of the completely split subfield, and one full H_ℓ-orbit of primitive roots is selected. Hence $\deg G_\ell = h_\ell$. It is squarefree, and its roots have exact order ℓ. Finally, every component of A is $\mathbb{F}_{p^d}$, and $\ell \mid p^d - 1$ makes G_ℓ split over that field. The asserted finite free étale presentation and component description follow. All computations have the claimed bit bound. □

Proof of Theorem 1.2

Proof. The elementary valuation identity gives $a = v_r(N)$, so the displayed cofactor is a positive integer. Sieve all primes up to L, test their divisibility in N, and compute the relevant valuations by repeated exact division. This costs a polynomial in $L + r \log p$, and $\sum_{\ell \in \mathcal{T}_{r,t,L}(p)} \ell \le L^2$.

Follow the proof of Theorem 5.1, with this residual cofactor in its projected rational-count step. The only additional ingredient is the treatment of a prime-order transition at a supported ℓ with $b = v_r(\ell - 1) > 0$. The refinement in that proof produces a nonempty base factor A_{good} on which the order-ℓ element

$$\varepsilon \in A_{\mathrm{good}}[X]/(X^r - z)$$

is nonconstant in every rank-r component. The original root procedure restarts on the complementary base factor. For $b = 0$, the full formal algebra $A_{\mathrm{good}}[T]/\Phi_\ell(T)$ has rank $\ell - 1$ prime to r, exactly as in that theorem. For $1 \le b \le t$, use Lemma 16.3 to construct G_ℓ and use

$$B = A_{\mathrm{good}}[T]/G_\ell(T), \qquad h_\ell = (\ell - 1)/r^b,$$

instead. This has a strict primitive order-ℓ element $\theta = T$, and its rank over A_{good} is prime to r.

The norm calculation requires a root on every component of the full auxiliary algebra. For $j = 0, \dots, \ell - 1$, compare ε with θ^j and perform the Euclidean steps on $X^r - z$ and $\varepsilon - \theta^j$. A nonunit coefficient splits the finite étale coefficient algebra into its multiplication kernel and image; continue the same ε-based calculation on both direct factors. In every field component, nonconstancy of ε ensures that some comparison has a gcd of degree strictly between zero and r. For a monic proper factor f of degree j, the product of its roots is $b = (-1)^j f(0)$ and satisfies $b^r = z^j$. Since $1 \le j < r$, choose u, v with $uj + vr = 1$; then $(b^u z^v)^r = z$. Patch all the roots by the computed idempotents. There are at most $\dim_{\mathbb{F}_p} B$ leaves, and this produces a simultaneous $y \in B$ with $y^r = z$ before taking a norm. The rational search is not restarted inside B, and another prime is not adjoined inside it.

Compute the full determinant norm $n = N_{B/A_{\mathrm{good}}}(y)$ on the original free basis $1, T, \dots, T^{h_\ell - 1}$. It satisfies $n^r = z^{h_\ell}$. Choose $u, v \in \mathbb{Z}$ with $uh_\ell + vr = 1$; then

$$x = n^u z^v \in A_{\mathrm{good}}, \qquad x^r = z.$$

Negative powers are of units, and a division-free determinant gives the norm in polynomial time. The same formula with rank $\ell - 1$ applies when $b = 0$.

The remaining primary-coboundary argument, the lifting of Z through $a-1$ root extractions, the tame projector and the fixed-algebra descent are those of Theorem 5.1. They give the simultaneous root interface, the specified maximal tame unit U, and the degree-r irreducible polynomial. During each simultaneous-root call, each base branch has at most one prime-order auxiliary algebra open at a time; the dimensions of the distinct base branches sum to at most $r - 1$. Thus no product of the degrees of different supported primes occurs.

For each $1 \le b \le t$, the field degree r^b in Lemma 16.3 belongs to a fixed finite list. The total cost of the residue maps is consequently $(L + \log p)^{O_{r,t}(1)}$. The rational trial count, algebra ranks, valuation chains, norms and restarts have polynomial cost in $r, L, C_{r,t,L}(p), \log p$ with an absolute exponent before this fixed-degree CM cost is included. Summing over the at most L supported primes preserves an exponent depending only on (r, t), not on their number or on their valuations in N. This proves the bound. For $t = 0$ no CM residue map is called, so no analytic assumption is used. □

Cubic specialization and its limits. For $r = 3$, $t = 1$ and $p \equiv 2 \pmod 3$, the removable odd primes other than three are precisely those $\ell \le L$ dividing $p + 1$ with $\ell \not\equiv 1 \pmod 9$. Thus the theorem is polynomial, with an absolute exponent, in $L + \log p$ and

$$\frac{p+1}{2^{v_2(p+1)} 3^{v_3(p+1)} \prod_{\substack{\ell \text{ odd prime},\ \ell \neq 3,\ \ell \not\equiv 1 \ (9)}} \ell^{v_\ell(p+1)}}.$$

The CM dimension is always three; the growing auxiliary algebra rank is charged polynomially in L. For $p = 2$, take $X^3 + X + 1$ and $U = Z$. For $p = 3$, the irreducible-polynomial task has $X^3 - X - 1$. For $p \equiv 1 \pmod 3$, rational Dirichlet GRH already gives a cubic nonresidue and hence a binomial irreducible cubic independently of this cofactor bound. No infinitude or density assertion about primes with a small residual cofactor is made. Allowing t or r to grow does not preserve the claimed exponent.

The rank restriction explains the limit of this particular extension. If only a subfield of degree r^j, with $j < b = v_r(\ell - 1)$, has been selected, the remaining cyclotomic rank is $r^{b-j}(\ell - 1)/r^b$,

still divisible by r. For a root $y^r = z$ in that algebra its full norm satisfies $N(y)^r = z^{\mathrm{rank}}$. Products of such norms and powers of z still have r-th power z^e with $r \mid e$; the required Bézout exponent one cannot result. Selecting the entire r-primary quotient removes this obstruction, but makes the CM dimension r^b. This is a limitation of these norm and CM steps, not a lower bound for other constructions.

Remark 16.4 (One valuation layer can remove a power of p). Theorem 1.2 can improve the bound of Theorem 1.1 by a power of p. Fix an odd prime r and an integer $t \geq 1$. Choose a prime $\ell \equiv 1 + r^t \pmod{r^{t+1}}$, so that $v_r(\ell - 1) = t$, and put $L = \ell$. For $k \geq 1$ let $a = \max(1, \lfloor\sqrt{k}\rfloor)$ and let p_k be the least prime with

$$p \equiv 1 \pmod 4, \quad p \equiv r^a - 1 \pmod{r^{a+1}}, \quad p \equiv \ell^k - 1 \pmod{\ell^{k+1}}, \quad p \equiv 1 \pmod q$$

for every prime $3 \leq q \leq L$ with $q \neq r, \ell$. Then $\mathrm{ord}_r(p_k) = 2$ and $N = p_k + 1$, with $v_2(N) = 1$, $v_r(N) = a$ and $v_\ell(N) = k$; every other odd prime $q \leq L$ has $N \equiv 2 \pmod q$. Hence, unconditionally,

$$C_{r,t-1,L}(p_k) = \frac{p_k + 1}{2r^a}, \qquad C_{r,t,L}(p_k) = \frac{C_{r,t-1,L}(p_k)}{\ell^k}.$$

For $t = 1$ the first is the cofactor C_L of Theorem 1.1. Under A the least-prime bound (Lamzouri et al., 2015, Corollary 1.2, corrected author p. 3) gives $p_k \leq (\varphi(M_k) \log M_k)^2$ for the modulus M_k of the displayed class, and $\log M_k = k \log \ell + O(\sqrt{k})$. Therefore

$$\liminf_{k \to \infty} \frac{\log(C_{r,t-1,L}(p_k)/C_{r,t,L}(p_k))}{\log p_k} \geq \frac{1}{2}.$$

Along this sequence $C_{r,t-1,L}(p_k) = p_k^{1-o(1)}$ and $C_{r,t,L}(p_k) \leq p_k^{1/2+o(1)}$, while $v_r(p_k^2 - 1) = a \to \infty$, so the fixed-valuation results do not apply. Both cofactors remain exponential in $\log p_k$: the remark compares numerical bounds and gives no new class of inputs solved in polynomial time. For $r = 3$ and $L = \ell = 7$, so $t = 1$:

k	$v_3(p_k + 1)$	p_k	$C_{3,0,7}(p_k)$	$C_{3,1,7}(p_k)$	ratio
1	1	1,721	287	41	7
2	1	147,881	24,647	503	7^2
3	1	430,121	71,687	209	7^3
4	2	33,666,821	1,870,379	779	7^4
5	2	95,900,741	5,327,819	317	7^5
6	2	1,738,616,921	96,589,829	821	7^6

The same rows were computed for $(r, \ell) = (3, 19)$ and $(5, 11)$, $1 \leq k \leq 6$. GPT-6.2 Astra at reasoning effort ultra (OpenAI Codex) found this remark after the text of version 1 was frozen. Claude Opus 5.5 at effort high checked its key step and recomputed all eighteen rows with an independent program.

17 Classical consequences and their exact hypotheses

Table 1 distinguishes a zero-location hypothesis from the algorithmic conclusion to which it contributes. Three distinctions recur. A short guaranteed search interval is not automatically a polynomial-time search in the bit length of the modulus. A set of class-group generators is not a presentation of that group. And a randomized or interactive complexity bound remains randomized or interactive when its analytic hypothesis is proved. The elementary arguments below make the positive transfers explicit; the remaining paragraphs identify the precise analytic families that survive.

17.1 Quadratic characters suffice for Miller's cutoff

The quadratic-character formulation is already explicit in Schoof (2008, Theorem 2.2, pp. 105–106); it is not a new strengthening claimed here. Handle $n \le 16$ by trial division; for larger n the cutoff is less than n, so it does not test a prime against its own zero residue. The following proof gives a short transfer with an elementary perfect-power precheck. Suppose n is odd, composite, and not a perfect power, and write $n-1 = 2^s d$ with d odd. It has two distinct prime divisors. For every prime $p \mid n$, write $p-1 = 2^{t_p} u_p$ with u_p odd. If a unit a passes the strong probable-prime test, either $a^d = 1$ at every prime divisor of n, or there is one $j < s$ for which $a^{2^j d} = -1$ at every such prime. In the first case the reduction of a has odd order at each prime. In the second, its 2-primary order is exactly 2^{j+1} at every prime.

If the t_p are not all equal, choose $q \mid n$ with $t_q > \min_{p \mid n} t_p$. Every passing unit then has $\left(\frac{a}{q}\right) = 1$. If the t_p are all equal, choose two distinct prime divisors p, q. Their Legendre symbols agree for every passing unit, so every such unit lies in the kernel of $\chi(a) = \left(\frac{a}{p}\right)\left(\frac{a}{q}\right)$. Thus in either case all passing units lie in the kernel of a nonprincipal primitive quadratic character of conductor $f \le n$. The small-character-witness theorem supplies $a \le 2(\log f)^2 \le 2(\log n)^2$ with $\chi(a) \ne 1$. If a is not a unit modulo n, its gcd detects compositeness; otherwise it fails the strong test. This proves the cutoff after the precheck.

The distinction between $\chi(a) \ne 1$ and $\chi(a) = -1$ matters for the constant 2: a small nonunit is permitted by the former statement and is handled by the gcd test. The published formulation without a perfect-power precheck first excludes repeated prime factors using a smooth-number count, then applies quadratic characters (Schoof, 2008, pp. 105–106). All these operations take polynomial bit time. This conclusion concerns Miller's particular test; primality already has an unconditional deterministic polynomial-time algorithm, also described in Schoof (2008, Section 3).

17.2 Nonresidues, primitive roots, and order information

For a given prime $r \mid p-1$, the kernel of a character of order r modulo p is precisely the subgroup of rth powers. The Dirichlet small-witness bound therefore gives an rth nonresidue among the positive integers at most $2(\log p)^2$ for $p \ge 17$, using the explicit least-prime bound outside a proper subgroup. A candidate is certified by $a^{(p-1)/r} \ne 1$. The finite number of small primes for which the cutoff reaches p can be treated directly; for $p \ge 17$ it lies below p. This uses rational Dirichlet L-functions, uniformly in the variable r (Lamzouri et al., 2015, pp. 1–2).

Suppose the full factorization $p-1 = \prod_i r_i^{e_i}$ is given. Choose an r_ith nonresidue a_i and put

$$b_i = a_i^{(p-1)/r_i^{e_i}}.$$

The r_i-part of the order of a_i is $r_i^{e_i}$, so b_i has exact order $r_i^{e_i}$. Their product has order $p-1$ because these orders are pairwise coprime. This constructs a primitive root in polynomial time and does not factor $p-1$.

The least-root theorem gives a different kind of information. Shoup proves $g(p) \ll t^4(1+\log t)^4 \log^2 p$, where $t = \omega(p-1)$. His proof invokes a weighted sum of rational Dirichlet characters, not number-field Hecke characters (Shoup, 1992, Theorem 3, author p. 4; Proposition 9, p. 10; Section 6, p. 13). Since $t! \le p-1$, we have $t(1+\log t) = O(\log p)$ and hence $g(p) = O(\log^6 p)$. A short interval containing a generator is useful, but identifying a generator requires an order certificate or an adequate order-computation procedure.

When a factored multiple M of an element's order is given, order computation is elementary and unconditional. Start with $d = M$ and remove a prime factor ℓ whenever $a^{d/\ell} = 1$. The invariant $\operatorname{ord}(a) \mid d$ persists; at termination no prime factor can divide $d/\operatorname{ord}(a)$, so equality holds. In contrast, Miller's general theorem states a polynomial-time Turing equivalence

between integer factorization and rational periods, hence integer-modulus order, under the rational Dirichlet RH used in his appendix (Miller, 1976, Theorem 4, pp. 311–312; pp. 315–316). Discharging that hypothesis establishes the equivalence, not a polynomial-time algorithm for either problem.

17.3 Randomized primality tests and their separate derandomization

Solovay–Strassen is unconditional as a randomized algorithm (Solovay and Strassen, 1977). For completeness, let n be odd and composite, and $e = (n-1)/2$. The Euler–Jacobi liars in $(\mathbb{Z}/n\mathbb{Z})^*$ form the kernel of the homomorphism $a \mapsto a^e \left(\frac{a}{n}\right)^{-1}$. This subgroup is proper. If $p^v \| n$ with $v \geq 2$, CRT gives a unit that is $1 + p^{v-1}$ modulo p^v and 1 at the other prime-power factors. Its Jacobi symbol is 1, whereas its eth power is not 1 modulo p^v. If n is squarefree composite, choose a nonresidue at one prime divisor and 1 at every other prime divisor. Its Jacobi symbol is -1, but its eth power is 1 at another prime divisor. At most half the units, and therefore at most half the candidate bases, pass. Independent trials give error at most 2^{-k} after k repetitions. The Dirichlet small-witness theorem can separately make this a deterministic small-base test by detecting an element outside the proper liar subgroup. It is this latter statement, not the original randomized error estimate, that consumes a Riemann hypothesis.

17.4 Distribution theorems: progression primes, Artin, and races

The corrected explicit bound $P(a,q) \leq (\varphi(q)\log q)^2$ for $q > 3$ and $\gcd(a,q) = 1$ uses only rational Dirichlet characters and therefore follows from A (Lamzouri et al., 2015, Corollary 1.2, corrected author p. 3). Enumerating an interval of length comparable to q^2 is nevertheless exponential in the binary length of q; this explicit theorem is not by itself a deterministic polynomial-time prime-search algorithm.

Hooley's proof of the positive-density version of Artin's primitive-root conjecture uses a different family. For a fixed integer $a \neq -1$ that is not a square, it counts complete splitting in the Kummer fields $\mathbb{Q}(\zeta_k, a^{1/k})$, for squarefree k, under RH for their Dedekind zeta functions (Hooley, 1967, translation pp. 40–45, 47). These fields are generally nonabelian over $\mathbb{Q}$. Factoring out the abelian pieces covered by A does not remove the remaining pieces from this proof. This identifies the surviving hypothesis of that proof; it is not a claim of logical independence from A, nor does it conflate an infinitude theorem with Hooley's full asymptotic.

Rubinstein–Sarnak's first limiting-distribution theorem does follow from Dirichlet RH (Rubinstein and Sarnak, 1994, Theorem 1.1, p. 174). The normalized prime-count errors have a limiting measure for bounded continuous test functions. A strict race is an indicator of an ordering region, so possible mass on its boundary must still be controlled. Their grand simplicity hypothesis is a separate rational-independence condition on the ordinates of Dirichlet zeros. It enters the explicit characteristic function and the familiar strictly ordered race densities (Rubinstein and Sarnak, 1994, pp. 176–177). Proving the zeros' real parts are 1/2 does not prove that independence.

The separate 2026 infinitude result. Theorem 1.1 of OpenAI (2026b, p. 2) states that, for each integer a that is neither -1 nor a square, there are constants $c_a > 0$ and x_a such that every $x \geq x_a$ satisfies

$$\#\{\ell \text{ prime} : x < \ell < 2x,\ \ell \nmid a,\ \operatorname{ord}_\ell(a) = \ell - 1\} \geq c_a \frac{x}{(\log x)^2}.$$

This is an infinitude theorem with a quantitative dyadic-interval lower bound. Hooley's conclusion is the stronger full asymptotic $\pi_a(x) \sim \delta_a \operatorname{Li}(x)$, with the predicted positive Artin density δ_a (Hooley, 1967, translation p. 47). The displayed $x/(\log x)^2$ lower bound does not itself imply that asymptotic or a positive density among the primes. The 2026 manuscript also states

the separate zero-free strip $\Re s > 1 - 10^{-6}$ for finite-order Hecke L-functions over cyclotomic fields containing μ_{12}, uniformly in conductor and height (OpenAI, 2026b, Theorem 1.2, p. 3). Its analytic input is therefore distinct from estimates for rational Dirichlet L-functions alone.

17.5 Class groups and elliptic algorithms

The bound $12(\log |D_K|)^2$ for prime-ideal generators of $\mathrm{Cl}(K)$ is often cited under the label GRH. Its analytic input includes ideal-class Hecke characters over K and ζ_K. The primary refinement by Belabas et al. (2008, author pp. 2–3) identifies these two zero sums explicitly; their own field-specific bound removes the need to assume RH for the trivial-character zero sum, but still requires the nontrivial ideal-class characters. The factorization of the Hilbert class field's Dedekind zeta function into these Hecke functions gives a convenient sufficient single-field package. Even when $K/\mathbb{Q}$ is quadratic, its Hilbert class field need not be abelian over $\mathbb{Q}$. Nor does a small generating set, by itself, bound the time needed to find relations or units.

The random-walk hitting estimate of Jao et al. (2009, Corollary 1.3, author p. 4; proof pp. 9–10) fixes the field degree and an exponent $B > 2$, and applies once its discriminant–modulus parameter is sufficiently large. Its proof displays the Hecke functions of the relevant class characters in equations (2.4)–(2.6), author p. 9. This is the random-walk statement classified in the table. Biasse and Erukulangara (2023, Proposition 1, author pp. 8–9) invoke that corollary for vectors chosen uniformly from an integer ℓ_1-sphere. That distribution differs from the exponent vector of independent random-walk steps: two inverse steps can cancel, and different vectors have different numbers of realizing step sequences. The cited walk theorem alone therefore does not establish the uniform-vector formulation. No assertion about that formulation or the full Hafner–McCurley runtime is needed here.

The same spectral theorem governs the ordinary elliptic isogeny application over imaginary quadratic orders (Jao et al., 2009, Theorems 1.1, 1.5–1.6). For base field $K = \mathbb{Q}$ its ray-class characters are Dirichlet characters and A suffices; the imaginary-quadratic class characters generally require more. The elliptic application fixes the finite field, the point count (hence the isogeny class), and the endomorphism level within that class. Its conclusion is expansion and rapid mixing; it does not provide a polynomial-time discrete-logarithm algorithm. Similarly, the split-prime selection estimate in the CM/CRT computation of Hilbert class polynomials uses RH for the ring class field's Dedekind zeta function (Sutherland, 2011, equation (21) and Lemma 2, corrected author pp. 21–22). This is a statement about that stage of the computation.

Schoof point counting occupies the unconditional side of the table. The original algorithm has deterministic $O(\log^9 q)$ complexity in characteristic different from $2, 3$ (Schoof, 1985, pp. 483, 485–490); Hasse's estimate is a proved finite-field theorem. ECPP certificate correctness is also unconditional. The practical Atkin-style algorithm described by Schoof has an explicitly heuristic runtime analysis involving prime and smoothness distributions (Schoof, 2008, Theorem 5.2, p. 121; pp. 122, 124–125). Supplying A does not silently supply those additional assumptions.

17.6 What the separate Dedekind strip gives for class generators

A fixed zero-free half-plane already gives a polynomial-logarithmic ray-class generator bound, although it does not give Bach's sharper quadratic logarithmic bound. The weighted discriminant argument is the same one that appears in the auxiliary-prime construction.

Proposition 17.1. *Fix $1/2 \le \theta < 1$. Suppose that for every number field J, ζ_J has no nontrivial zero with real part greater than θ. Let K be a number field of degree n and absolute discriminant D, and let $\mathfrak{m}$ be a modulus with finite part $\mathfrak{m}_0$. Put*

$$B = 1 + n + \log D + \log N\mathfrak{m}_0.$$

There is a constant C_θ, independent of K and $\mathfrak{m}$, such that the ray class group modulo $\mathfrak{m}$ is generated by prime ideals coprime to $\mathfrak{m}_0$ of norm at most

$$C_\theta B^{1/(1-\theta)}.$$

In particular, accepting CT*, the 2026 preprint theorem with Tang's revised Lemma 8.1, gives the exponent* 32*. Only that preprint's Dedekind-zeta conclusion is used.*

Proof. Fix once and for all a nonnegative smooth function w, supported in $(1, 2)$, with integral $c_w > 0$, and use the smoothed prime-ideal sum $\Psi_J(x)$ defined in Lemma 18.3. The stated zero-free half-plane and the explicit formula give

$$\Psi_J(x) = c_w x + O_{w,\theta}(x^\theta(\log D_J + [J : \mathbb{Q}])).$$

Here the constant is independent of J; the rapid decay of the Mellin transform makes the zero-counting sum bounded by a constant times $\log D_J + [J : \mathbb{Q}]$.

Let H be the subgroup generated by the prime ideals coprime to $\mathfrak{m}_0$ of norm at most $2x$. Suppose that H is proper. A maximal subgroup containing it has prime index ℓ in the ray class group. Class field theory gives a cyclic extension L/K of degree ℓ for that quotient, in which every prime counted in H splits completely. Every character of this quotient has conductor dividing $\mathfrak{m}$. The conductor–discriminant formula therefore gives

$$\log D_L \le \ell \log D + (\ell - 1) \log N\mathfrak{m}_0, \qquad [L : \mathbb{Q}] = \ell n.$$

This formula follows equally by factoring the completed zeta function of L into the completed Hecke functions of the characters and comparing their finite conductor factors.

Remove the Euler factors over prime ideals dividing $\mathfrak{m}_0$ from both smoothed sums, and denote the results by $\Psi_K^{(\mathfrak{m})}$ and $\Psi_L^{(\mathfrak{m})}$. Their weighted difference is

$$\Psi_K^{(\mathfrak{m})}(x) - \frac{1}{\ell}\Psi_L^{(\mathfrak{m})}(x) = \left(1 - \frac{1}{\ell}\right) c_w x + O_{w,\theta}(Bx^\theta).$$

To justify removing the Euler factors uniformly, for a fixed prime ideal $\mathfrak{p}$ of K the interval $(x, 2x)$ contains at most one power of $N\mathfrak{p}$. The sum of the residue degrees of primes above $\mathfrak{p}$ is at most ℓ. After the weight $1/\ell$, the removed contribution is consequently $O_w(\log N\mathfrak{p})$ on either side. Summing gives $O_w(\log N\mathfrak{m}_0)$, which is absorbed in the displayed error.

Every unremoved prime ideal that can contribute has norm at most $2x$, so its class belongs to H and it splits completely in L. Its entire prime-power contribution to the two weighted sums cancels. The left side is therefore zero. On the other hand, $1 - 1/\ell \ge 1/2$, and choosing $x = T_\theta B^{1/(1-\theta)}$ with a sufficiently large fixed T_θ makes the right side positive. This is a contradiction. Thus H is the whole ray class group, and $C_\theta = 2T_\theta$ proves the assertion. □

For the ordinary ideal class group take $\mathfrak{m} = 1$. If an integral-basis presentation of O_K is given, the resulting prime ideals can be enumerated by considering rational primes up to the stated norm bound and factoring the finite algebras O_K/pO_K. Unconditional algorithms polynomial in the numerical prime p suffice here. This constructs a generating list; it makes no assertion about the time needed to find all relations or compute the class-group structure.

17.7 Koiran's AM inclusion needs one solution-field zeta function

Koiran's extended report states the inclusion of complex polynomial-system feasibility in AM, with its probabilistic and interactive meaning explicit (Koiran, 1996b, Theorem 2, pp. 5–6). The journal's prime-density step explicitly notes that only the effective prime ideal theorem is used (Koiran, 1996a, Theorem 9, p. 282). To see the precise field, choose one algebraic solution, an integral primitive element θ for its coordinate field K, and the monic minimal polynomial f of θ. Exclude primes dividing its discriminant and the denominators of the coordinate formulas. For the remaining primes, the number $W(p)$ of roots of f modulo p is the number of degree-one prime ideals above p. Higher-degree prime ideals contribute only $O([K:\mathbb{Q}]\sqrt{x})$ to the prime-ideal count below x. RH for ζ_K therefore estimates $\sum_{p\le x} W(p)$, and

$$W(p) \le [K:\mathbb{Q}]\,\mathbf{1}_{\{f \text{ has a root modulo } p\}}$$

gives the necessary lower bound for primes at which the original system has a solution.

No normal closure is required for this argument. For example, every solution of $X^3-2=0$ has a nonnormal degree-three coordinate field. No such root lies in an abelian extension of $\mathbb{Q}$: every subfield of an abelian Galois extension is normal over $\mathbb{Q}$, whereas $\mathbb{Q}(\sqrt[3]{2})$ and its conjugates are not. Thus A does not in general discharge the solution-field zeta hypothesis. Choosing an integral primitive element also avoids silently omitting leading-coefficient bad primes in a nonmonic formulation. The corrected quantitative bounds in the author's erratum must be used in the subsequent bit-size analysis (Koiran, 1997). None of these refinements changes the classification of the conclusion as an AM inclusion rather than a deterministic polynomial-time decision algorithm.

17.8 A fixed uniform zero-free half-plane suffices for the AM protocol

The AM inclusion also follows from the separate uniform Dedekind zero-free half-plane. The following proof spells out the transfer, including the polynomial bit length of the prime range.

Proposition 17.2. *Suppose that a fixed $\theta < 1$ has the following property: for every number field K, the Dedekind zeta function ζ_K has no nontrivial zero with real part greater than θ. Assume the standard uniform functional equations and local zero-counting estimates for these zeta functions. Then feasibility over $\mathbb{C}$ of a finite system of integer polynomials belongs to* AM*, for the usual dense or sparse binary encoding. In particular this conclusion follows upon accepting* CT*, with the preprint and revised-lemma provenance specified in Section 1.*

Proof. We identify explicitly the effective algebraic input used below. For an input of bit length N, the standard effective algebraic-point and arithmetic Nullstellensatz bounds supply a computable polynomial $Q(N)$ with these two properties. Write $M = 2^{Q(N)}$. If the system is feasible, it has a point with coordinates in a number field K and with denominators dividing an integer $B \neq 0$, such that

$$n_K \le M, \qquad \log|D_K| \le M, \qquad \log|B| \le M.$$

More precisely, the coordinates belong to $O_K[1/B]$. If the system is infeasible, there is a nonzero integer A and integer polynomials G_i with

$$\sum_i G_i F_i = A, \qquad \log|A| \le M.$$

These effective algebraic bounds are the inputs in Koiran (1996b, Theorems 4, 5, and 7, author pp. 8–10), with the corrected quantitative bounds of Koiran (1997). For a primitive element θ with minimal polynomial of degree m and leading coefficient c, the element $\alpha = c\theta$ is integral: $c^{m-1}f(X/c)$ is monic integral. Its polynomial discriminant is a multiple of D_K. Substitution

$\theta = \alpha/c$ in the coordinate polynomials clears the additional denominators by a power of c. The degree and logarithmic heights of these quantities fit the same bound after enlarging the polynomial Q. The polynomial Q may be increased by a fixed constant below. For sparse binary exponents the total degree can be exponential in N; the permitted bounds are still exponential in a polynomial in N.

Choose once and for all

$$\max(\theta, 1/2) < \sigma < 1.$$

We first prove the required uniform prime-density estimate. For $x \geq 2$, define

$$W_K(x) = \sum_{N\mathfrak{a} \leq x} \Lambda_K(\mathfrak{a})(1 - N\mathfrak{a}/x).$$

Mellin inversion and a contour shift give

$$W_K(x) = \frac{x}{2} + O_{\theta,\sigma}\big(x^\sigma(\log|D_K| + n_K)\big). \tag{20}$$

Here the implicit constant is independent of K. For completeness, start with

$$\frac{1}{2\pi i}\int_{(2)} -\frac{\zeta_K'}{\zeta_K}(s)\frac{x^s}{s(s+1)}\,ds.$$

On the line $\Re s = \sigma$, the standard Hadamard formula and local zero count imply

$$\left|\frac{\zeta_K'}{\zeta_K}(\sigma + it)\right| \ll_{\theta,\sigma} \log|D_K| + n_K \log(|t| + 3).$$

One way to see the uniformity is to subtract the Hadamard formula at $2 + it$. Zeros with $|\Im\rho - t| \leq 1$ are separated from the new line by $\sigma - \theta$ and number $O(\log|D_K| + n_K \log(|t| + 3))$. The difference of the reciprocal terms for the remaining zeros is summable with quadratic decay in $|\Im\rho - t|$. The gamma terms satisfy the same bound, the pole at 1 is a fixed distance from the new line, and $\zeta_K'/\zeta_K(2 + it) = O(n_K)$ by its absolutely convergent Dirichlet series. The horizontal integrals vanish because the kernel is $O(|t|^{-2})$. Only the pole at $s = 1$ is crossed, with residue $x/2$. The integral on the new line is bounded by the error in (20), because the displayed logarithmic factor is integrable against $|(\sigma + it)(\sigma + 1 + it)|^{-1}$.

For each rational prime power p^j, the coefficient of p^{-js} in $-\zeta_K'/\zeta_K(s)$ is at most $n_K \log p$. Therefore all terms with $j \geq 2$ contribute at most

$$n_K \sum_{p \leq \sqrt{x}} \log x \leq n_K \sqrt{x} \log x.$$

The remaining terms come from prime ideals of norm a rational prime. A given rational prime contributes at most $n_K \log x$. Thus, when the error in (20) is at most $x/4$ and $n_K\sqrt{x}\log x \leq x/8$, there are at least

$$\frac{x}{8n_K \log x}$$

rational primes at most x having a residue-degree-one prime in K.

Choose a fixed integer

$$C \geq \max\{8, 4/(1 - \sigma)\}, \qquad x = M^C = 2^{CQ(N)}.$$

After increasing the constant term of Q, both preceding error inequalities hold for every input. If the system is feasible, a degree-one prime of K over a rational prime $p \nmid B$ reduces the given algebraic point to a solution in $\mathbb{F}_p$. There are at most $\log|B|/\log 2 \leq 2M$ excluded primes. Consequently the set

$$S = \{p \leq x : p \text{ prime and the input system has a solution in } \mathbb{F}_p\}$$

has size at least $256M$, again by a fixed increase of the constant term of Q. If the system is infeasible, the Nullstellensatz identity implies $p \mid A$ for every $p \in S$, so $|S| \le 2M$. No normal closure of K was taken: this is why the degree and discriminant parameters stay within the stated exponential bounds.

It remains to realize the cardinality gap by an AM protocol. Let $m = CQ(N) + 1$, so all integers up to x have fixed-length m-bit encodings, and let $k = Q(N) + 4$, so $R = 2^k = 16M$. Arthur samples a uniformly random binary matrix $U \in \mathbb{F}_2^{k \times m}$ and vector $b \in \mathbb{F}_2^k$, and sends them to Merlin. The affine hash is $h(z) = Uz + b$. Merlin replies with a prime $p \le x$ and a solution of the input system in $\mathbb{F}_p$, satisfying $h(p) = 0$. Arthur verifies the hash, primality, and all equations. Sparse powers are evaluated by binary exponentiation, so the verification is polynomial in N. The hash is applied only to the prime, not to its possibly many solution witnesses.

For an infeasible instance, a union bound gives

$$\Pr[\exists p \in S : h(p) = 0] \le |S|/R \le 1/8.$$

For a feasible instance, let Z count the primes of S in the zero bucket. Pairwise independence gives $\mathbb{E}Z = |S|/R \ge 16$ and $\operatorname{Var} Z \le \mathbb{E}Z$. Chebyshev's inequality therefore gives

$$\Pr[Z = 0] \le 1/\mathbb{E}Z \le 1/16.$$

On the complementary event Merlin can provide an accepted witness. Both messages and Arthur's computation have polynomial length and time, since C is fixed and Q is polynomial. This is a two-message public-coin AM protocol with completeness at least $15/16$ and soundness at most $1/8$. □

Consequently, accepting CT supplies the hypothesis of Proposition 17.2. The AM conclusion is available through that separate input, while the original solution-field RH hypothesis is not discharged by the rational Dirichlet transfer alone. The proof above uses one solution-field zeta function for each feasible instance and no normal closure.

17.9 Prime-field nonresidues and extension-field hypotheses

The distinction between a rational Dirichlet character and a ray-class character over another number field also matters in statements about finite-field nonresidues. The weak-GRH estimate of Bhargava et al. (2017, Theorem 6.7, p. 8) concerns $\mathbb{F}_p^\times$: its character is explicitly $x \mapsto x^{(p-1)/r}$. Theorem 1.8 proves the extension transfer when $r \mid p - 1$. Neither statement by itself covers $r \nmid p - 1$.

The wider generating-set theorem of Buchmann and Shoup (1996, author pp. 2–3) has complexity $(n \log p)^{O(n)}$ and hence polynomial complexity for fixed n. Its definition of ERH explicitly covers Dedekind zeta functions of all number fields (Buchmann and Shoup, 1996, author p. 5). There are two relevant uses: Fact 1 on that page invokes the Adleman–Lenstra auxiliary prime to construct a cyclic degree-n number field in which p is inert; Fact 6 on author p. 9 invokes Bach's ray-class generator bound over that field, with modulus p. Thus this theorem does not discharge the remaining extension-field nonresidue step merely by assuming A.

One can state the algorithmic input without any convention about the letters GRH: require deterministic polynomial-time construction of $\mathbb{F}_{p^d}$ and deterministic construction of a kth nonresidue in a presented such field, for a supplied prime $k \mid p^d - 1$, in time polynomial in k and $d \log p$. This is exactly Hypothesis $(*)$ in Guo (2017, §1.1, p. 3). Its two clauses must both be justified in an application. In particular, the prime-field result in Theorem 6.7 just cited does not alone prove that full hypothesis. The finite-family formulation in Theorem 1.10 records the additional input needed by the particular factorization route analyzed here.

Remark 17.3. The explicit extension-field calls also occur in Guo (2017, Caltech dissertation, Appendix B, Algorithm 17, p. 241, lines 9 and 12). We use those calls to identify the inputs, without relying on the literal pseudocode as a complete implementation. In particular, the loop described in Lemma B.1, p. 240, requires additional digit state. To see this, take $A = \mathbb{F}_{19}$, $\ell = 3$, $x = 4$ and $y = 9 = x^4$. Both units have order nine. Trying x^k, $0 \le k < 3$, and then replacing the current x by x^3 gives the successive test sets

$$\{1, 4, 16\}, \qquad \{1, 7, 11\}, \qquad \{1\}.$$

The base thereafter stays at one. None of the differences from 9 is zero or a zero divisor in this field. Thus this literal iteration misses a promised solution; a prime-power discrete-logarithm routine must accumulate digits and update its residual target. Algorithm 17 also omits the increment of n that its following prose describes. These local omissions concern the printed routines, not the existence of a correct digit algorithm or the thesis's abstract factoring theorem. For the promised prime-power comparison there is a direct repair. Suppose x has order ℓ^a, $a > 0$, and y has order dividing ℓ^a in a commutative finite semisimple algebra. Retain the components on which $x^{\ell^{a-1}} \neq 1$. They form a nonempty direct factor and x has full order in each. There y is a power of x since a finite field has cyclic multiplicative group. Starting at $k_0 = 0$, find a nonempty component support where

$$(yx^{-(k_i + j\ell^i)})^{\ell^{a-i-1}} = 1, \qquad 0 \le j < \ell.$$

The ℓ tests partition the current support by the next digit of the unique exponent there. Set $k_{i+1} = k_i + j\ell^i$ on one nonempty support. Multiplication kernels and images compute each direct factor by linear algebra. After a rounds $x^{k_a} - y$ vanishes on at least one original component. It is a nonzero zero divisor if y is not globally a power. The procedure uses at most $a\ell$ digit tests; $a = 0$ gives $x = y = 1$.

18 The zeta function of a pure field

This section proves the analytic equivalence and then identifies the finite computation which an auxiliary prime makes possible.

Lemma 18.1 (The affine permutation identity)**.** *For F, K, E in Theorem 1.6,*

$$\zeta_E(z) = \zeta_K(z) \left(\frac{\zeta_F(z)}{\zeta_{\mathbb{Q}}(z)} \right)^{r-1} \tag{21}$$

as meromorphic functions, including their ramified Euler factors. Moreover, for a nontrivial character χ of $\mathrm{Gal}(E/K)$*,*

$$L_K(z, \chi) = \frac{\zeta_F(z)}{\zeta_{\mathbb{Q}}(z)}. \tag{22}$$

Here L_K is the finite-order Hecke L-function identified with the Artin L-function by abelian class field theory over K.

Proof. Only r ramifies in $K/\mathbb{Q}$, so p is unramified there. At every prime above p, the polynomial $X^r - p$ is Eisenstein. Thus $[E : K] = r$, and the splitting group over $\mathbb{Q}$ is $G = \mathbb{F}_r \rtimes \mathbb{F}_r^{\times}$, acting by $t \mapsto bt + c$ on the r roots. Let $N = \mathbb{F}_r$ be its translation subgroup, let P be the permutation character of that action, and let Q_0 be the regular character of G/N, inflated to G. Then

$$\mathrm{reg}_G = (r - 1)(P - 1) + Q_0.$$

At the identity both sides equal $r(r-1)$. At a nonzero translation the right side is $(r-1)(-1) + (r-1) = 0$. At an element with $b \neq 1$, the affine action has exactly one fixed point and Q_0

vanishes. These cases prove the identity at every element. Artin formalism applied to this identity proves (21). It applies at a ramified prime as well: an equality of representations remains an equality after taking inertia invariants, and therefore gives the same local determinants.

The induced character $\operatorname{Ind}_N^G \chi$ is $r-1$ at the identity, -1 at a nonzero translation, and zero off N. Indeed conjugation runs through all nontrivial characters of N, whose sum is -1 on a nonidentity element. Consequently $\operatorname{Ind}_N^G \chi = P - 1$. Induction of Artin L-functions and class field theory give (22). No general Artin-holomorphy conjecture is used. □

Proof of Theorem 1.6. Under A, both $\zeta_{\mathbb{Q}}$ and ζ_K are nonzero at every point of the open critical strip off the line $\Re z = 1/2$. For ζ_K this follows from its factorization into Dirichlet L-functions, since $K/\mathbb{Q}$ is abelian. Equation (21) therefore says that ζ_E vanishes at such a point if and only if ζ_F does. This proves the RH equivalence and also identifies its extra hypothesis as RH for the single function in (22).

The cyclotomic discriminant is $|D_K| = r^{r-2}$. The relative discriminant of the order $\mathcal{O}_K[p^{1/r}]$ is generated, up to a unit, by $r^r p^{r-1}$. The relative field discriminant divides this ideal. The discriminant tower formula consequently gives

$$|D_E| \le r^{r(2r-3)} p^{(r-1)^2}, \qquad [E:\mathbb{Q}] = r(r-1).$$

Apply Bach–Sorenson's theorem to the conjugacy class of a nonidentity translation. Under the zeta hypotheses just established it gives an unramified rational prime of size at most $(4\log|D_E| + \frac{5}{2}[E:\mathbb{Q}] + 5)^2$ with that Frobenius class. Substitution proves (2).

Its restriction to K is trivial, so $\ell \equiv 1 \pmod r$. At a chosen prime $\mathfrak{L}$ of E above ℓ, write the action of its Frobenius on $p^{1/r}$ as $p^{1/r} \mapsto \zeta_r^c p^{1/r}$, where $c \neq 0$. The defining Frobenius congruence gives

$$p^{(\ell-1)/r} \equiv \zeta_r^c \not\equiv 1 \pmod{\mathfrak{L}}.$$

The left side belongs to $\mathbb{F}_\ell$, so it is different from one modulo the rational prime ℓ. Both p and r ramify, so neither is ℓ. The remainder follows from Lemma 18.2, proved next. Enumerating positive integers up to (2), testing primality, and testing the two displayed conditions has polynomial bit cost. □

Lemma 18.2 (A Gaussian-period construction)**.** *Let $p \neq \ell$ be primes, let the prime r divide $\ell - 1$, and suppose $p^{(\ell-1)/r} \not\equiv 1 \pmod \ell$. A presentation of $\mathbb{F}_{p^r}$ can be constructed deterministically in time polynomial in ℓ and $\log p$.*

Proof. Put $G = (\mathbb{Z}/\ell\mathbb{Z})^\times$ and let H be its subgroup of order $(\ell-1)/r$. We can enumerate G, its multiplication table, and H in polynomial time in ℓ; factoring integers at most ℓ by trial division also has this cost. Form the separable algebra $B = \mathbb{F}_p[Z]/\Phi_\ell(Z)$ and let G act by $Z \mapsto Z^b$. Compute B^H by linear algebra. Over an algebraic closure, the points of B are indexed by G, and those of B^H by G/H. Frobenius acts on these cosets by multiplication by p. The displayed nonresidue condition is exactly $p \notin H$. Since G/H has prime order r, Frobenius is a single r-cycle. A finite reduced algebra over a finite field is a field precisely when Frobenius is transitive on its geometric points. Hence $B^H \cong \mathbb{F}_{p^r}$. All matrices have dimension at most $\ell - 1$, and modular arithmetic has polynomial bit cost.

There is also an explicit generator. The elements $Z, Z^2, \dots, Z^{\ell-1}$ form a basis of B: a relation would be a scalar multiple of Φ_ℓ with zero constant coefficient, hence would vanish. The sums $\theta_{bH} = \sum_{h\in H} Z^{bh}$ have disjoint supports in this basis, so they are linearly independent and distinct. They form a basis of B^H. Frobenius cyclically permutes them; therefore θ_H has degree r. Its minimal polynomial, found by linear dependence among its powers, gives a polynomial presentation of the field. □

18.1 The sharper classical bound and a one-sided criterion

The least-nonsplit-prime theorem gives a better dependence on r than the explicit general conjugacy-class bound. Murty's result, in the form recorded by Zaman (2018, p. 1, equation (1.2)), gives $P(E/K) \ll (r^{-1}\log|D_E|)^2$ under RH for ζ_E. Here $P(E/K)$ is the least norm of a degree-one prime of K which is unramified and does not split completely in E. It gives the scale $r^2\log^2(rp)$ in our situation. We prove the needed specialization directly, including the exclusion of the two rational primes which ramify, and then state its weaker one-sided sufficient condition.

Fix once and for all a nonzero nonnegative smooth function w with compact support in $(1,2)$, and write

$$c_w = \int_1^2 w(t)\,dt, \qquad \Psi_J(x) = \sum_{\mathfrak{l}}\sum_{j\ge 1} \log N\mathfrak{l}\; w((N\mathfrak{l})^j/x)$$

for a number field J. Its absolute degree and discriminant are denoted by n_J and D_J.

Lemma 18.3 (Uniform smoothed estimate)**.** *Let $1/2 \le \theta < 1$ be fixed. If ζ_J has no nontrivial zero with real part greater than θ, then, uniformly for number fields J and real $x \ge 2$,*

$$\Psi_J(x) = c_w x + O_w\big(x^{\theta}(\log|D_J| + n_J)\big).$$

Proof. Put $\widehat{w}(z) = \int_1^2 w(t)t^{z-1}\,dt$. The smoothed explicit formula gives

$$\Psi_J(x) = c_w x - \sum_{\rho} x^{\rho}\widehat{w}(\rho) + O_w(\log|D_J| + n_J),$$

where ρ runs over nontrivial zeros with multiplicity. This is the uniform formula used in OpenAI (2026a, Lemma 11.2, p. 45). For completeness, the uniformity of its zero term follows as follows. Three integrations by parts give $|\widehat{w}(\sigma+it)| \ll_w (1+|t|)^{-3}$, uniformly for $0 \le \sigma \le 1$. The standard unit-height zero count is

$$N_J(T+1) - N_J(T) \ll \log|D_J| + n_J\log(T+3).$$

Summing over integer height intervals therefore bounds $\sum_\rho |\widehat{w}(\rho)|$ by $O_w(\log|D_J| + n_J)$. The zero-free assumption bounds $|x^\rho|$ by x^θ. The pole at one gives the displayed main term, and the gamma factors and trivial zeros give the stated remainder. These estimates prove the lemma. Thus the proof is valid for any fixed θ, even though OpenAI's stated lemma selects one particular zero-free width. □

Proposition 18.4 (The pure-field prime detector)**.** *With F, K, E as in Theorem 1.6, put $\mathcal{H}(x) = \Psi_K(x) + \Psi_{\mathbb{Q}}(x) - \Psi_F(x)$. Its prime-power coefficients are nonnegative, including at ramified primes. At an unramified rational prime ℓ, its coefficient at ℓ^j is*

$$r\log\ell\;\mathbf{1}_{\mathcal{C}}(g_\ell^j),$$

where g_ℓ is a Frobenius element and $\mathcal{C} = N \setminus \{1\}$ is the class of nontrivial translations. In particular, if no auxiliary prime occurs in $(x, 2x)$, then $\mathcal{H}(x) \ll_w r\sqrt{x}$.

Proof. The virtual character $Q_0 + 1 - P$ in Lemma 18.1 equals r on $\mathcal{C}$ and zero elsewhere. The coefficient assertion follows from the local logarithmic derivative of the corresponding Artin Euler factors. At a ramified prime, the coefficient is the trace of a power of Frobenius on inertia invariants. If I is inertia and g a lift of Frobenius, that trace for this virtual representation is

$$\frac{1}{|I|}\sum_{h\in I}(Q_0 + 1 - P)(g^j h),$$

which is nonnegative and at most r.

If no auxiliary prime occurs in $(x, 2x)$, all unramified first-power terms in the support of w vanish. The other unramified terms have $j \geq 2$, so their total is at most

$$r\|w\|_\infty \sum_{j\geq 2} \vartheta((2x)^{1/j}) \ll_w r\sqrt{x}.$$

The last inequality follows from the elementary bound $\vartheta(y) \ll y$: the square term costs $O(\sqrt{x})$ and the remaining $O(\log x)$ terms cost $O(x^{1/3}\log x) = O(\sqrt{x})$. Only p and r can ramify in E. For each, summing its powers in the support of w costs $O_w(r\log(2x))$, also absorbed by $O_w(r\sqrt{x})$. This proves the upper bound. □

Theorem 18.5 (Polynomial search from one zero-free half-plane)**.** *Assume A and let $1/2 \leq \theta < 1$ be fixed. If $\zeta_{\mathbb{Q}(p^{1/r})}$ has no nontrivial zero with real part greater than θ, then there is an auxiliary prime*

$$\ell \ll_\theta \left(r\log(rp)\right)^{1/(1-\theta)}.$$

In particular, RH for that pure-field zeta function gives $\ell \ll r^2\log^2(rp)$ and deterministic field construction in time polynomial in r and $\log p$.

Proof. The discriminant of $X^r - p$ is, up to sign, $r^r p^{r-1}$, so $\log|D_F| \leq r\log r + (r-1)\log p$. Under A, Lemma 18.3 applies to K and $\mathbb{Q}$ with $\theta = 1/2$, and by hypothesis it applies to F with the stated θ. Thus, with $B = r\log(rp)$,

$$\mathcal{H}(x) = c_w x + O_w(Bx^\theta).$$

If there is no auxiliary prime in $(x, 2x)$, Proposition 18.4 bounds the left side by $O_w(r\sqrt{x})$. Set $x = TB^{1/(1-\theta)}$. After division by x, the analytic error is $O_w(T^{-(1-\theta)})$, and $r/\sqrt{x} \leq T^{-1/2}$ up to an absolute constant. A sufficiently large fixed T, depending only on w and θ, makes both errors smaller than $c_w/3$, a contradiction. Hence a suitable prime is less than $2x$. □

Proposition 18.6 (A finite one-sided sufficient condition)**.** *Fix $0 < \varepsilon < 1$. Under A, if*

$$x \geq C_{w,\varepsilon} r^2 \log^2 r, \qquad \Psi_F(x) \leq (2-\varepsilon)c_w x,$$

for a sufficiently large constant $C_{w,\varepsilon}$, then an auxiliary prime occurs in $(x, 2x)$. Consequently a uniform polynomial upper bound on some such $x = x(p, r)$ would suffice for the classical field-construction algorithm.

Proof. Under A, the two abelian fields give

$$\Psi_K(x) + \Psi_{\mathbb{Q}}(x) = 2c_w x + O_w(r\log r\sqrt{x}).$$

The assumed upper bound for Ψ_F therefore gives $\mathcal{H}(x) \geq \varepsilon c_w x - O_w(r\log r\sqrt{x})$. If no auxiliary prime occurs, Proposition 18.4 gives the contradictory upper bound $O_w(r\sqrt{x})$ once $C_{w,\varepsilon}$ is large enough. □

The positivity in Proposition 18.4 gives only $\Psi_F(x) \leq 2c_w x + O_w(r\log r\sqrt{x})$ under A. It reaches the threshold two without a fixed gap. No argument here proves the strict one-sided estimate from A. This is the precise remaining step for this route; a finite-scale estimate would be enough, so global RH for ζ_F is a stronger sufficient premise than the construction itself requires.

19 Nonresidues from Frobenius eigenspaces

This section checks the extension-field step and gives a direct proof of its unconditional even-degree quadratic case.

Proof of Theorem 1.8 for odd r. Under A, the Dirichlet character of order r modulo p has a nontrivial value at an integer of size $O((\log p)^2)$, by the small-character bound. Testing these integers gives $b \in \mathbb{F}_p^\times$ with $\zeta = b^{(p-1)/r} \neq 1$. As r is prime, ζ has order r. If $r \nmid m$, then

$$b^{(p^m-1)/r} = \zeta^{1+p+\cdots+p^{m-1}} = \zeta^m \neq 1,$$

so the constant b is already a nonresidue in the given field.

Suppose $r \mid m$ and put $Q = p^{m/r}$. In the cyclic group $\mathbb{F}_{p^m}^\times$, the image of the map $x \mapsto x^{Q-1}$ has order $S = 1 + Q + \cdots + Q^{r-1}$. Because $Q \equiv 1 \pmod r$, $r \mid S$, so this image contains ζ. The $\mathbb{F}_p$-linear equation $u^Q = \zeta u$ therefore has a nonzero solution. Compute one by linear algebra. Writing $Q = 1 + rc$ and expanding modulo r^2 gives

$$S \equiv r + rc\frac{r(r-1)}{2} \equiv r \pmod{r^2}.$$

Thus $u^{(p^m-1)/r} = \zeta^{S/r} = \zeta \neq 1$. All exponents have $O(m \log p)$ bits; in this branch $r \leq m$. The claimed uniform polynomial bound follows. □

The quadratic case. For odd m, a prime-field nonsquare remains a nonsquare, so the preceding small-character bound applies. Let now m be even and set $Q = p^{m/2}$. If $p \equiv 1 \pmod 4$ or $4 \mid m$, then $Q \equiv 1 \pmod 4$. For a polynomial generator α of the given field, $u = \alpha^Q - \alpha \neq 0$ and $u^Q = -u$. Therefore

$$u^{(p^m-1)/2} = (-1)^{(Q+1)/2} = -1.$$

It remains that $p \equiv 3 \pmod 4$ and $m \equiv 2 \pmod 4$. Compute $K = \ker(z \mapsto z^{p^2} - z)$ and choose $v \in K \setminus \mathbb{F}_p$. Then $u = v^p - v \neq 0$, $u^p = -u$, and $c = u^2$ is a nonsquare in $\mathbb{F}_p$. Since $-c$ is a square, $t = (-c)^{(p+1)/4}$ satisfies $t^2 = -c$. Put $i = u/t$; thus $i^2 = -1$ and $K = \mathbb{F}_p(i)$.

Start with $z = i$, which has order four and norm one. While $N_{K/\mathbb{F}_p}(z) = 1$, write $z = a + bi$. Its order is a power of two at least four; since the 2-part of $\mathbb{F}_p^\times$ has order two, $b \neq 0$. The nonzero numbers $(a+1)/2$ and $(a-1)/2$ have product $-b^2/4$, a nonsquare. Exactly one, say δ, is a square. Test by Euler's criterion, put $x = \delta^{(p+1)/4}$, and set $y = b/(2x)$. The relation $a^2 + b^2 = 1$ verifies $(x + yi)^2 = z$. The norm of this root is 1 when $\delta = (a+1)/2$, and -1 when $\delta = (a-1)/2$. Replace z by the root.

Each step doubles the order. The norm-one group has order $p+1$, so the process must reach norm -1 within $O(\log p)$ steps. At that point

$$z^{(p^2-1)/2} = N_{K/\mathbb{F}_p}(z)^{(p-1)/2} = -1.$$

The extension degree $[\mathbb{F}_{p^m} : K] = m/2$ is odd, so this element remains a nonsquare in the given field. The entire construction uses polynomially many field operations and linear-algebra operations, with no analytic hypothesis. In characteristic two every element is a square, so no quadratic-nonresidue assertion is made there. □

20 Zero divisors and the earlier quasipolynomial route

This section proves the coefficient-unit step explicitly and shows how it completes the earlier factorization algorithm. The final polynomial-time algorithm of Tang uses a different prime-field theorem; the two degree dependences must be kept separate.

Proposition 20.1 (Splitting from a maximal coefficient unit)**.** *Let A_0 be an explicitly presented split algebra $\mathbb{F}_p^t$, with $t > 1$, and let τ be an automorphism of prime order $r \neq p$. Put $s = \mathrm{ord}_r(p)$, $R = r^{v_r(p^s-1)}$, and $D = \mathbb{F}_p[T]/\Phi_r(T)$. Suppose a maximal Teichmüller unit $\eta \in D^\times$ is known. A nonzero zero divisor in A_0 can be constructed in time polynomial in $t, r, \log p$. The same assertion holds for $r = 2$, with $D = \mathbb{F}_p$ and a scalar unit of the full 2-primary order.*

Proof. Compute the ideal generated by $\tau(b) - b$ as b runs through a basis of A_0. In split coordinates, its support is precisely the set of points moved by τ. It is nonzero. If it is proper, its identity element is the required zero divisor. Otherwise all cycles have length r; assume this from now on.

Work in $C = A_0 \otimes_{\mathbb{F}_p} D$ and extend τ trivially on D. Write $\zeta = T \bmod \Phi_r$. We first construct a unit $y \in C$ with $\tau(y) = \zeta y$. On any nonzero τ-invariant component ideal eC, the projection

$$P_e(z) = \sum_{i=0}^{r-1} \zeta^{-i} \tau^i(z)$$

is not the zero map. Over an algebraic closure this follows by applying it to a point idempotent in a length-r orbit: the r terms have disjoint supports. Since the map is $\mathbb{F}_p$-linear, a basis search finds $x = P_e(z) \neq 0$. The ideal xC has an identity f, computable by solving its multiplication equations. This identity is fixed by τ, because $\tau(x) = \zeta x$ preserves the ideal. The element x is a unit on fC. Replace e by $e - f$ and repeat. At each step the remaining dimension decreases. Adding the vectors on their disjoint supports gives the required unit y in at most $\dim C$ steps.

Every field factor of C has cardinality p^s. Write $p^s - 1 = Rw$, and choose a positive h with $wh \equiv -1 \pmod r$. Then $z = y^{wh}$ has r-power order dividing R and $\tau(z) = \zeta^{-1} z$. For odd r, use the lifts ω modulo R and put

$$x = \prod_{b \in \mathbb{F}_r^\times} \rho_b^{-1}(z^{\omega(b)}).$$

Reindexing proves $\rho_c(x) = x^{\omega(c)}$. Because τ commutes with ρ_b and $b^{-1}\omega(b) \equiv 1 \pmod r$, it also gives $\tau(x) = \zeta^{-(r-1)} x = \zeta x$. For $r = 2$ take $x = z$; the same properties hold.

The two Teichmüller elements x and η generate a noncyclic r-group. Indeed its exponent divides R, and a cyclic such group containing η of order R would be generated by η. Then x would be a power of η and fixed by τ, contradicting $\tau(x) = \zeta x$.

Here is a concrete way to turn this discrepancy into a zero divisor. In every field component of C, η generates the entire r-primary group, so write $x = \eta^{c_i}$ in that component for the correctness argument. The algorithm does not know these exponents. Let $R = r^a$ and $\xi = \eta^{R/r}$. Starting with $x_0 = x$, recover common base-r digits successively. At step k, $0 \leq k < a$, the previous k common digits have been removed, and compute $v = x_k^{r^{a-1-k}}$. Each component of v is one of ξ^j, $0 \leq j < r$. Test the multiplication matrix of every $v - \xi^j$. If one is zero, the kth digit is common to all components; replace x_k by $x_k \eta^{-jr^k}$ and continue. If none is zero, at least one is nonzero and singular, because some coordinate value occurs. This gives a nonzero zero divisor z_0 in C. All a steps cannot finish with a common digit, since that would give $x \in \langle \eta \rangle$.

The vanishing support of $z_0 = v - \xi^j$ is invariant under every ρ_b. Indeed v and ξ^j are Teichmüller elements, and raising two r-primary units to the common exponent $\omega(b)$ preserves their equality, since that exponent is prime to r. The identity of $\mathrm{Ann}_C(z_0)$ is therefore a nonzero proper idempotent fixed by all ρ_b. It lies in $C^{\mathbb{F}_r^\times} = A_0$, and is the required zero divisor. All ideals, identities, and fixed spaces are computed by linear algebra in dimension at most $t(r-1)$. Also $a \log r \leq s \log p$. This proves the polynomial bit bound. □

This proof specializes the constructions of Ivanyos et al. (2012, Lemma 4.2, author p. 16; Proposition 4.6, pp. 18–19). It makes explicit both the nonvanishing eigenvector and the descent

of the idempotent from the cyclotomic algebra. For an arbitrary nonidentity automorphism, its order can first be determined using the known exponent $t!$ of the permutation group on t points. The factorization of $t!$ is obtained by trial division up to t; successively remove prime factors while a matrix power is the identity. Taking a suitable power then gives an automorphism of prime order.

20.1 A complete transfer from a factor-count witness

The next proof gives the form of the factor-count theorem needed for the cyclotomic and elliptic constructions below. It is the prime-degree consequence of Bhargava et al. (2017, Theorems 1.2 and 1.4, p. 2). We include the nonvanishing construction of the resolvent and its relative norm, so that the conclusion uses a deterministically chosen unit in every component.

Proposition 20.2 (A factor-count witness gives the coefficient unit)**.** *Let k be a presented finite field of cardinality q and characteristic p, let $r \neq p$ be an odd prime, and let $f \in k[X]$ be squarefree of degree N. Suppose that, for some d divisible by r, the number n of irreducible degree-d factors of f is not divisible by r. In deterministic time polynomial in N and $\log q$, one can construct the maximal Teichmüller unit in $D = k[T]/\Phi_r(T)$, and hence a presentation of $\mathbb{F}_{q^r}$. Neither the degree-d factors nor the factorization of Φ_r is an input to this algorithm.*

Proof. Distinct-degree factorization uses only the gcds of successive $X^{q^j} - X$ with the remaining polynomial, for $1 \leq j \leq N$. It determines the counts and gives the product g of the degree-d factors in polynomial time. Set $A = k[X]/(g)$, so that, for the correctness argument, A is a product of n fields of degree d. Put

$$s = \operatorname{ord}_r(q), \quad h = \gcd(d, s), \quad Q = q^s, \quad R = r^{v_r(Q-1)}, \quad C = A \otimes_k D.$$

Above each component $\mathbb{F}_Q$ of D, the algebra C has nh components, each of degree d/h over $\mathbb{F}_Q$. This follows by decomposing the tensor product of finite fields: their common subfield has degree h, and their compositum has degree ds/h over k.

Let F denote Qth-power Frobenius on C. It fixes D. The automorphism $\sigma = F^{d/(rh)}$ has order r and has only length-r cycles on the geometric points of C. Here $r \nmid h$, since $h \mid s \mid r - 1$. With $\zeta = T \bmod \Phi_r$, construct a unit $\beta \in C$ satisfying $\sigma(\beta) = \zeta\beta$. The support construction in Proposition 20.1 does this directly: on each remaining invariant ideal apply $\sum_{i=0}^{r-1} \zeta^{-i}\sigma^i$ to basis vectors, take the identity of the nonzero image ideal, and continue on its complementary support. The projection is nonzero on each length-r orbit, and the sum of the vectors on disjoint supports is a unit. This uses no factorization or division by the order of an auxiliary group.

Set $M = Q^{d/(rh)}$. On each field component, $\beta^{M-1} = \zeta$. Since $M \equiv 1 \pmod r$ and r is odd,

$$1 + M + \cdots + M^{r-1} \equiv r \pmod{r^2}.$$

Consequently

$$\beta^{(Q^{d/h}-1)/r} = \zeta$$

in every component. Take the algebra norm $b = N_{C/D}(\beta)$. For each of the nh components over $\mathbb{F}_Q$, the norm to $\mathbb{F}_Q$ has rth-power character ζ. Their product therefore gives

$$b^{(Q-1)/r} = \zeta^{nh}.$$

Since $r \nmid nh$, this is a primitive rth root in every component of D.

Write $Q - 1 = Rw$ and set $\delta = b^w$. Then $\delta^{R/r} = \zeta^{nh}$, so δ has full order R. Using the Teichmüller lifts modulo R, put

$$\eta = \prod_{c \in \mathbb{F}_r^\times} \rho_c(\delta)^{\omega_a(c)^{-1}}, \qquad a = v_r(Q - 1).$$

Reindexing proves $\rho_c(\eta) = \eta^{\omega_a(c)}$, whereas $\eta^{R/r} = \zeta^{nh(r-1)} = \zeta^{-nh}$ proves that its order is R. Corollary 3.1 now constructs the required field.

All algebras have dimension at most $N(r-1)$, and $r \le N$. The norm is the determinant of a multiplication matrix over D; a division-free determinant algorithm computes it over this commutative algebra in polynomial time. The exponents have polynomially many bits. This proves the uniform bit bound. □

The nonvanishing choice is essential. In $\mathbb{F}_7[\alpha]/(\alpha^3 - 2)$, with $\zeta = 4$, one has $\alpha^7 = 4\alpha$, $\alpha^{49} = 2\alpha$, and $\alpha + 4\alpha^7 + 2\alpha^{49} = 0$. Thus a resolvent of a fixed polynomial generator can vanish. The support construction above, or a basis search in a single field eigenspace, proves the form used here without assuming that a prescribed resolvent is nonzero.

20.2 Proof of Theorem 1.10

Proof. Squarefree decomposition and inverse Frobenius retain all multiplicities with polynomial cost. For a squarefree factor g, set $A = \mathbb{F}_{p^m}[X]/(g)$ and compute $B = \ker(z \mapsto z^p - z)$ by $\mathbb{F}_p$-linear algebra. If g has t irreducible factors, then $B \cong \mathbb{F}_p^t$; this follows from the Chinese remainder theorem and does not require those factors as input. When $t = 1$, g is irreducible.

If $p \le d$ and $t > 1$, choose $b \in B \setminus \mathbb{F}_p 1$ and test $b - c$ for $c \in \mathbb{F}_p$. One of these is a nonzero zero divisor, so at most $p \le d$ tests split the polynomial. Repeating gives the first claim.

For $p > d$, apply the unconditional algebra theorem of Ivanyos et al. (2012, Theorem 5.8, author pp. 28–29) to B over its scalar subalgebra $\mathbb{F}_p$. It gives either a nonzero zero divisor or a semiregular automorphism of order t. Choose a prime $r \mid t$ and take its order-r power. The moved-point ideal in Proposition 20.1 either splits B or reduces to its semiregular case. Here $r \le d < p$.

There are four ways to obtain the maximal coefficient unit for that prime. If r is odd and $r \mid m$, extract the degree-r subfield of the given coefficient field and apply Theorem 1.7. If $r = 2$ and m is even, use the unconditional part of Theorem 1.8 to find a nonsquare z in that field. Its norm b to $\mathbb{F}_p$ is a nonsquare because $b^{(p-1)/2} = z^{(p^m-1)/2} = -1$. Raising b to the odd part of $p - 1$ gives the required scalar unit. If $r \mid p - 1$, find a prime-field nonresidue under A. For odd r, construct the irreducible binomial $X^r - c$ and apply Theorem 1.7; for $r = 2$, raise the nonsquare to the odd part of $p - 1$. If $v_r(p^{\mathrm{ord}_r(p)} - 1) = 1$, use the tautological unit ζ directly. Every remaining r lies in $\mathcal{B}$; its one pure-field zeta hypothesis gives the field by Theorem 1.6, and then the unit by Theorem 1.7. Proposition 20.1 now splits B.

For a zero divisor $b \in B \subset A$, represented by a polynomial of degree below $\deg g$, the gcd $\gcd(g, b)$ is a proper nonconstant factor. There are at most $2d - 1$ nodes in the splitting tree. The IKRS cost at each node is $d^{O(\log d)}(\log(p^m))^{O(1)}$, and every other step has polynomial cost. This proves the stated bound. Under the final promise, select the given-subfield or valuation-one branch before any analytic branch; no analytic hypothesis is then used. □

The earlier construction through Adleman–Lenstra and Bach et al. (2001, Theorem 2, printed p. 7) uses the sufficient family $\zeta_{\mathbb{Q}(\zeta_{2\nu}, p^{1/\nu})}$, with construction degrees $\nu \le d^2$, together with $\zeta_{\mathbb{Q}}$. Theorem 1.10 states the smaller prime-degree family obtained by applying Lenstra's equivalence directly. Both are sufficient hypotheses for the earlier quasipolynomial algorithm; neither is a necessity claim, and neither is Tang's polynomial-time theorem.

21 Tang's polynomial algorithm and its analytic input

This section separates Tang's two input interfaces and identifies the zero-free information that makes the prime-field calls polynomial. Tang's Theorem 1.1 (Tang, 2026a, pp. 1–2) factors a nonzero dense polynomial over a presented finite field and, as a separate operation, constructs a presentation from an integer promised to be a prime power. In the second interface the

polynomial coefficients are entered after the power basis has been constructed. A list of integers called coefficients, without a field presentation, does not specify the first problem.

The algebraic input is OpenAI's Theorem 1.2 (p. 3) and its explicit bit bound, Proposition 10.1 (pp. 41–44) OpenAI (2026a). For a degree-n prime-field input put

$$B = 20 + (n+1)(\lceil \log_2 p \rceil + 1).$$

When $n \geq 2$ and $p > B^{200000}$, the input table contains, for each prime $r \leq n$, a prime ℓ such that

$$\ell \notin \{2, 3, p, r\}, \qquad \ell \equiv 1 \pmod{12r}, \qquad p^{(\ell-1)/r} \not\equiv 1 \pmod{\ell}. \tag{23}$$

The other branch uses the elementary prime-field algorithm and an empty table. If E is the largest table entry, or 2 for an empty table, the bound is

$$O(B^{500000} + B^{100} E^{20}). \tag{24}$$

This charges the numerical value of E, rather than its bit length. The algebraic theorem contains no zero-free hypothesis. Its proof and bit bound are external inputs here.

21.1 Two prime-ideal sums and one sufficient zeta hypothesis

For the cyclic Kummer extension M/K used below, abelian factorization gives $\zeta_M = \zeta_K \prod_{\chi \neq 1} L_K(s, \chi)$. The nontrivial finite-order factors are entire. Hence a zero of ζ_K in the strip is also a zero of ζ_M. A zero-free half-plane for ζ_M suffices for both prime-ideal estimates.

The following proof records why a fixed half-plane suffices, including the prime $r = 2$. It is Tang's Lemma 2.1 (Tang, 2026a, pp. 2–4), using the smoothed explicit formula of (OpenAI, 2026a, Lemma 11.2, p. 45) and (Cureton and Thorner, 2026, Theorem 1.1, p. 3).

Proposition 21.1. *Let $n \geq 2$, let $p > \max(n, 3)$ be prime, and let $r \leq n$ be prime. Set*

$$K = \mathbb{Q}(\mu_{12r}), \qquad M = K(p^{1/r}).$$

If both ζ_K and ζ_M have no nontrivial zero with real part greater than $31/32$, then a prime satisfying (23) exists with $\ell = O(B^{32})$. The implied constant is absolute.

Proof. The prime p is unramified in K. At a prime above p, the valuation of p is one, so $X^r - p$ is Eisenstein. Thus $[M : K] = r$; since $\mu_r \subset K$, the extension is cyclic. Write $d_F = [F : \mathbb{Q}]$ and D_F for the absolute discriminant of a number field F. The cyclotomic degree and the power-basis discriminants give

$$d_K \leq 4r, \quad \log D_K \leq d_K \log(12r), \quad d_M = r d_K,$$

and

$$\log D_M \leq r \log D_K + d_K\big(r \log r + (r-1)\log p\big).$$

Consequently

$$(\log D_K + d_K) + r^{-1}(\log D_M + d_M) = O(r(1 + \log p)) = O(B). \tag{25}$$

Here $r < p$, and the constants do not depend on the two fields.

Fix a nonzero nonnegative $w \in C_c^\infty((1,2))$, and let $c_w = \int_1^2 w(t)\,dt > 0$. Define

$$\Psi_F(x) = \sum_{\mathfrak{p}} \sum_{j \geq 1} \log N\mathfrak{p}\, w((N\mathfrak{p})^j / x).$$

The smoothed explicit formula expresses this as $c_w x$, minus $\sum_\rho x^\rho \widehat{w}(\rho)$, and an error $O_w(\log D_F + d_F)$. Uniform zero counting in unit height intervals and three integrations by parts give $\sum_\rho |\widehat{w}(\rho)| = O_w(\log D_F + d_F)$. The stated zero-free half-plane therefore gives

$$\Psi_F(x) = c_w x + O_w(x^{31/32}(\log D_F + d_F)).$$

Although the statement of OpenAI's Lemma 11.2 fixes a different width, its displayed proof establishes precisely this implication for a fixed width. Applying it to K and M, with their indicated weights, yields

$$\Psi_K(x) - r^{-1}\Psi_M(x) = (1 - r^{-1})c_w x + O_w(Bx^{31/32}). \tag{26}$$

Suppose no prime satisfying (23) lies in $(x, 2x)$. A degree-one prime of K above $\ell \notin \{2, 3, p, r\}$ has $\ell \equiv 1 \pmod{12r}$. The supposition says that p is an rth power modulo ℓ. Thus $X^r - p$ has r distinct roots in the residue field, and this prime of K splits completely in M. Its first-power contribution to Ψ_K cancels with the r contributions to $r^{-1}\Psi_M$. Every uncancelled term from M has a negative sign and can be discarded in an upper bound. Apart from the four exceptional rational primes, the remaining positive terms lie over rational primes at most $\sqrt{2x}$. For any prime ideal, the sum over eligible exponents is at most $\|w\|_\infty \log(2x)$. Hence

$$\Psi_K(x) - r^{-1}\Psi_M(x) \le C_w r\sqrt{x}\log(2x). \tag{27}$$

Put $x = TB^{32}$ with T an absolute constant. After division by x, the error in (26) is $O_w(T^{-1/32})$, whereas the right side of (27) is $O_w(T^{-1/2}(1 + \log T))$, uniformly in B. For sufficiently large fixed T, both are less than $c_w/4$, contradicting the main term $(1 - r^{-1})c_w \ge c_w/2$. Thus $\ell < 2TB^{32}$. □

Accepting CT provides exactly the two zero-free statements in Proposition 21.1. Their Hecke-character corollary is not required. Nor does this proof invoke the stronger analytic input quoted as OpenAI's Theorem 11.1: a half-plane of width 10^{-6} for every finite-order Hecke character of every cyclotomic field containing μ_{12}. Substituting $E = O(B^{32})$ in (24) gives $O(B^{500000})$, since $100 + 20 \cdot 32 = 740$. Upward search finds the primes without knowing the implied constant.

For $r = 2$, the fields are $K = \mathbb{Q}(\mu_{24})$ and $M = K(\sqrt{p})$. They are abelian over $\mathbb{Q}$, so Dirichlet GRH supplies RH for both zeta functions. For odd r, the additional radical extension is generally nonabelian over $\mathbb{Q}$. Tang's proof supplies a polynomial search bound by the fixed Dedekind zero-free theorem; it does not derive that theorem, or RH for the radical field, from Dirichlet GRH.

21.2 Factoring in a field that has already been presented

The absolute Frobenius fixes a much smaller scalar field than the qth-power map. This gives the following explicit reduction, due to Tang (Tang, 2026a, Proposition 3.1, pp. 4–5), in the tradition of Berlekamp's reduction to prime-field root finding.

Proposition 21.2. *Given $\mathbb{F}_q = \mathbb{F}_p[U]/(h)$, with h monic irreducible of degree m, a nonzero degree-n polynomial over $\mathbb{F}_q$ can be completely factored using at most $\max(0, n-1)$ prime-field factorization calls. Each query is split and squarefree of degree at most n. All remaining bit operations have cost polynomial in $n+1$ and $\log_2 q$.*

Proof. First let g be monic and squarefree, and put $A = \mathbb{F}_q[X]/(g)$. The matrix of $b \mapsto b^p - b$ over $\mathbb{F}_p$ is obtained by binary powering the $m \deg g$ elements of the tensor power basis. Compute its kernel B by Gaussian elimination. If g has s irreducible factors, the Chinese remainder theorem identifies A with a product of s finite fields and B with $\mathbb{F}_p^s$. Thus $\dim B = s$.

When $s = 1$, g is irreducible. Otherwise choose a basis element $b \in B$ outside $\mathbb{F}_p 1$. Its minimal polynomial is found from the first dependence among $1, b, b^2, \ldots$. It is the product of $Y - a$ over the distinct coordinate values of b, so it is split, squarefree, and has degree between 2 and s. Factor it by one prime-field call. For each root a, the polynomial $\gcd(g, b - a)$ selects precisely the components with coordinate a. These gcds partition the factors into at least two nonempty groups. Recurse on each group. A rooted splitting tree with s leaves and at least two children at each internal vertex has at most $s - 1$ internal vertices, hence at most $s - 1$ queries.

For a general monic f, the polynomial $g = f/\gcd(f, f')$ is the product of the irreducible factors whose multiplicities are not divisible by p. Indeed an irreducible u is separable over a finite field, and differentiating $u^e v$, with $u \nmid v$, shows that its multiplicity in $\gcd(f, f')$ is $e - 1$ if $p \nmid e$ and e if $p \mid e$. Factor g, remove the full multiplicity of each factor by exact division, and take the pth root of the remainder. This root is obtained by dividing monomial exponents by p and applying $a \mapsto a^{p^{m-1}}$ to the coefficients. Repeat. Different squarefree layers contain disjoint sets of irreducibles. Their splitting trees therefore use at most $n - 1$ queries altogether when $n > 0$. Constants use none.

The matrices have dimension at most mn, all quotient arithmetic is polynomial in that dimension, and every exponent has $O(m \log p)$ bits. There are only polynomially many divisions and root steps. Restoring the leading coefficient and multiplicities therefore preserves the claimed bound. □

21.3 Constructing a presentation from its order

Here the prime-field oracle also supplies the roots needed to build the coefficient field. The reduction from construction to factoring goes back to Shoup (1990, Theorem 3.1). The query count below is Tang's Proposition 4.1 (Tang, 2026a, pp. 5–7); its arithmetic construction is included to expose the precise point at which a maximal root of unity is obtained.

Proposition 21.3. *For a prime p and an integer $m \geq 1$, a monic irreducible polynomial of degree m over $\mathbb{F}_p$ can be constructed with $O(m^2(1 + \log_2 p))$ prime-field factorization calls of degree at most m and polynomial additional bit cost in m and $\log_2 p$.*

Proof. The case $m = 1$ is immediate. Factor m by trial division, whose cost is polynomial in m, and treat each exact prime-power divisor r^a separately.

For odd $r \neq p$, factor Φ_r over $\mathbb{F}_p$ and choose one irreducible factor. This constructs a field E of degree at most $r - 1$ containing a primitive rth root u. Let $t = v_r(|E| - 1)$. An element of order r^j, for $j < t$, has an rth root in the cyclic group $E^\times$, and every such root has order r^{j+1}. Starting at u, factor $V^r - b$ over E by Proposition 21.2 at each step. A linear factor gives the next root. After $t - 1$ steps this produces an element c of order r^t.

The polynomial $V^{r^a} - c$ is irreducible over E. To prove this, any one of its roots has order r^{t+a}. Its degree over E is therefore the least e for which r^{t+a} divides $|E|^e - 1$. The geometric sum formula for the part of e prime to r, followed by the binomial theorem for each factor of r, gives

$$v_r(|E|^e - 1) = t + v_r(e).$$

The least such e is r^a. Thus this binomial constructs a field containing the degree-r^a extension of $\mathbb{F}_p$.

For $r = 2 \neq p$ and $p \equiv 1 \pmod 4$, use $E = \mathbb{F}_p$, start with -1, and take successive square roots until their order is $2^{v_2(p-1)}$. The same valuation identity is valid because $v_2(p - 1) \geq 2$, and $V^{2^a} - c$ is irreducible. For $p \equiv 3 \pmod 4$, first use $E = \mathbb{F}_p[i]$, $i^2 = -1$. This already completes $a = 1$. For $a \geq 2$, start at i, take square roots to the maximal 2-power order in $E^\times$, and use $V^{2^{a-1}} - c$. Its degree over $\mathbb{F}_p$ is 2^a.

For $r = p$, build an Artin–Schreier tower. In each current field E, the $\mathbb{F}_p$-linear map $z \mapsto z^p - z$ has kernel $\mathbb{F}_p$ and image of codimension one. Choose a basis vector c outside that

image. Then $V^p - V - c$ has no root in E; its roots in a splitting field are the translates of one root by $\mathbb{F}_p$. Its splitting field is generated by one root and its Galois group is a nontrivial subgroup of the additive group $\mathbb{F}_p$. The group therefore has order p, proving irreducibility. Repeat a times.

In each ambient field constructed above, find the degree-r^a subfield as the kernel of $z \mapsto z^{p^{r^a}} - z$. Choose a basis element outside the degree-r^{a-1} subfield. Every proper subfield of the first is contained in the second, so this element has degree r^a. A dependence among its successive powers gives its minimal polynomial over $\mathbb{F}_p$.

The degrees r^a for distinct primes are coprime. Adjoining these generators successively therefore gives fields throughout: a degree-r^a irreducible polynomial remains irreducible over an extension of coprime degree, by the length of its Frobenius orbit. The resulting compositum has degree m. The sum of its chosen generators is primitive. Indeed, if its degree is $e \mid m$, then for each $r^a \| m$ the field of degree $\operatorname{lcm}(e, m/r^a)$ contains the sum and all the other generators, hence this generator as well. Thus $r^a \mid \operatorname{lcm}(e, m/r^a)$, and coprimality gives $r^a \mid e$. Therefore $e = m$. Compute its minimal polynomial.

The root chain for a prime r has $O(r(1+\log_2 p))$ steps, and each uses at most $r-1$ prime-field queries of degree at most r. The initial cyclotomic query has degree $r-1$. Since $\sum_{r \mid m} r \le m$, the sum of the squares of these primes is at most m^2, giving the stated query bound. The ambient degrees are at most m^2. Multiplication tables, linear systems, minimal polynomials, and binary powering in these tables consequently take polynomial bit time. □

For a binary integer q promised to be a prime power greater than one, test exact integer roots with exponents from $\lfloor \log_2 q \rfloor$ down to 1. The largest successful exponent gives $q = p^m$ with p prime, by unique factorization. Integer-root tests are polynomial in $\log_2 q$. Since $m \log_2 p = \log_2 q$, Proposition 21.3 and the prime-field algorithm prove the presentation interface in polynomial bit time. Proposition 21.2 separately proves the factoring interface. Credit for this polynomial algorithm belongs to Tang and to the cited prime-field and zero-free results.

A recent sparse-factorization extension makes a related separation of interfaces. Kriebel (2026, Theorem 2, pp. 5–6; Section 4 and Corollary 3, p. 7) reduces represented-field factoring and prescribed-degree construction to prime-field factoring. Its polynomial conclusion uses OpenAI's prime-field theorem and explicitly invokes its cyclotomic Hecke input; it therefore supplies no auxiliary-prime bound from A alone. We use none of its sparse-factorization assertions here.

21.4 The two algorithmic interfaces for almost all characteristics

This section uses A, the Riemann hypothesis for rational Dirichlet L-functions, as its only new analytic input. Its algorithmic deductions also use the purely algebraic auxiliary-prime reduction of OpenAI and the two algebraic oracle interfaces of Tang. No Hecke or Dedekind zero-free strip is used.

Proof of Theorem 1.14. Fix a prime $r \le D$, put $Q = 12r$, and set $B_r = C_0 r^2 \log^2(Pr)$. One has $\varphi(Q) \le 4r$, also for $r = 2, 3$. Character orthogonality and Lemma 22.1 give

$$\sum_{\substack{u \le B_r \\ u \equiv 1 \pmod{Q}}} \Lambda(u) \log(B_r/u) = \frac{B_r}{\varphi(Q)} + O(\sqrt{B_r} \log(QB_r)).$$

The contribution from proper prime powers is $O(\sqrt{B_r} \log B_r)$, by the elementary Chebyshev bound $\psi(t) - \vartheta(t) \ll \sqrt{t}$. Uniformly for $r \le (\log P)^{a_0}$, $\log(QB_r) = O_{a_0}(\log\log P) = o(\log P)$. Consequently the sum over primes alone is at least $B_r/(2\varphi(Q))$ for sufficiently large P. Every prime contributes at most $(\log B_r)^2/4$. If

$$\mathcal{S}_r = \{\ell \le B_r : \ell \text{ prime}, \ \ell \equiv 1 \pmod{Q}\}, \qquad k_r = |\mathcal{S}_r|,$$

then

$$k_r \ge \frac{2B_r}{\varphi(Q)(\log B_r)^2} \ge \frac{B_r}{2r(\log B_r)^2}.$$

Since $\log B_r = O_{a_0}(\log\log P)$ and $r\log r \ge 2\log 2$, this implies, uniformly in this range,

$$k_r \log r \ge 2\log P \tag{28}$$

once P is large enough. Any fixed sufficiently large absolute C_0 works, with the threshold enlarged if needed.

For each $\ell \in \mathcal{S}_r$, choose a character χ_ℓ of exact order r modulo ℓ, and put $M_r = \prod_{\ell\in\mathcal{S}_r} \ell$. These choices are for the counting proof; the algorithm below does not compute these characters. Define

$$I_r(u) = r^{-k_r} \sum_{(j_\ell)\in\{0,\dots,r-1\}^{\mathcal{S}_r}} \prod_{\ell\in\mathcal{S}_r} \chi_\ell^{j_\ell}(u),$$

where a zeroth power denotes the principal character of its modulus. For $(u, M_r) = 1$, this is the indicator that u is an r-th power modulo every $\ell \in \mathcal{S}_r$; for nonunits it is zero. The Chinese remainder theorem shows that exactly one character in the sum is principal. All coefficients are the positive real number r^{-k_r}. Applying Lemma 22.1 at $y = 2P$ therefore yields

$$\sum_{u\le 2P} \Lambda(u) I_r(u) \log(2P/u) = \frac{2P}{r^{k_r}} + O(\sqrt{P}\log(2PM_r)).$$

The left side is nonnegative. Each bad prime $p \in (P/2, P]$ contributes at least $\log(P/2)\log 2$, because all candidate primes are smaller than $P/2$ for sufficiently large P. Moreover $\log M_r \le \vartheta(B_r) \ll B_r$, and (28) gives $2P/r^{k_r} \le 2/P$. Thus the number of bad primes for this r is

$$O\left(\frac{\sqrt{P}(B_r + \log P)}{\log P}\right) = O\left(\frac{\sqrt{P}\, r^2 \log^2(Pr)}{\log P}\right).$$

Summing over $r \le D$, and using $\sum_{r\le D,\ r \text{ prime}} r^2 \le D^3$, proves (3). Finally $B_r < P/2$ uniformly in the prescribed range, so $\ell \ne p$; the congruence gives $\ell > 12r$, excluding $2, 3, r$. For $D < 2$ the auxiliary conditions are empty. □

The exceptional proportion among primes in $(P/2, P]$ is

$$O\left(\frac{D^3 \log^2(PD)}{\sqrt{P}}\right) = o(1).$$

Thus the simultaneous statement holds for a set of characteristics of relative density tending to one. No assertion that a specified prime belongs to that set is needed for the algorithms to check their auxiliary tables.

Proof of Corollary 1.15. First consider a prime-field factorization query of degree $t \le n$, where n is the outer input degree. With $L = \lceil \log_2 p \rceil$, write $\mathcal{B}_n = 20 + (n+1)(L+1)$. If $n \le 1$, no table is needed. If $p \le \mathcal{B}_n^{200000}$, Berlekamp's scalar-enumeration branch factors all prime-field queries by testing at most p scalars per split. Its cost is polynomial in the outer input parameters, and its auxiliary table is empty. This branch also covers large degrees relative to p; it does not require the almost-all estimate.

Otherwise $p > \mathcal{B}_n^{200000} > n$. Since $P < 2p$, the auxiliary primes promised by (4), for every prime $r \le n$, all lie below

$$E_n = C_1 n^2 \log^2(pn)$$

for an absolute C_1. Enumerate candidates up to $\lceil E_n \rceil$, test primality by trial division, and test the congruence, exclusions, and modular power in (4). This requires polynomial time in

$n+1, L$, and the residue test needs no factorization of $\ell - 1$. The table works for every query of degree at most n. The quoted algebraic algorithm and its bound

$$O(\mathcal{B}_n^{500000} + \mathcal{B}_n^{100} E_n^{20})$$

therefore give a uniform polynomial prime-field procedure. Any table that is returned is checked directly. On inputs outside the stated promise, bounded search may report that the required table was not found; no incorrect factorization is asserted.

For a field of order p^e presented by an irreducible quotient or an equivalent multiplication table, Tang's Proposition 3.1 reduces complete factorization in degree n to at most $\max(0, n-1)$ prime-field factorizations of split square-free polynomials of degree at most n (Tang, 2026a, pp. 4–5). The query-degree bound is independent of e; the remaining bit cost is polynomial in $n + 1, e \log p$. This is why the same set of good characteristics works for every presented extension, with no bound on e. The input is promised to be a field and its polynomial is nonzero. A nonzero constant has an empty factor list.

For the presentation problem, exact integer-root tests recover p, m from the promised prime power q: among exponents from $\lfloor \log_2 q \rfloor$ down to 1, the largest exponent giving an integer root at least 2 is m, and that root is p. Binary search and capped powering take polynomial time in $\log q$. Tang's Proposition 4.1 constructs a degree-m irreducible polynomial using $O(m^2(1 + \log_2 p))$ prime-field factorization queries of degree at most m, with polynomial additional bit cost (Tang, 2026a, pp. 5–7). Apply the preceding construction with the parameter m in place of n. Since $m \log p = \log q$, the complete cost is polynomial in the input length. No factorization of the integer $p - 1$, primitive-root oracle, or nonabelian analytic estimate is used. □

This is an average consequence in the characteristic. It leaves the requested uniform auxiliary-prime assertion for every characteristic unchanged.

21.5 The example $p = 17$, $r = 3$

This example shows the exact step provided by each construction. Here $17 \equiv -1 \pmod 3$ and $17^2 - 1 = 288 = 2^5 \cdot 3^2$, so $\mathrm{ord}_3(17) = 2$ and $v_3(17^2 - 1) = 2$.

For the auxiliary-prime construction, $\ell = 7$ works: $17^{(7-1)/3} \equiv 2 \pmod 7$. Moreover $\mathrm{ord}_7(17) = 6$, so Φ_7 is irreducible over $\mathbb{F}_{17}$. If z is its root, the involution $z \mapsto z^{-1}$ fixes a degree-three subfield. Its element $z + z^{-1}$ has minimal polynomial

$$X^3 + X^2 - 2X - 1.$$

The auxiliary prime supplies a cyclotomic algebra from which the degree-three invariant field is extracted.

Tang instead begins with $E = \mathbb{F}_{17}[u]/(u^2 + u + 1)$. The order of u is three, whereas the maximal 3-power order in $E^\times$ is nine. The root chain has exactly one step. It can choose

$$c = 2 + 8u, \qquad c^3 = u, \qquad c^{-1} = 11 + 9u, \qquad c + c^{-1} = 13.$$

The cube roots of u are $2 + 8u$, $6 + 15u$, and $9 + 11u$. Their absolute traces are 13, 14, and 7. In the split algebra $E[V]/(V^3 - u)$, the element $b = V + u^2V^2 = V + V^{17}$ is fixed by absolute Frobenius and has minimal polynomial

$$Y^3 - 3Y + 1 = (Y - 7)(Y - 13)(Y - 14) \quad \text{over } \mathbb{F}_{17}.$$

Thus one cubic prime-field query separates all three roots and recovers c by a gcd. This is the step that constructs the maximal 3-power unit.

Now adjoin v with $v^3 = c$. It has order 27 and gives a degree-six field over $\mathbb{F}_{17}$. Since $17^3 \equiv -1 \pmod{27}$, the element $\alpha = v + v^{-1}$ lies in its degree-three subfield. The identity $(v + v^{-1})^3 - 3(v + v^{-1}) = c + c^{-1}$ gives

$$\alpha^3 - 3\alpha - 13 = 0.$$

It has degree three, so $X^3 - 3X - 13$, or $X^3 + 14X + 4$ in standard residues, is the resulting irreducible polynomial. For this small prime the prime-field calls use the empty-table branch: the example requires no zero-free theorem. If one nevertheless wants a prime meeting Tang's stronger congruence, $\ell = 37$ satisfies $\ell \equiv 1 \pmod{36}$ and $17^{12} \equiv 26 \pmod{37}$. The calculations above can also be reproduced by the programs *Prime-degree Kummer examples* and *Tang's construction at seventeen.*

21.6 The local revision to Cureton–Thorner

Tang's review remark (Tang, 2026a, Remark 1.3, p. 2) distinguishes the official version of Cureton and Thorner (2026) from a separately posted revision Tang (2026b). Their difference concerns the residual norm estimate in Lemma 8.1, official pp. 21–22, equation (8.4). The original passage appeals to a compact boundary pairing. The revision replaces that passage by an invariant-form argument; the theorem statement and the subsequent Sobolev argument are unchanged.

The comparison begins on the space V^{fin} of finite sums of spherical harmonics in the complex local model. For a unitary central character, Kazhdan–Patterson already realizes its global residues in L^2, before the extension to smooth local vectors (Kazhdan and Patterson, 1984, Theorem II.1.4, pp. 114–117; Theorem II.2.1 and the preceding paragraph, p. 118). With the other local vectors fixed, the form $B(u, v) = \langle \Theta(u), \Theta(v) \rangle_{L^2}$ is therefore well defined on V^{fin}. This use of the L^2 norm does not assume the continuity estimate being proved.

Let $\mathcal{J}$ be the complex standard intertwiner. In the compact picture its eigenvalue on harmonics of degree ℓ is

$$a_\ell = a_0 \prod_{j=1}^{\ell} \frac{j - 1/3}{j + 1/3} > 0, \qquad a_\ell^{-1} \le a_0^{-1}(1 + \ell).$$

The latter bound follows by comparing each inverse factor with $(j + 1)/j$. The pairing between normalized inducing representations with opposite real parameters is invariant, so $Q(u, v) = \langle \mathcal{J}^{-1} u, v \rangle_{L^2(\mathbb{S}^2)}$ is a positive invariant form on the target representation. Both forms are invariant under the maximal compact subgroup and the real Lie algebra. The compact types are distinct and have multiplicity one. Represent B relative to Q on each type and take the algebraic direct sum of these operators, obtaining T. Invariance gives $T\pi(X) = \pi(X)T$ for each Lie algebra element X. The one-dimensional spherical type is an eigenspace of T, and it generates the irreducible local module (Kazhdan and Patterson, 1984, Theorem I.6.4(c), p. 104). Consequently $T = cI$ for a finite $c \ge 0$, and

$$\|\Theta(u)\|_{L^2}^2 = c \sum_{\ell \ge 0} a_\ell^{-1} \|u_\ell\|_2^2 \le c a_0^{-1} \sum_{\ell \ge 0} (1 + \ell)^2 \|u_\ell\|_2^2.$$

Only finite sums occur at this stage, so no convergence or unbounded-operator assertion is hidden in the comparison. The product argument handles simultaneous variation at the complex places. Differentiation and local Sobolev inequalities then give uniform bounds for every derivative on compact sets, as required to extend the realization continuously to smooth vectors.

This verifies the local replacement relative to the cited classical representation theorems. It does not constitute a new proof of the remaining analytic arguments in Cureton and Thorner (2026). The official text's account of its Lean formalization concerns the official version; it is not

a certificate for the separately posted replacement. The proof of Tang's transfer uses the fixed zero-free conclusion, and its provenance must be kept separate from a check of that conclusion's entire proof.

22 Auxiliary primes for almost all characteristics

This section proves the fixed-degree short search and the uniform selected-menu bound of Theorem 1.12. The averaging variable is the input prime p. Neither conclusion gives a bound for every individual input. All logarithms are natural.

Lemma 22.1. *Assume A. There is an absolute constant C_0 such that, for every Dirichlet character χ modulo $q \geq 1$ and every $x \geq 2$,*

$$\left| \operatorname{Re} \sum_{n \leq x} \Lambda(n)\chi(n) \log(x/n) - \delta_\chi x \right| \leq C_0 \sqrt{x} \log(2q),$$

where $\delta_\chi = 1$ for the principal character and is zero otherwise.

Proof. We include the uniformity argument because the modulus below can be larger than any fixed power of x. For a primitive nonprincipal character of conductor $q \geq 3$, use Lemma 2.2 and the real-part identity in the corrected author version of Lamzouri et al. (2015, corrected author pp. 5–7, Lemmas 2.1–2.3). Their Lemma 2.3, evaluated at the fixed value $x_0 = 100$, gives

$$|\operatorname{Re} B(\chi)| \ll \log(2q).$$

Indeed the denominator in that identity is at least $1 - 2/10 + 1/100 = 81/100$, and its finite prime sum and elementary terms are bounded absolutely. Lemma 2.2 therefore gives the asserted bound: its remaining terms are $O(\log(2q) \log x + (\log x)^2 + 1)$, which are absorbed in $O(\sqrt{x} \log(2q))$. The primitive principal character is the Riemann zeta case, covered by their Lemma 2.1.

If χ is induced by a primitive character χ^*, the two sums differ only at prime powers whose prime divides q. For each such prime v the absolute difference is at most

$$\log v \sum_{1 \leq j \leq \log x / \log v} \log(x/v^j) \leq (\log x)^2.$$

There are at most $\log(2q)/\log 2$ such primes. Finally $(\log x)^2 \ll \sqrt{x}$ uniformly for $x \geq 2$. This proves the assertion, including imprimitive principal characters, with an absolute constant and without a restriction on q. □

Proposition 22.2. *Let r be an odd prime, $B \geq 2$, and put*

$$\mathcal{L} = \{\ell \leq B : \ell \text{ prime}, \ \ell \equiv 1 \pmod r\}, \qquad m = |\mathcal{L}|, \qquad M = \prod_{\ell \in \mathcal{L}} \ell.$$

For an integer n, let $G(n)$ be one if n is a nonzero rth power modulo every $\ell \in \mathcal{L}$, and zero otherwise. The empty product is interpreted as one. Under A, uniformly in $r, B, x \geq 2$,

$$\sum_{n \leq x} \Lambda(n) G(n) \log(x/n) = \frac{x}{r^m} + O\left(\sqrt{x} \log(2M)\right), \tag{29}$$

with an absolute implied constant.

For $P \geq 4$, let $E_r(P, B)$ count primes $p \in (P/2, P]$ with $p \nmid M$ for which no prime $\ell \leq B$ satisfies

$$\ell \equiv 1 \pmod r, \qquad p^{(\ell-1)/r} \not\equiv 1 \pmod \ell.$$

Then

$$E_r(P,B) \ll \frac{P/r^m + \sqrt{P}\log(2M)}{\log P}, \tag{30}$$

again with an absolute implied constant.

For the cubic hard congruence, define $G^-(n) = 1_{n\equiv -1 \pmod 9} G(n)$ *with* $r = 3$*, and put* $Q = 9M$*. The corresponding exact main term is*

$$\sum_{n\le x} \Lambda(n) G^-(n) \log(x/n) = \frac{x}{6\cdot 3^m} + O\big(\sqrt{x}\log(2Q)\big). \tag{31}$$

In particular, for exceptional primes $p \equiv -1 \pmod 9$ *the* P/r^m *term in* (30) *can be replaced by* $P/(6\cdot 3^m)$*. Such primes never divide* Q*.*

Proof. For each $\ell \in \mathcal{L}$, choose a Dirichlet character η_ℓ of exact order r modulo ℓ. Extend every character, including the principal character η_ℓ^0, by zero on nonunits. Orthogonality in the cyclic quotient of order r gives

$$1_{\{n\neq 0 \text{ and } n \text{ an } r\text{th power mod } \ell\}} = \frac{1}{r}\sum_{j=0}^{r-1} \eta_\ell^j(n).$$

Consequently

$$G(n) = r^{-m} \sum_{\boldsymbol{j}\in\{0,\dots,r-1\}^m} \prod_{\ell\in\mathcal{L}} \eta_\ell^{j_\ell}(n).$$

These are distinct characters modulo M, by the Chinese remainder theorem, and exactly one is principal. All coefficients are real and their absolute values sum to one. Apply Lemma 22.1 to each character and take real parts. The total error is bounded by the largest individual error; it is not multiplied by r^m. This proves (29).

Every prime counted by $E_r(P,B)$ has $G(p) = 1$ and contributes at least $\log(P/2)\log 2$ to the nonnegative sum with $x = 2P$. The same sum can contain additional primes and prime powers, all with nonnegative coefficients. Thus no prime-power subtraction is needed, and the first bound follows directly from (29).

For the cubic refinement, expand

$$1_{n\equiv -1 \pmod 9} = \frac{1}{6}\sum_{\psi \bmod 9} \overline{\psi(-1)}\psi(n).$$

Every coefficient $\overline{\psi(-1)}$ is real, equal to 1 or -1. Since $(9, M) = 1$, multiplying the two expansions gives $6\cdot 3^m$ distinct characters modulo Q, exactly one principal, with coefficient $1/(6\cdot 3^m)$. Their coefficient absolute values again sum to one. The same argument proves (31) and the refined exception bound. A prime $p \equiv -1 \pmod 9$ is neither 3 nor any $\ell \equiv 1 \pmod 3$, which proves the final assertion. □

Proof of Theorem 1.11. The prime number theorem in the fixed progression $1 \pmod r$ gives

$$m = \frac{B}{(r-1)\log B}(1+o(1)) = \left(\frac{C_r}{r-1} + o(1)\right)\log P.$$

The strict inequality for C_r implies $m \ge \log P/(2\log r)$ for all sufficiently large P, so $r^m \ge \sqrt{P}$. The elementary Chebyshev bound gives $\log M \le \sum_{\ell\le B}\log\ell \ll B$. Also $B < P/2$ eventually, so no prime in the interval divides M. Substitution into (30) proves the asserted bound. For $r = 3$, $1 > 1/\log 3$.

Enumerating primes at most B and testing $p^{(\ell-1)/r} \not\equiv 1 \pmod \ell$ use a number of bit operations polynomial in B, r and $\log p$. A successful test is a certificate independent of A. The auxiliary-prime reduction then uses an explicitly represented cyclotomic algebra of dimension $\ell - 1$, so its established running time is polynomial in these parameters. Since $B = O_{r,C_r}(\log p \log\log p)$ on the dyadic interval, this gives the stated algorithmic conclusion. □

Corollary 22.3. *There is a sufficiently large absolute constant C such that, under A, for every odd prime r and every $P \geq 4$, setting*

$$L = \log(Pr), \qquad B = Cr^2L^2$$

gives $m \geq \log P/\log r$. The number of hard primes $p \in (P/2, P]$ having no good auxiliary prime at most B is

$$O\left(\frac{\sqrt{P}\, r^2 \log^2(Pr)}{\log P}\right),$$

with an absolute implied constant. Relative to the number of all prime inputs in the interval, this is negligible in ranges where $r^2 \log^2(Pr) = o(\sqrt{P})$. No uniform relative-density assertion within the thinner set of hard inputs follows from this coarse estimate alone. Theorem 1.12 gives a stronger bound and a separate hard-input denominator. On every successful input the auxiliary-prime search and field construction take time polynomial in r and $\log p$, uniformly in these parameters.

Proof. We first obtain a uniform progression count with a positive weight. Let

$$S(y; r, 1) = \sum_{\substack{n \leq y \\ n \equiv 1 \bmod r}} \Lambda(n) \log(y/n).$$

Character orthogonality and Lemma 22.1 give $S(y; r, 1) = y/(r-1) + O(\sqrt{y}\log(2r))$. Form

$$T(B) = S(B; r, 1) - 2S(B/2; r, 1) + S(B/4; r, 1).$$

The coefficient of $\Lambda(n)$ is zero for $n \leq B/4$, equals $\log(4n/B)$ for $B/4 < n \leq B/2$, and equals $\log(B/n)$ for $B/2 < n \leq B$. It is therefore nonnegative and at most $\log 2$. Also

$$T(B) = \frac{B}{4(r-1)} + O(\sqrt{B}\log(2r)).$$

The total contribution of proper prime powers is $O(\sqrt{B})$: Chebyshev's bound gives $\sum_{k \geq 2} \vartheta(B^{1/k}) \ll \sqrt{B} + B^{1/3}\log B \ll \sqrt{B}$. For an absolute C_1, if $B \geq C_1 r^2 \log^2(2r)$, it follows that the prime part of $T(B)$ is at least $B/[8(r-1)]$. As each such prime contributes at most $(\log 2)\log B$,

$$m \geq \frac{B}{8(r-1)(\log 2)\log B}.$$

Now take $B = Cr^2L^2$ with $L = \log(Pr)$. The preceding size condition holds for all $P \geq 4$, $r \geq 3$, when the absolute C is sufficiently large. Moreover $\log B \leq c_C L$, where $c_C = O(1 + \log C)$. Thus, after increasing the same absolute C if needed,

$$m \geq \frac{CrL}{8(\log 2)c_C} \geq \frac{\log P}{\log r}.$$

The elementary bound $\log M \ll B$ and (30) now prove the conclusion. Every hard prime has $\mathrm{ord}_r(p) > 1$, whereas every prime factor of M is $1 \pmod r$. Thus the exclusion $p \nmid M$ causes no loss, even when $B \geq P/2$. □

Proof of Theorem 1.12. The progression count in Corollary 22.3 gives at least $\log P/\log r$ available primes after choosing C sufficiently large. Since their number is an integer, this supplies the stated k primes for every $P \geq 4$ and $r \geq 3$. Fix any such set $\mathcal{S}$, and put

$$M_{\mathcal{S}} = \prod_{\ell \in \mathcal{S}} \ell.$$

The proof of Proposition 22.2 applies to this subset: the product of the k subgroup indicators has r^k character coefficients, all equal to r^{-k}, with exactly one principal coefficient. Their absolute values sum to one. Denote this product indicator by $G_{\mathcal{S}}$, with value zero on nonunits. Then Lemma 22.1, taking real parts, gives

$$\sum_{n\le 2P} \Lambda(n)G_{\mathcal{S}}(n)\log(2P/n) = \frac{2P}{r^k} + O\Big(\sqrt{P}\log(2M_{\mathcal{S}})\Big).$$

All summands are nonnegative. An exceptional prime not dividing $M_{\mathcal{S}}$ contributes at least $\log(P/2)\log 2 \ge \frac{1}{2}\log P\log 2$. There are at most k input primes dividing $M_{\mathcal{S}}$. Consequently their total exceptional count $E_{r,\mathcal{S}}(P)$ satisfies

$$E_{r,\mathcal{S}}(P) \ll k + \frac{P/r^k + \sqrt{P}\log(2M_{\mathcal{S}})}{\log P}. \tag{32}$$

Prime powers need not be subtracted from this upper bound.

Write $t = \log P$, $u = \log r$, and $L = t + u$. The choice of k gives $r^k \ge \sqrt{P}$ and $k \le t/(2u) + 1$. Also

$$\log(2M_{\mathcal{S}}) \le \log 2 + k\log B, \qquad \log B = \log C + 2u + 2\log L.$$

Since $t, u > 1$, division by t yields

$$\frac{\log(2M_{\mathcal{S}})}{t} \ll 1 + \frac{u}{t} + \frac{\log L}{u}.$$

For example the additional term $\log L/t$ is bounded by $1 + u/t$, because $\log L \le L$. Furthermore $k \le t/2 + 1 \le e^{t/2} = \sqrt{P}$. Substitution in (32) proves the asserted absolute bound. Failure of the full menu implies failure of every selected subset, proving the full-menu assertion. If $r \le P^{A_0}$, then

$$A_r(P) \ll_{A_0} 1 + \log\log P.$$

Dividing by the $\asymp P/\log P$ primes in the interval proves the asserted uniform density among all prime inputs.

The thinner hard subclass requires its own denominator. For $p \ne r$, the quotient $(r-1)/d$ is prime to r; factoring a difference of powers gives

$$v_r(p^{r-1} - 1) = v_r(p^d - 1).$$

Let T_r be the order-$(r-1)$ subgroup of $(\mathbb{Z}/r^2\mathbb{Z})^\times$. Reduction modulo r maps it bijectively to $(\mathbb{Z}/r\mathbb{Z})^\times$. Thus the hard residues are exactly $T_r \setminus \{1\}$, with density

$$\delta_r = \frac{r-2}{r(r-1)} \ge \frac{1}{2r}.$$

Their indicator h_r, extended by zero on multiples of r, is the subgroup indicator minus the singleton indicator. Its Dirichlet coefficients are real and have total absolute value at most two, independently of r.

For a dyadic lower bound use the fixed nonnegative weight

$$w(y) = \begin{cases} \log(2y), & 1/2 \le y \le 1/\sqrt{2}, \\ \log(1/y), & 1/\sqrt{2} \le y \le 1, \\ 0, & \text{otherwise.} \end{cases}$$

If $F_\chi(x) = \sum_{n\le x} \Lambda(n)\chi(n)\log(x/n)$, then

$$\sum_n \Lambda(n)\chi(n)w(n/P) = F_\chi(P) - 2F_\chi(P/\sqrt{2}) + F_\chi(P/2).$$

The integral $c_w = \int w(y)\,dy = (1 - 1/\sqrt{2})^2$ is positive. Lemma 22.1 and the bounded coefficient sum therefore give

$$\sum_n \Lambda(n) h_r(n) w(n/P) = c_w \delta_r P + O\Big(\sqrt{P}\log(2r)\Big).$$

The proper prime powers contribute $O(\sqrt{P})$, as in the preceding progression count. Each prime contributes at most $\|w\|_\infty \log P$. Hence

$$H_r(P) \geq \frac{c_w \delta_r P - C'\sqrt{P}\log(2r)}{\|w\|_\infty \log P}$$

for an absolute C'. Increasing C_1 makes the numerator at least $c_w P/(4r)$ whenever $P \geq C_1 r^2 \log^2(2r)$, giving the stated lower bound. The hard exceptions are a subset of all exceptions, so their proportion is

$$O\Big(\frac{r\log P}{\sqrt{P}} A_r(P)\Big). \tag{33}$$

For the precise uniform limit, let $\epsilon(P) \to 0$ be positive. Every $r \leq \epsilon(P)\sqrt{P}/\log P$ eventually satisfies the denominator condition, since

$$\frac{C_1 r^2 \log^2(2r)}{P} \leq C_1 \epsilon(P)^2 \left(\frac{1}{2} + \frac{\log 2}{\log P}\right)^2 = o(1).$$

If $r \leq P^{1/4}$, the expression in (33) is $O(P^{-1/4}\log P(1 + \log\log P))$, uniformly. For $P^{1/4} < r \leq \epsilon(P)\sqrt{P}/\log P$, the factor $A_r(P)$ is bounded and $r\log P/\sqrt{P} \leq \epsilon(P)$. Both bounds tend to zero. For $r = 3$, the hard class is $p \equiv -1 \pmod 9$, so this includes that earlier congruence.

Finally, let C_0 be a constant sufficient for the progression count and choose the theorem's $C \geq 9C_0$. For an integer or dyadic scale put

$$N_0 = \lceil \log_2 P \rceil + \lceil \log_2 r \rceil, \qquad T = C_0 r^2 N_0^2,$$

increasing the fixed C_0 to an integer if necessary. Since $\log(Pr) \leq N_0 < 3\log(Pr)$, a sieve or trial division up to the integer $T \leq B$ finds the required menu. The integer k is found by comparing r^{2k} with P. Modular powering tests each prime; a successful one is an unconditional certificate for the algebraic construction in Lemma 18.2. All these costs are polynomial in $r, \log P$, with absolute exponents. The r^k characters are a proof device and are never enumerated by the algorithm. □

The sieve argument averages fixed rational Dirichlet characters over the varying input prime. It proves neither a pointwise upper bound for the least good auxiliary prime nor the required positive nonabelian prime-detection gap for an arbitrarily specified p.

23 Average constructions from explicit cubic polynomials

The character averages become an explicit field-construction search when the cyclic cubic fields are indexed by a small integer parameter.

23.1 Logarithmic normalized Tate parameters for almost all inputs

The quadratic norms in the normalized Tate family provide linearly many independent rational cubic characters. This improves the parameter range in the average construction to a constant times $\log P$. The product estimate is the classical quadratic-value mechanism; compare Cilleruelo (2011, Lemmas 1, 3, and 4 and §2.4, pp. 1131–1134). We give its specialization in full to track the private conductor primes and the leading constant.

Lemma 23.1 (Large primes in the quadratic norms)**.** *For an integer $K \geq 4$, put*

$$N(k) = k^2 + 3k + 9, \qquad \mathcal{S}_K = \{q > K+3 : q \text{ prime},\ q \mid N(k) \text{ for some } 1 \leq k \leq K\}.$$

Then

$$|\mathcal{S}_K| \geq \left(\frac{1}{2} - o(1)\right) K \qquad (K \longrightarrow \infty). \tag{34}$$

Each $N(k)$ has at most one prime divisor exceeding $K+3$, and that divisor has valuation one. A prime $q \in \mathcal{S}_K$ divides at most two of the values; if it divides two, then $q \leq 2K+3$.

Proof. Stirling's formula gives

$$\sum_{k=1}^{K} \log N(k) = 2\log K! + \sum_{k=1}^{K} \log(1 + 3/k + 9/k^2) = 2K\log K - 2K + O(\log K). \tag{35}$$

For a prime $q \neq 3$, the congruence $N(k) \equiv 0 \pmod q$ has two simple roots if $q \equiv 1 \pmod 3$, and no roots otherwise. Indeed, division by nine transforms it into $u^2 + u + 1 = 0$; its roots are the nonidentity cube roots of unity. For $q = 2$ there are no roots. Hensel lifting therefore gives exactly two roots modulo every q^a in the split case.

Set $L_q = \lfloor \log N(K)/\log q \rfloor$. For $q \equiv 1 \pmod 3$,

$$\sum_{k=1}^{K} v_q(N(k)) \leq 2 \sum_{a=1}^{L_q} \left(\frac{K}{q^a} + 1\right) \leq \frac{2K}{q-1} + 2L_q. \tag{36}$$

The accumulated floor error after multiplication by $\log q$ and summation over $q \leq K+3$ is at most

$$2\pi(K+3)\log N(K) = O(K).$$

The prime number theorem in the fixed progression $1 \pmod 3$, followed by partial summation, yields

$$\sum_{\substack{q \leq x \\ q \equiv 1 \bmod 3}} \frac{\log q}{q-1} = \frac{1}{2}\log x + o(\log x).$$

The prime three contributes only $O(K)$: if $3 \nmid k$ its valuation is zero, while for $k = 3u$ one has $N(k) = 9(u^2+u+1)$ and $v_3(u^2+u+1) \leq 1$. Thus the total logarithmic contribution from primes at most $K+3$ is at most $K\log K + o(K\log K)$. Subtracting this from (35) gives

$$\sum_{q > K+3} \log q \sum_{k=1}^{K} v_q(N(k)) \geq K\log K + o(K\log K). \tag{37}$$

Since

$$N(k) \leq K^2 + 3K + 9 < (K+3)^2,$$

no $N(k)$ contains two prime factors greater than $K+3$, and every such factor occurs once. A prime $q > K+3$ divides at most two values, because the parameters are distinct modulo q and the quadratic has two roots. If it divides $N(a)$ and $N(b)$ with $a \neq b$, then $a + b \equiv -3 \pmod q$. The bounds $1 \leq a, b \leq K < q - 3$ force $a + b = q - 3$, hence $q \leq 2K + 3$.

A prime occurring once consequently contributes at most $\log N(K) < 2\log(K+3)$ to (37); one occurring twice contributes at most $2\log(2K+3)$. Therefore

$$2|\mathcal{S}_K|\log(2K+3) \geq K\log K + o(K\log K),$$

which proves (34) and the remaining assertions. □

Lemma 23.2 (Cyclic cubic characters and ramification)**.** *For every positive integer k, there is a primitive cubic Dirichlet character θ_k of conductor $f_k \mid N(k)$ such that, for every prime $p \nmid N(k)$,*

$$h_k(Y) := Y^3 - (k+3)Y^2 + kY + 1 \text{ is irreducible over } \mathbb{F}_p \quad \Longleftrightarrow \quad \theta_k(p) \neq 1. \tag{38}$$

If $q \neq 3$ and $v_q(N(k)) = 1$, then $q \mid f_k$.

Proof. The rational-root test proves that h_k is irreducible over $\mathbb{Q}$, since

$$h_k(1) = -1, \qquad h_k(-1) = -2k - 3.$$

Its discriminant is $N(k)^2$, so a root generates a cyclic cubic splitting field $L_k/\mathbb{Q}$. The Kronecker–Weber theorem and the conductor–discriminant formula give a primitive cubic Dirichlet character of conductor f_k, with

$$\operatorname{disc}(L_k) = f_k^2;$$

see Silvester et al. (2006, p. 472) for these cyclic-cubic facts. If α_k is a root and $I_k = [\mathcal{O}_{L_k} : \mathbb{Z}[\alpha_k]]$, the polynomial discriminant identity gives

$$N(k)^2 = f_k^2 I_k^2,$$

and hence $f_k \mid N(k)$. No equality between f_k and $N(k)$ is assumed. For $p \nmid N(k)$, Frobenius has the cycle type of the factorization of h_k modulo p. The only possibilities in the cyclic group of order three are the identity and a three-cycle. A nontrivial character is faithful on that group, proving (38).

Finally, direct expansion gives

$$h_k\!\left(Z + \frac{k+3}{3}\right) = Z^3 - \frac{N(k)}{3}Z - \frac{(2k+3)N(k)}{27}. \tag{39}$$

If $q \neq 3$ and $v_q(N(k)) = 1$, the identity $(2k+3)^2 = 4N(k) - 27$ implies $q \nmid 2k+3$. All denominators in (39) are q-adic units, and the translated polynomial is Eisenstein over $\mathbb{Q}_q$. Thus q is totally ramified in L_k, so $q \mid \operatorname{disc}(L_k)$ and $q \mid f_k$. □

Corollary 23.3 (A linear-size independent parameter set)**.** *For every sufficiently large K, there is a set $\mathcal{K}_K \subset \{1, \dots, K\} \setminus \{3\}$ of size at least $(1/2 - o(1))K$ whose associated characters θ_k are linearly independent over $\mathbb{F}_3$. Each such character has a conductor prime greater than $K+3$ which divides none of the other selected conductors. For any m selected parameters, their conductors satisfy*

$$\log \prod_{i=1}^{m} f_{k_i} \leq m \log N(K) < 2m \log(K+3).$$

Proof. For each $q \in \mathcal{S}_K$, choose one $k(q) \leq K$ with $q \mid N(k(q))$. These parameters are distinct, because a value has at most one prime divisor greater than $K+3$. The valuation at q is one, so Lemma 23.2 gives $q \mid f_{k(q)}$. For a different selected parameter $k(q')$, the value $N(k(q'))$ contains $q' > K+3$ and therefore cannot also contain q. Consequently $q \nmid f_{k(q')}$. Restriction to the q-component of the common modulus forces the exponent of $\theta_{k(q)}$ in any principal character product to be zero modulo three. Doing this for every selected q proves independence. The parameter three cannot occur, since $N(3) = 27$ has no prime factor exceeding $K+3$. The size assertion follows from Lemma 23.1, and the conductor bound follows from $f_k \mid N(k) \leq N(K)$. □

Proof of Theorem 1.13. Put $K = \lfloor C \log P \rfloor$ with $C > 1/\log 3$ and let

$$m = \left\lceil \frac{\log P}{2 \log 3} \right\rceil .$$

By Corollary 23.3, there are at least $(C/2 - o(1)) \log P$ independent characters from parameters at most K. The strict condition on C therefore allows us to select exactly m of them for all sufficiently large P. Write them as $\theta_{k_1}, \dots, \theta_{k_m}$ and put

$$Q = \prod_{i=1}^{m} f_{k_i}.$$

Then

$$\log Q \le m \log N(K) = O_C(\log P \log\log P), \qquad 3^m \ge \sqrt{P}. \tag{40}$$

Also $N(K) < P/2$ eventually, so primes in the input interval divide none of the polynomial discriminants or conductors.

Regard the characters as characters modulo Q, extended by zero on nonunits, including each exponent-zero character. The indicator $G(n)$ of common triviality is

$$G(n) = \frac{1}{3^m} \sum_{\boldsymbol{e} \in \{0,1,2\}^m} \prod_{i=1}^{m} \theta_{k_i}(n)^{e_i}.$$

Here $G(n) = 0$ when $(n, Q) > 1$. Independence implies that the displayed sum contains 3^m distinct characters, exactly one principal. Its coefficients are real and have absolute values summing to one. Lemma 22.1, applied termwise and followed by taking real parts, consequently gives

$$\sum_{n \le x} \Lambda(n) G(n) \log(x/n) = \frac{x}{3^m} + O\big(\sqrt{x} \log(2Q)\big) .$$

If every admissible polynomial fails at a prime $p \in (P/2, P]$, then all the selected polynomials fail and (38) gives $G(p) = 1$. At $x = 2P$ each such prime contributes at least $\log(P/2) \log 2$ to this nonnegative sum. Thus the number of exceptional primes is

$$O\left(\frac{P/3^m + \sqrt{P} \log(2Q)}{\log P} \right) = O_C(\sqrt{P} \log\log P)$$

by (40). This counts all primes in the interval, so it also bounds the exceptions among $p \equiv -1 \pmod 9$.

Every selected parameter is different from three. Since $K < p$ eventually, it also satisfies $k \not\equiv 3 \pmod p$. The deterministic algorithm simply scans all $1 \le k \le K$, skipping that residue, and tests

$$\gcd\big(h_k(Y), Y^p - Y\big) = 1.$$

For a cubic this is equivalent to irreducibility, since a reducible cubic has a linear factor. Repeated squaring computes $Y^p \bmod h_k$ in polynomial bit time in $\log p$. There are $O_C(\log p)$ tests. A successful test gives the presentation $\mathbb{F}_p[Y]/(h_k)$ of $\mathbb{F}_{p^3}$; its correctness is independent of A. The independent character set and its conductors are used only in the proof and need not be computed by the algorithm. □

The conclusion is an average over the varying input prime. It retains the displayed exceptional set and supplies no pointwise bound for a prescribed exceptional prime.

Part II
Routes that do not give a uniform construction

The results in this part identify additional inputs, exact reformulations, and limits of specific constructions. These constructions isolate particular missing inputs for a uniform algorithm.

24 Analytic windows and the unresolved zero contribution

The explicit formulas below isolate the finite-scale estimate that would make the classical auxiliary-prime search uniform.

24.1 An averaged window estimate and a limitation of positive tests

The one-sided condition in Proposition 18.6 can be forced by an unconditional averaged estimate, but the resulting range is a power of the discriminant. We give the estimate and isolate the conductor term responsible for that scale. The limitation proved below concerns a specified use of positive transforms; it does not exclude other deductions from A.

Let J be a number field of degree n_J, with $r_{1,J}$ real embeddings and absolute discriminant D_J. For an even real compactly supported test function f, define

$$\begin{aligned}
\Phi_f(s) &= \int_{\mathbb{R}} f(t) e^{(s-1/2)t}\, dt, \\
T_J(f) &= \sum_{\mathfrak{p}} \sum_{k\ge 1} \frac{\log N\mathfrak{p}}{(N\mathfrak{p})^{k/2}} f(k \log N\mathfrak{p}), \\
Z_J(f) &= \lim_{T\to\infty} \sum_{|\operatorname{Im}\rho_J|<T} \Phi_f(\rho_J).
\end{aligned}$$

The zeros in Z_J are nontrivial zeros, counted with multiplicity. We take f to satisfy the admissibility conditions in Belabas et al. (2008, author pp. 4–5, equations (3)–(4)): f is continuous, $f(t)e^{(1/2+\eta)t}$ has bounded variation and is integrable on $[0,\infty)$ for some $\eta > 0$, and $(f(0) - f(t))/t$ has bounded variation there. Their unconditional explicit formula, specialized to the trivial character, is

$$2T_J(f) = 2\Phi_f(1) + f(0)\log D_J + \mathcal{A}_J(f) - Z_J(f), \tag{41}$$

where, with γ_E denoting Euler's constant,

$$\begin{aligned}
\mathcal{A}_J(f) = f(0)\left(-n_J\gamma_E - n_J\log(8\pi) - \frac{r_{1,J}\pi}{2}\right) \\
+ r_{1,J}\int_0^\infty \frac{f(0)-f(t)}{2\cosh(t/2)}\, dt + n_J\int_0^\infty \frac{f(0)-f(t)}{2\sinh(t/2)}\, dt.
\end{aligned}$$

The factor 2 in the pole term is important: for an even test, $\Phi_f(0) = \Phi_f(1)$, so the two poles contribute $\Phi_f(0) + \Phi_f(1) = 2\Phi_f(1)$.

Lemma 24.1 (A compact test positive throughout the strip)**.** *For $S \geq 1$, put*

$$h_S(t) = \left(1 - \frac{|t|}{S}\right)_+, \qquad f_S(t) = e^{-|t|/2} h_S(t).$$

Then $\operatorname{Re} \Phi_{f_S}(\beta + i\gamma) \geq 0$ *for every* $0 \leq \beta \leq 1$ *and* $\gamma \in \mathbb{R}$*. Moreover,*

$$\Phi_{f_S}(1) = \frac{S}{2} + 1 - \frac{1 - e^{-S}}{S}. \tag{42}$$

Proof. Use the Fourier convention $\widehat{g}(\gamma) = \int_{\mathbb{R}} g(t) e^{i\gamma t}\, dt$. The real part in question is the Fourier transform of

$$f_S(t) \cosh((\beta - 1/2)t) = \frac{h_S(t)}{2} \left(e^{-\beta|t|} + e^{-(1-\beta)|t|}\right).$$

The two functions

$$\widehat{h}_S(\gamma) = \frac{2(1 - \cos(S\gamma))}{S\gamma^2}, \qquad \widehat{e^{-a|t|}}(\gamma) = \frac{2a}{a^2 + \gamma^2} \quad (a > 0)$$

are nonnegative, with the first expression interpreted continuously at zero. A product of these original functions has a Fourier transform equal to their convolution divided by 2π, and is therefore nonnegative. The cases $\beta = 0, 1$ follow by continuity. Finally,

$$\Phi_{f_S}(1) = \int_0^S (1 + e^{-t}) \left(1 - \frac{t}{S}\right) dt,$$

which gives (42). The test is admissible in (41). It is continuous and piecewise smooth, and its transform is $O_S((1 + |\gamma|)^{-2})$ uniformly in the closed strip. Thus its zero sum converges absolutely by the usual unit-height zero count $O(\log D_J + n_J \log(|\gamma| + 3))$. □

Proposition 24.2 (An averaged one-sided estimate)**.** *Fix the nonzero nonnegative $w \in C_c^\infty((1,2))$ of Proposition 18.6 and put $c_w = \int_1^2 w(v)\, dv$. For every number field J, with degree n_J and absolute discriminant D_J, and every $S \geq 1$,*

$$I_J(S) := \int_1^{e^S} \frac{\Psi_J(x)}{x^2} \left(1 - \frac{\log x}{S}\right) dx \leq \frac{c_w S}{2} + \frac{c_w}{4} \log D_J + O_w(n_J). \tag{43}$$

The estimate is unconditional, with the implied constant independent of J and S.

Proof. Write $h_S(t) = (1 - |t|/S)_+$. Use the test of Belabas et al. (2008, §5, equations (17)–(22)),

$$g_S(t) = \frac{h_S(t)}{\cosh(t/2)}.$$

With the conventions above, $g_S(0) = 1$ and $\Phi_{g_S}(1) = S$. On the line $\Re s = 1$, the real part of its transform equals the nonnegative Fourier transform of h_S. Evenness gives the same conclusion on $\Re s = 0$. The transform tends to zero uniformly in this closed strip as $|\Im s| \to \infty$, by integration by parts on the two smooth pieces of the support. The minimum principle therefore gives $\Re \Phi_{g_S} \geq 0$ in the strip. A second integration by parts, including the one-sided derivatives at the break points, gives $O_S((1 + |\Im s|)^{-2})$, so the zero sum converges absolutely.

Let Λ_J be the ideal von Mangoldt function. In equation (20) of the cited paper the pole term for this test is $2S$ and the ideal sum has coefficient 4. The remaining archimedean terms are $O(n_J)$, uniformly in $S \geq 1$. Indeed, the terms involving $1 - h_S(t)$ are bounded by using $1 - h_S(t) \leq t/S$ and the integrability of $t/\sinh t$ and $t/(2\cosh^2(t/2))$ on the positive half-line. Dropping the nonnegative zero sum gives

$$4 \sum_{\mathfrak{a}} \frac{\Lambda_J(\mathfrak{a})}{1 + N\mathfrak{a}} h_S(\log N\mathfrak{a}) \leq 2S + \log D_J + O(n_J).$$

At each rational prime power the coefficient of $-\zeta_J'/\zeta_J$ is at most n_J times its rational von Mangoldt coefficient. In the sum before multiplication by 4, changing $1/(1+N\mathfrak{a})$ to $1/N\mathfrak{a}$ costs at most $n_J \sum_{m\ge 2} \Lambda(m)/m^2 = O(n_J)$. Hence

$$\sum_{\mathfrak{a}} \frac{\Lambda_J(\mathfrak{a})}{N\mathfrak{a}} h_S(\log N\mathfrak{a}) \le \frac{S}{2} + \frac{1}{4}\log D_J + O(n_J). \tag{44}$$

Put $X = e^S$ and $c_{w,1} = \int_1^2 w(v)\log v\,dv$. For each ideal prime-power norm $2 \le m \le X$, substituting $v = m/x$ in the smoothed integral gives

$$\int_1^X \frac{w(m/x)}{x^2}\left(1 - \frac{\log x}{S}\right)dx = \frac{1}{m}\left(c_w h_S(\log m) + \frac{c_{w,1}}{S}\right).$$

Rational Chebyshev bounds and partial summation give

$$\sum_{N\mathfrak{a}\le X} \frac{\Lambda_J(\mathfrak{a})}{N\mathfrak{a}} \ll n_J(1+\log X), \qquad \sum_{X<N\mathfrak{a}<2X} \frac{\Lambda_J(\mathfrak{a})}{N\mathfrak{a}} \ll n_J$$

with absolute constants. For $X < m < 2X$, the weight after substitution is at most $(\log 2)/S$, and no larger norm contributes. Consequently $I_J(S)$ equals c_w times the sum in (44) plus $O_w(n_J)$. This proves (43). □

Let $F = \mathbb{Q}(p^{1/r})$ for distinct primes p, r with r odd. If $0 < \varepsilon < 1$, $1 \le u_0 \le S$, and $\Psi_F(x) \ge (2-\varepsilon)c_w x$ throughout $[e^{u_0}, e^S]$, then

$$I_F(S) \ge \frac{(2-\varepsilon)c_w}{2S}(S-u_0)^2.$$

Since $n_F = r$, comparison with (43) gives the necessary inequality

$$(1-\varepsilon)S \le \tfrac{1}{2}\log D_F + 2(2-\varepsilon)u_0 + O_w(r). \tag{45}$$

Violation of this inequality forces a strict one-sided window.

Under A, direct integration of the error term gives an auxiliary prime satisfying

$$\ell \le D_F^{1/2}\exp(C_w r) \tag{46}$$

for a constant C_w independent of p, r. Indeed, suppose there is no auxiliary prime up to $2e^S$. For $2 \le x \le e^S$, Proposition 18.4 gives $\mathcal{H}(x) \ll_w r\sqrt{x}$. The smoothed estimates for $\mathbb{Q}$ and $K = \mathbb{Q}(\zeta_r)$ under A give

$$\Psi_F(x) = \Psi_K(x) + \Psi_{\mathbb{Q}}(x) - \mathcal{H}(x) \ge 2c_w x - C_w' r\log(2r)\sqrt{x}.$$

Put $x_0 = (r\log(2r))^2$ and $u_0 = \log x_0$. For $S \ge u_0$, restrict the nonnegative integral $I_F(S)$ to $x_0 \le x \le e^S$. The integrated error is at most $C_w' r\log(2r)\int_{x_0}^{\infty} x^{-3/2}dx = 2C_w'$. The integrated main term is $c_w(S - 2u_0 + u_0^2/S)$. It follows that

$$S \le \tfrac{1}{2}\log D_F + 4u_0 + O_w(r).$$

Here $u_0 = 2\log r + 2\log\log(2r) = O(\log(2r))$. Choosing $S = \frac{1}{2}\log D_F + C_w'' r$, with C_w'' large enough, gives $S \ge u_0$ and a contradiction. The factor 2 from the endpoint $2e^S$ is absorbed in C_w. The scale of this argument is a power of the absolute field discriminant D_F.

The next estimate uses the classical interval case of Turán's extremal problem for positive definite functions; see Kolountzakis and Révész (2006, author version, Theorem 6, equation (29), p. 14, and Remark 4, p. 15). We include a periodization proof to fix the normalization and the support convention.

Lemma 24.3 (Turán bound for tests positive throughout the strip)**.** *Let f be a continuous real even function supported on $[-S,S]$. Suppose*

$$\operatorname{Re}\Phi_f(\beta+i\gamma)\geq 0 \quad (0<\beta<1,\ \gamma\in\mathbb{R}).$$

Then $f(0)\geq 0$, and

$$0\leq \Phi_f(1)\leq Sf(0), \qquad \Phi_f(0)+\Phi_f(1)\leq 2Sf(0). \tag{47}$$

If $f(0)=0$, then f is identically zero.

Proof. Continuity in β gives positivity also at $\beta=1$. The function $g(t)=f(t)\cosh(t/2)$ therefore has nonnegative Fourier transform. Here no smoothness assumption or informal Fourier inversion is needed: Gaussian regularization gives

$$\frac{1}{2\pi}\int_{\mathbb{R}}\widehat{g}(\gamma)e^{-a\gamma^2}\,d\gamma\longrightarrow g(0) \qquad (a\downarrow 0).$$

Monotone convergence shows that $\widehat{g}$ is integrable, with integral $2\pi g(0)$. Fourier inversion now implies

$$|g(t)|\leq\frac{1}{2\pi}\int_{\mathbb{R}}\widehat{g}(\gamma)\,d\gamma=g(0)=f(0).$$

The integral of g is $\Phi_f(1)=\widehat{g}(0)\geq 0$. To obtain the stronger interval bound, fix $T>S$ and periodize:

$$G_T(t)=\sum_{k\in\mathbb{Z}}g(t+kT).$$

This is a continuous T-periodic function, and $G_T(0)=g(0)$ because g is supported on $[-S,S]$. Its Fourier coefficients are

$$c_n=\frac{1}{T}\widehat{g}(2\pi n/T)\geq 0.$$

The Fejér sums at zero therefore satisfy

$$\sum_{|n|\leq N}\left(1-\frac{|n|}{N+1}\right)c_n\geq c_0=\frac{\widehat{g}(0)}{T}.$$

By Fejér's theorem, the left side converges to $G_T(0)=g(0)$. Hence $\Phi_f(1)\leq Tf(0)$. Letting $T\downarrow S$ proves the first bound. Evenness gives $\Phi_f(0)=\Phi_f(1)$ and the second. If $f(0)=0$, the preceding pointwise bound forces $g=f=0$. □

The test g_S in the proof of Proposition 24.2 attains the bound: $g_S(0)=1$ and $\Phi_{g_S}(1)=S$.

The exact conductor for the detector is smaller than D_F, but has the same dependence on p. Put

$$K=\mathbb{Q}(\zeta_r), \qquad L(s)=\frac{\zeta_F(s)}{\zeta_{\mathbb{Q}}(s)}, \qquad Q_\chi=\frac{D_F}{D_K}.$$

The Artin induction identities identify $L(s)$ with a nontrivial finite-order Hecke L-function over K, of finite conductor $\mathfrak{f}_\chi$, and give $D_K N\mathfrak{f}_\chi=D_F$. Thus $Q_\chi=N\mathfrak{f}_\chi$ is a positive integer. For $T_{\mathcal{H}}=T_K+T_{\mathbb{Q}}-T_F$, subtracting the three explicit formulas gives the exact identity

$$\boxed{2T_{\mathcal{H}}(f)=2\Phi_f(1)-f(0)\log Q_\chi-Z_K(f)+Z_L(f).} \tag{48}$$

Here Z_L is the zero sum of the entire nontrivial Hecke function L, so $Z_F=Z_{\mathbb{Q}}+Z_L$. There is no archimedean remainder in (48):

$$n_K+n_{\mathbb{Q}}-n_F=(r-1)+1-r=0, \qquad r_{1,K}+r_{1,\mathbb{Q}}-r_{1,F}=0+1-1=0.$$

Proposition 24.4 (A limitation of discarding the unknown zero sum)**.** *Let f be a nonzero admissible test satisfying the hypotheses of Lemma 24.3. The lower bound obtained from* (48) *by using only $Z_L(f) \geq 0$ cannot be strictly positive if*

$$S \leq \frac{1}{2} \log Q_\chi.$$

Proof. Both $Z_L(f)$ and $Z_K(f)$ are nonnegative by the transform hypothesis. The stated use of positivity gives only

$$2T_{\mathcal{H}}(f) \geq 2\Phi_f(1) - f(0) \log Q_\chi - Z_K(f).$$

The expression on the right is at most $f(0)(2S - \log Q_\chi)$ by Lemma 24.3, and $f(0) > 0$. It therefore cannot certify a positive value in the stated range, even if the abelian zero sum $Z_K(f)$ is known exactly. □

Only p and r ramify in F, and p is tamely totally ramified, so $v_p(D_F) = r - 1$. Since $D_F \leq r^r p^{r-1}$ and $D_K = r^{r-2}$, the integral conductor Q_χ satisfies

$$p^{r-1} \leq Q_\chi \leq r^2 p^{r-1}, \qquad (r-1)\log p \leq \log Q_\chi \leq (r-1)\log p + 2\log r. \tag{49}$$

In the difficult scalar case, $s = \mathrm{ord}_r(p)$ and $v_r(p^s - 1) \geq 2$, one has the exact identity $Q_\chi = p^{r-1}$. Indeed, then $p^{r-1} \equiv 1 \pmod{r^2}$; the decomposition $\mathbb{Z}_r^\times = \mu_{r-1} \times (1 + r\mathbb{Z}_r)$ and the r-adic logarithm show that p is an rth power in $\mathbb{Q}_r$. Since K contains μ_r, the extension $K(p^{1/r})/K$ splits completely at its prime above r. The character conductor is consequently supported only above p, and its norm has the already determined p-adic exponent $r - 1$.

For a search bounded by a fixed polynomial in r and $\log p$, the logarithmic support is $S = O(\log r + \log\log p)$. Equations (47) and (49) explain why the strategy in Proposition 24.4 cannot supply that range uniformly. This is already visible with r fixed. It is a limitation of discarding the nonabelian zero sum, not an impossibility assertion about the one-sided window condition or A.

For comparison, suppose transform positivity is required only for $\delta \leq \beta \leq 1 - \delta$, where $0 < \delta \leq 1/2$. The same argument applied to $g_\delta(t) = f(t)\cosh((1/2 - \delta)t)$ gives

$$|\Phi_f(1)| \leq 4f(0)\frac{e^{\delta S} - 1}{\delta}.$$

Indeed, $|g_\delta(t)| \leq f(0)$ and $\cosh(t/2)/\cosh((1/2 - \delta)t) \leq 2e^{\delta t}$ for $t \geq 0$. With fixed $\delta > 0$, this upper bound permits a pole response large enough at S proportional to $\log\log Q_\chi$; it does not construct such a test. If instead $\delta \asymp 1/\log Q_\chi$ and $S = O(\log\log Q_\chi)$, the bound remains $O(Sf(0))$. This comparison does not cover a height-dependent positivity region with a separately estimated tail, nor a quantitative lower bound for $Z_L(f)$, nor a sign-changing transform whose zero terms can be controlled by additional information.

The common multiplicity in the normal closure does not create an additional conductor discount in these formulas. The identities

$$\zeta_E = \zeta_K L^{r-1}, \qquad D_E = D_K D_F^{r-1}$$

give

$$Z_E(f) - Z_K(f) = (r-1)Z_L(f), \qquad \frac{\log D_E - \log D_K}{r-1} = \log D_F.$$

The same factor $r - 1$ has therefore already been removed from both the zero sum and the conductor in passing to L. Combining its explicit formula with that of K leaves exactly $\log Q_\chi$ in (48). Any further saving would require additional information on the remaining zero sum, rather than the multiplicity identity alone.

24.1.1 Rational twists and the full tensor-positivity cone

This calculation determines the constraints imposed by rational twists and positive tensor powers. Let $r \geq 3$ and $p \neq r$ be primes, and put

$$K = \mathbb{Q}(\zeta_r), \qquad F = \mathbb{Q}(p^{1/r}), \qquad E = K(p^{1/r}), \qquad n = r - 1.$$

Write $G = \mathrm{Gal}(E/\mathbb{Q}) = N \rtimes \mathbb{F}_r^\times$, where $N \simeq C_r$, and let ρ denote its irreducible representation of degree n. Let Q_0 be the regular representation of G/N, inflated to G. Thus

$$\rho(g) = \begin{cases} n & g = 1, \\ -1 & g \in N \setminus \{1\}, \\ 0 & g \notin N, \end{cases} \qquad Q_0(g) = n\, 1_N(g), \qquad L(s,\rho) = \frac{\zeta_F(s)}{\zeta_{\mathbb{Q}}(s)}. \tag{50}$$

The same notation denotes a representation and its character. For an unramified prime ℓ, the value $(Q_0 - \rho)(\mathrm{Frob}_\ell)$ is r exactly when $\ell \equiv 1 \pmod r$ and p is not an r-th power modulo ℓ, and is zero otherwise. These are the auxiliary primes. The assertions below concern particular positivity and averaging arguments. They do not assert a limitation on every possible deduction from rational Dirichlet GRH.

24.1.2 The inequalities from all tensor powers

Fix nonnegative weights $w_{\ell,k}$, only finitely many of which are nonzero. For a class function h on G, define

$$\mathcal{P}(h) = \sum_{\ell,k} w_{\ell,k} \log \ell \, \frac{1}{|I_\ell|} \sum_{u \in I_\ell} h(\phi_\ell^k u),$$

where I_ℓ is an inertia group and ϕ_ℓ is a lift of arithmetic Frobenius. For a representation this is the corresponding weighted logarithmic-derivative prime-power sum: averaging over inertia is its projector onto inertia invariants. In particular, $\Psi_K = \mathcal{P}(Q_0)$ and $\Psi_L = \mathcal{P}(\rho)$. Pointwise nonnegative class functions have nonnegative $\mathcal{P}$. This formulation includes ramified primes.

Proposition 24.5. *For every integer $m \geq 1$,*

$$\rho^{\otimes m} = u_m Q_0 + v_m \rho, \qquad u_m = \frac{n^{m-1} + (-1)^m}{r}, \qquad v_m = \frac{n^m - (-1)^m}{r}. \tag{51}$$

The coefficients are nonnegative integers. Odd powers are not pointwise nonnegative. Even powers give the bounds

$$\Psi_L \geq -\frac{u_{2k}}{v_{2k}} \Psi_K, \qquad \frac{u_{2k}}{v_{2k}} \searrow \frac{1}{n}.$$

Consequently all even powers improve the lower bound from $\rho^{\otimes 2} = Q_0 + (r-2)\rho$, but give no upper bound. Their limiting lower bound is already the positivity bound from $\mathrm{Reg}_G = Q_0 + n\rho$, equivalently $\zeta_E = \zeta_K L(s,\rho)^n$.

Proof. At 1 and at a nonidentity element of N, respectively, the required equations are

$$n(u_m + v_m) = n^m, \qquad n u_m - v_m = (-1)^m.$$

They have the displayed solution, which also agrees with the zero values outside N. Integrality follows from $n \equiv -1 \pmod r$; nonnegativity is immediate, including $u_1 = 0$. At a nonidentity translation an odd power has value -1. For an even power every value is nonnegative, including each value in the inertia average. Alternatively, writing $S = \rho^{\otimes k}$, its local trace is

$$\mathrm{tr}(\phi_\ell^j \mid (S \otimes S^*)^{I_\ell}) = \frac{1}{|I_\ell|} \sum_{u \in I_\ell} |\,\mathrm{tr}\, S(\phi_\ell^j u)|^2 \geq 0.$$

Finally

$$\frac{u_{2k}}{v_{2k}} = \frac{1}{n} + \frac{n+1}{n(n^{2k}-1)},$$

which decreases to $1/n$. The regular representation has character rn at the identity and zero elsewhere, proving the last assertion. □

Proposition 24.6. *The pointwise nonnegative functions in* $\mathrm{span}_{\mathbb{R}}\{Q_0, \rho\}$ *form the cone generated by* Reg_G *and* $Q_0 - \rho$*. All inequalities obtained by applying* $\mathcal{P}$ *to this cone are exactly*

$$-\frac{\Psi_K}{n} \le \Psi_L \le \Psi_K. \tag{52}$$

Moreover, after fixing the nonnegative mass outside N*, every polynomial-square moment condition in* $1, Q_0, \rho$*, of arbitrary degree, is already implied by the positive semidefiniteness of a single* 2*-by-*2 *Gram matrix. In particular these conditions give no strict improvement of the upper endpoint in* (52)*.*

Proof. By (50), $aQ_0 + b\rho \ge 0$ precisely when $a+b \ge 0$ and $na-b \ge 0$. The exact decomposition is

$$aQ_0 + b\rho = \frac{a+b}{r}\,\mathrm{Reg}_G + \frac{na-b}{r}(Q_0 - \rho).$$

For the moment assertion set $q = Q_0/n = 1_N$. In the algebra of functions supported on N,

$$q^2 = q, \qquad q\rho = \rho, \qquad \rho^2 = nq + (n-1)\rho.$$

Its two disjoint idempotents are

$$e_0 = \frac{q+\rho}{r} = 1_{\{1\}}, \qquad e_1 = \frac{nq-\rho}{r} = 1_{N\setminus\{1\}}.$$

Every polynomial supported on N reduces to its two values, and the square of its modulus is a nonnegative linear combination of e_0, e_1. A general polynomial in $1, q, \rho$ has the additional orthogonal idempotent $1-q$, whose nonnegative mass is part of the fixed abelian marginal and does not involve $\mathcal{P}(\rho)$. Put $M_0 = \mathcal{P}(q)$ and $M_1 = \mathcal{P}(\rho)$. The Gram matrix for q, ρ is

$$\begin{pmatrix} M_0 & M_1 \\ M_1 & nM_0 + (n-1)M_1 \end{pmatrix}, \qquad \det = (nM_0 - M_1)(M_0 + M_1).$$

For $M_0 > 0$ it is positive semidefinite exactly when $-M_0 \le M_1 \le nM_0$; these inequalities make the other diagonal entry nonnegative as well. For $M_0 = 0$, positivity requires $M_1 = 0$, giving the same conclusion. Since $\Psi_K = nM_0$, this is (52). Conversely, these inequalities say $\mathcal{P}(e_0), \mathcal{P}(e_1) \ge 0$; together with $\mathcal{P}(1-q) \ge 0$, they imply every higher polynomial-square condition. A positive measure supported at 1 attains the upper endpoint. This last statement concerns the moment constraints; it does not claim that an arbitrary prescribed measure comes from the primes of a fixed number field. □

For example, the upper comparison between any even power $m \ge 4$ and the second power is exactly a multiple of the existing detector:

$$n^{m-2}\rho^2 - \rho^m = \frac{n^{m-2}-1}{r}(Q_0 - \rho).$$

Thus the full moment audit is more precise than saying that all powers give the same inequality: they improve the lower endpoint, while the upper endpoint is unchanged.

24.1.3 Rational twists retain the nonabelian part and its conductor

Proposition 24.7. *Let ψ be any finite-order rational Dirichlet character. If $d = \mathrm{ord}_r(p) > 1$, then its twist satisfies*

$$a_p(\rho \otimes \psi) \geq n, \qquad \mathfrak{f}(\rho \otimes \psi) \geq p^n, \tag{53}$$

where a_p is the local Artin-conductor exponent and $\mathfrak{f}$ the positive integer finite conductor. This holds even when p divides the conductor of ψ. If ψ is unramified at p, equality holds in the first bound. If $d = 1$, the weaker uniform bound $a_p(\rho \otimes \psi) \geq n - 1$ holds.

Proof. At p, the inertia image in G is N; the decomposition group is $N \rtimes C_d$, and a Frobenius lift acts on N by $\tau \mapsto \tau^p$. If $d > 1$, the commutators of a Frobenius lift with inertia map onto N, because $p - 1$ is invertible modulo r. They lie in absolute inertia and are killed by every one-dimensional character ψ of the absolute local Galois group. The restriction of $\rho \otimes \psi$ to these elements therefore has the same invariant vectors as the restriction of ρ to N, which has none. Its inertia-invariant space is zero. The conductor formula

$$a_p(V) = \dim V - \dim V^{I_p} + \mathrm{Swan}_p(V)$$

gives (53). An unramified twist preserves the tame representation at p, giving exponent n.

For $d = 1$, if ψ is nontrivial on the kernel of the inertia action in N, there are again no invariants. Otherwise it factors through N. The restriction $\rho|_N$ is the sum of the $r - 1$ distinct nontrivial characters of N; a scalar twist makes at most one summand trivial. Thus the invariant dimension is at most one and the claimed bound follows. □

Proposition 24.8. *Let $B = \mathbb{Q}(\zeta_M)$, and use complete Artin local factors for the characters ψ of* $\mathrm{Gal}(B/\mathbb{Q})$. *Then*

$$\prod_{\psi} L(s, \rho \otimes \psi) = \frac{\zeta_{FB}(s)}{\zeta_B(s)}. \tag{54}$$

Every character η of G/N is a self-twist of ρ: $\rho \otimes \eta \simeq \rho$. Every diagonal Rankin–Selberg product of a rational twist is likewise unchanged:

$$(\rho \otimes \psi) \otimes (\rho \otimes \psi)^* \simeq \rho \otimes \rho^*.$$

Consequently these two averaging operations do not replace the nonabelian factor by rational Dirichlet L-functions.

Proof. The commutator subgroup of G is N; hence $E \cap B \subseteq K$. The subgroup $\mathrm{Gal}(EB/B)$, acting on the roots of $X^r - p$, contains N and is transitive. Therefore $[FB : B] = r$. The permutation representation attached to FB is the tensor product of those attached to F and B. The former is $1 + \rho$, and the latter is $\bigoplus_\psi \psi$. Artin formalism, at every local factor, gives

$$\zeta_{FB}(s) = \prod_{\psi} L(s, \psi) L(s, \rho \otimes \psi),$$

which proves (54). The self-twist statement follows immediately from the character values: $\eta = 1$ on N, and $\rho = 0$ outside N. The Rankin–Selberg identity follows by cancellation of $\psi\overline{\psi}$. In Hecke notation the general twist is

$$L(s, \rho \otimes \psi) = L_K\big(s, \chi \cdot (\psi \circ N_{K/\mathbb{Q}})\big).$$

The norm-pullback factor is Galois invariant, so multiplication by it does not shorten the full Galois orbit of the order-r character χ. It remains an induced factor of degree n. □

There is also an exact finite moment statement. Enlarge B to contain K. On an abelian Frobenius fiber whose multiplier in $\mathbb{F}_r^\times$ is 1, every rational twist has value $\psi(a)\rho(g)$, with the scalar $\psi(a)$ fixed on that fiber. The remaining values of $\rho(g)$ are just $n, -1$. All mixed tensor-square conditions therefore reduce, on each such fiber, to the two-idempotent algebra in Proposition 24.6. On the other fibers $\rho = 0$. Knowledge of the total mass in every abelian fiber, even exact knowledge, does not force a positive mass at nonidentity translations: all mass in a multiplier-1 fiber may satisfy the moment constraints at the value n. This proves a limitation of these finite moment conditions, not an impossibility of using additional analytic relations between the twisted L-functions.

24.1.4 A quantitative attempt using twist zero density

The preceding conductor obstruction can be compared with a uniform zero-density theorem. Cyclic automorphic induction gives a cuspidal representation π_ρ of $\mathrm{GL}_n(\mathbb{A}_\mathbb{Q})$ with $L(s, \pi_\rho) = L(s, \rho)$; its cuspidality follows from the full orbit of χ (Arthur and Clozel, 1989, Chapter III, §6, Theorem 6.2 and Lemma 6.4, pp. 215–218; §7, Proposition 7.2 and the local comparison, pp. 220–221). Write C_ρ for its analytic conductor at height zero, in the normalization of Lemke Oliver and Thorner. Their Corollary 1.4 gives absolute constants $c_1 > 0$ and $C \geq 1$ such that

$$\sum_{\substack{q \leq T \\ (q, \mathfrak{f}(\rho))=1}} \sideset{}{^*}\sum_{\psi \bmod q} N_{\pi_\rho \otimes \psi}(1-\delta, T) \leq C(C_\rho T)^{c_1 n \delta}, \qquad T \geq 1, \quad 0 \leq \delta \leq \tfrac{1}{2}. \tag{55}$$

Here the star restricts to primitive characters; N counts zeros with real part at least $1-\delta$ and imaginary part of absolute value at most T. The source is Lemke Oliver and Thorner (2019), author version dated 19 August 2017, Corollary 1.4 on printed/PDF page 5; the analytic-conductor definition is equation (2.3), page 12. This theorem is unconditional. Applying rational Dirichlet GRH to the abelian factors does not remove C_ρ from its stated bound.

In the hard case $p^{r-1} \equiv 1 \pmod{r^2}$, the previously computed discriminant is $\mathfrak{f}(\rho) = D_F = r^{r-2}p^{r-1}$. The archimedean shifts are $n/2$ copies of 0 and $n/2$ copies of 1. Thus, in that source's normalization,

$$C_\rho = 12^{n/2} r^{r-2} p^{r-1}, \qquad \log C_\rho = n \log p + (r-2)\log r + \tfrac{n}{2}\log 12. \tag{56}$$

In particular $C_\rho \geq p^n$, a feature that persists under the arbitrary rational twists in Proposition 24.7 when $d > 1$.

One possible use of (55) is to seek a twist having no zero in the fixed finite strip by comparing the zero count with the number of twists. Even allowing the optimistic upper bound T^2 for the family size, this pigeonhole certificate requires

$$C(C_\rho T)^{c_1 n \delta} < T^2.$$

Hence a necessary condition for this particular certificate is

$$(2 - c_1 n \delta)\log T > c_1 n \delta \log C_\rho + \log C. \tag{57}$$

If $c_1 n \delta \geq 2$, this condition cannot hold. Otherwise it requires

$$\log T > \frac{c_1 n \delta}{2 - c_1 n \delta} \log C_\rho.$$

For fixed $r, \delta > 0$ this is a positive power of p, not a fixed power of $\log p$. Conversely, for $T \leq (r \log p)^A$, a fixed A, the available strip width for this certificate is at most

$$\delta < \frac{2 \log T}{c_1 n (\log C_\rho + \log T)} \leq \frac{2A \log(r \log p)}{c_1 (r-1)^2 \log p}.$$

These inequalities describe what the quoted bound can certify; they do not assert the existence of zeros in any twist. Moreover a zero-free estimate for one complex twist would still need a separate untwisting argument to yield the positive auxiliary-prime detector.

For completeness, the elementary finite mean square has the same absence of a strict saving. Let $a_j = \Lambda_\rho(j)w(j/x)$, where w is bounded and supported in $[0,2]$, and set

$$S_\psi = \sum_{\substack{j\le 2x\\(j,M)=1}} a_j\psi(j).$$

Orthogonality and Cauchy–Schwarz give the exact identity and bound

$$\frac{1}{\varphi(M)}\sum_{\psi \bmod M}|S_\psi|^2 = \sum_{a\in(\mathbb{Z}/M\mathbb{Z})^\times}\left|\sum_{\substack{j\le 2x\\ j\equiv a \bmod M}} a_j\right|^2 \le \left(1+\frac{2x}{M}\right)\sum_{j\le 2x}|a_j|^2.$$

Since $|\Lambda_\rho(j)| \le n\Lambda(j)$, the elementary bound $\sum_{j\le 2x}\Lambda(j)^2 \ll x\log x$ yields

$$|S_1| \ll_w n\sqrt{\varphi(M)\left(1+\frac{2x}{M}\right)x\log x}.$$

This does not by itself force the strict upper margin needed for the detector. Adding the exact polynomial-square constraints gives precisely the endpoint in Proposition 24.6. Thus this rational-twist attempt has an explicit residual requirement: an analytic estimate must distinguish the two translation classes inside an abelian Frobenius fiber, with conductor dependence stronger than (55) at the requested polynomial-logarithmic scale. No such estimate is established here.

The classical conductor-uniform zero-free theorem also preserves the dependence on p: applied over K to the nonreal character χ and the trivial character, it gives width of order

$$\big(\log D_F + \log D_K + r + (r-1)\log(|t|+3)\big)^{-1},$$

with no exceptional real zero in this application (Humphries and Brumley, 2019, publisher online-first PDF, Appendix A, Theorem A.1, p. 12, and Remark A.4, p. 13). For fixed r its product with $O(\log\log p)$ tends to zero, so this estimate does not supply the fixed width used in Theorem 18.5.

24.2 An exact residual zero criterion in the difficult scalar case

Let r be an odd prime, let $p \ne r$ be prime, and suppose

$$d = \operatorname{ord}_r(p), \qquad v_r(p^d-1) \ge 2.$$

Put $F = \mathbb{Q}(p^{1/r})$, $K = \mathbb{Q}(\zeta_r)$, $E = K(p^{1/r})$, and

$$L(z) = \frac{\zeta_F(z)}{\zeta_{\mathbb{Q}}(z)} = L_K(z,\chi), \qquad Q_\chi = p^{r-1},$$

where χ is any nontrivial character of $\mathrm{Gal}(E/K)$, interpreted as a primitive Hecke character. Call a rational prime ℓ auxiliary if

$$\ell \nmid pr, \qquad \ell \equiv 1 \pmod r, \qquad p^{(\ell-1)/r} \not\equiv 1 \pmod \ell.$$

The following criterion concerns the zero sum of L. It is an equivalent finite-scale detector, not a proof that the required strict inequality has been derived with A as the sole new analytic input; the separate fixed-strip theorem is excluded from this particular derivation.

Lemma 24.9 (The local factors of the detector)**.** *For* $\operatorname{Re} z > 1$,

$$\frac{\zeta_K(z)}{L(z)} = (1-p^{-dz})^{-(r-1)/d} \prod_{\ell \text{ auxiliary}} \frac{1-\ell^{-rz}}{(1-\ell^{-z})^r}. \tag{58}$$

In particular, the negative logarithmic derivative on the right has nonnegative coefficients. Apart from the displayed factor at p, its support consists precisely of powers ℓ^k of auxiliary primes for which $r \nmid k$.

Proof. The Galois group of $E/\mathbb{Q}$ is the affine group $\mathbb{F}_r \rtimes \mathbb{F}_r^\times$, acting on the r roots of $X^r - p$. At an unramified prime, a Frobenius with nontrivial multiplier of order e has one fixed point and $(r-1)/e$ cycles of length e. Its local factor in ζ_F is therefore the product of the local factors in $\zeta_K\zeta_\mathbb{Q}$. For multiplier 1, Frobenius is either the identity, when the same cancellation occurs, or a nontrivial translation, which is an r-cycle. The latter case is exactly the auxiliary-prime condition and gives the displayed factor. Equivalently, no power of an affine element outside the nontrivial translations becomes a nontrivial translation: if its multiplier $c \neq 1$ satisfies $c^k = 1$, then its translation part is multiplied by $1 + c + \cdots + c^{k-1} = 0$.

At p, the polynomial $X^r - p$ is Eisenstein, so F has one prime of residue degree 1. Its factor cancels that of $\zeta_\mathbb{Q}$. The field K has $(r-1)/d$ primes of residue degree d, giving the displayed factor at p. At r, the assumption implies $p^{r-1} \equiv 1 \pmod{r^2}$. The decomposition $\mathbb{Z}_r^\times = \mu_{r-1} \times (1 + r\mathbb{Z}_r)$ and the r-adic logarithm imply that p is an rth power in $\mathbb{Q}_r$. Consequently

$$F \otimes \mathbb{Q}_r \simeq \mathbb{Q}_r \times \mathbb{Q}_r(\zeta_r).$$

Both factors have residue degree 1, so the local factors of $\zeta_K\zeta_\mathbb{Q}/\zeta_F$ cancel also at r. Finally, logarithmic differentiation gives

$$-\frac{d}{dz}\log\frac{\zeta_K(z)}{L(z)} = (r-1)\log p \sum_{m\geq 1} p^{-dmz} + r \sum_{\ell \text{ auxiliary}} \log\ell \sum_{\substack{k\geq 1 \\ r\nmid k}} \ell^{-kz}.$$

All operations are justified by absolute convergence in $\operatorname{Re} z > 1$. □

Use the transforms and zero sums of Section 24.1. For an admissible even real test f supported on $[-S, S]$, set

$$P_p(f) = (r-1)\log p \sum_{m\geq 1} p^{-dm/2} f(dm\log p),$$
$$\mathcal{E}_f = Z_L(f) - Z_K(f) - f(0)\log Q_\chi + 2\Phi_f(1) - 2P_p(f).$$

Proposition 24.10 (Exact zero excess)**.** *One has the unconditional identity*

$$\mathcal{E}_f = 2r \sum_{\ell \text{ auxiliary}} \log\ell \sum_{\substack{k\geq 1 \\ r\nmid k}} \ell^{-k/2} f(k\log\ell). \tag{59}$$

If $f \geq 0$ and $f(t) > 0$ for $0 < t < S$, then $\mathcal{E}_f \geq 0$, and $\mathcal{E}_f > 0$ if and only if there is an auxiliary prime $\ell < e^S$.

Proof. The logarithmic derivative in Lemma 24.9 shows that $T_K(f) + T_\mathbb{Q}(f) - T_F(f)$ equals $P_p(f)$ plus the auxiliary-prime sum on the right of (59), divided by 2. Substitute this in (48). The positivity statement follows term by term. If the sum is positive, some power $\ell^k < e^S$ occurs, hence $\ell < e^S$. Conversely such a prime contributes a strictly positive $k = 1$ term. A prime at e^S contributes zero because continuity and compact support imply $f(S) = 0$. □

For the strip-positive triangle $f_S = e^{-|t|/2}(1-|t|/S)_+$, $S \geq 1$, this becomes

$$P_p(S) = (r-1)\log p \sum_{m\geq 1} p^{-dm}\left(1 - \frac{dm\log p}{S}\right)_+, \tag{60}$$

$$\begin{aligned}\mathcal{E}(S) &= Z_L(f_S) - Z_K(f_S) - (r-1)\log p + S + 2 - \frac{2(1-e^{-S})}{S} - 2P_p(S) \\ &= 2r \sum_{\ell \text{ auxiliary}} \log \ell \sum_{\substack{k\geq 1 \\ r\nmid k}} \frac{1}{\ell^k}\left(1 - \frac{k\log \ell}{S}\right)_+.\end{aligned} \tag{61}$$

The ramified correction is uniformly small:

$$0 \leq P_p(S) \leq \frac{(r-1)\log p}{p^d - 1} \leq \frac{3\log r}{r}. \tag{62}$$

Indeed $B = p^d \geq r^2 + 1$, and $\log B/(B-1)$ decreases for $B > 1$. Thus the middle expression is at most $(r-1)\log(r^2+1)/r^2 \leq 3\log r/r$. There is also a quantitative witness to strictness. If an auxiliary prime satisfies $\ell \leq X$, $X \geq 2$, its first power contributes at $S = \log(2X)$ at least

$$\mathcal{E}(\log(2X)) \geq \frac{2r(\log 2)^2}{X\log(2X)}. \tag{63}$$

Here $1/\ell \geq 1/X$, $\log \ell \geq \log 2$, and $1 - \log \ell/\log(2X) \geq \log 2/\log(2X)$.

Consequently a fixed polynomial auxiliary-prime bound is equivalent to a strict inequality

$$Z_L(f_S) > (r-1)\log p + Z_K(f_S) - 2\Phi_{f_S}(1) + 2P_p(S)$$

at some $S = O(\log(r + \log p + 3))$, with a uniform implied constant. Assumption A concerns the rational Dirichlet factors of ζ_K and hence the abelian sum Z_K. The uncompleted step of this particular derivation is a lower bound for the zero sum Z_L with a strict excess over the displayed threshold. Nonnegativity of Z_L alone is much weaker.

24.2.1 A quantitative attempt using the high zeros

The high zeros do yield an unconditional lower bound for $Z_L(f_S)$, with all conductor dependence visible. For $0 \leq a \leq 1$, let

$$A_{a,S}(\gamma) = \int_0^S (1 - t/S)e^{-at}\cos(\gamma t)\,dt.$$

For $|\gamma| \geq 3$ and $S \geq 1$,

$$\operatorname{Re}\Phi_{f_S}(\beta + i\gamma) = A_{\beta,S}(\gamma) + A_{1-\beta,S}(\gamma) \geq \frac{1}{1+\gamma^2} \qquad (0 \leq \beta \leq 1). \tag{64}$$

To check this directly, put $z = a - i\gamma$ and integrate:

$$A_{a,S}(\gamma) = \operatorname{Re}\left(\frac{1}{z} - \frac{1 - e^{-Sz}}{Sz^2}\right).$$

Writing $v = e^{-aS}$, the real numerator of $(1 - e^{-Sz})/z^2$ is

$$(a^2 - \gamma^2)(1 - v\cos(S\gamma)) + 2a\gamma v\sin(S\gamma).$$

For $a > 0$ it is at most $av[-S(\gamma^2 - a^2) + 2|\gamma|] \leq 0$; use $1 - v \geq aSv$ and $|\gamma| \geq 3$. The case $a = 0$ is immediate. Hence $A_{a,S}(\gamma) \geq a/(a^2+\gamma^2)$. If $u = \beta(1-\beta)$ and $t = \gamma^2$, the numerator of

$$\frac{\beta}{\beta^2 + t} + \frac{1-\beta}{(1-\beta)^2 + t} - \frac{1}{1+t}$$

over its positive common denominator is $u(1-u+3t) \geq 0$. This proves (64).

Let $n = r-1$ and $N_L(T) = \#\{\rho_L : |\operatorname{Im}\rho_L| \leq T\}$, with multiplicity. The identity $\zeta_F = \zeta_{\mathbb{Q}} L$ and the explicit zero count of Hasanalizade et al. (2022, Corollary 1.2, p. 278) give, for $T \geq 1$,

$$\begin{aligned} N_L(T) &= \frac{T}{\pi}\log D_F + \frac{nT}{\pi}\log\frac{T}{2\pi e} + E_L(T), \\ |E_L(T)| &\leq B(T) := 0.228\log D_F + 0.228(r+1)\log T + 23.108(r+1) + 9.040. \end{aligned} \tag{65}$$

For any $T \geq 3$, positivity of the omitted low-zero terms and partial summation therefore imply

$$\begin{aligned} Z_L(f_S) \geq \frac{\log D_F}{\pi}\arctan(1/T) + \frac{n}{\pi}\int_T^\infty \frac{\log(t/(2\pi))}{1+t^2}\,dt \\ - \frac{2B(T)}{1+T^2} - 0.114(r+1)\log(1+T^{-2}). \end{aligned} \tag{66}$$

For completeness, with $w(t) = (1+t^2)^{-1}$, the error in $\int_{(T,\infty)} w\,dN_L$ is bounded by

$$w(T)B(T) + \int_T^\infty (-w'(t))B(t)\,dt = \frac{2B(T)}{1+T^2} + 0.114(r+1)\log(1+T^{-2}).$$

The formula remains valid when a zero has height exactly T, because the integral uses (T,∞) and $N_L(T)$ includes the endpoint.

This is a positive conductor contribution, but it does not prove the required excess. Before its error is subtracted, the coefficient of $\log D_F$ furnished by (66) is at most $\pi^{-1}\arctan(1/3) < 0.103$. For bounded T, the zero-counting error also has order $\log D_F$. In particular, this estimate does not cancel the coefficient 1 of $\log Q_\chi$ to an error of size $O(\log r + \log\log p)$, even with r fixed. It is weaker than the exact non-strict inequality already contained in (59).

One may try to recover more conductor mass with a test concentrated close to zero in the logarithmic variable. The following exact observation isolates what that recovery does. If g is any admissible even test supported strictly inside $(-\log 2, \log 2)$, no prime-ideal power enters its explicit formula. In particular $P_p(g) = 0$, and

$$Z_L(g) - Z_K(g) = g(0)\log Q_\chi - 2\Phi_g(1), \qquad \mathcal{E}_g = 0. \tag{67}$$

Thus adding such a narrow logarithmic test can recover a known conductor contribution to the zero sums, but changes the residual excess by exactly zero. This observation does not exclude other transforms or a lower bound using arithmetic information about the low zeros. It identifies why conductor recovery by itself does not establish strictness at the desired finite scale.

25 Shifts, norms, and rational twists at a specified characteristic

Throughout this section $r \geq 3$ and $p \neq r$ are primes, $K = \mathbb{Q}(\zeta_r)$, $E = K(p^{1/r})$, $F = \mathbb{Q}(p^{1/r})$, and $n = r-1$. An auxiliary prime means a prime $\ell \equiv 1 \pmod r$ for which p is not an r-th power modulo ℓ. The stronger progression 1 mod $12r$ used in the algebraic factorization interface is not imposed in the statements below. The purpose is to examine specific ways of transferring an average over characteristics to a specified p. None of the following restrictions rules out every possible use of rational Dirichlet GRH.

25.1 Which rational substitutions preserve the Kummer extension?

Proposition 25.1. *For $a \in \mathbb{Q}^\times$,*

$$K(a^{1/r}) = E \quad \Longleftrightarrow \quad a = p^j c^r \quad (1 \le j \le r-1,\ c \in \mathbb{Q}^\times).$$

In particular, the only positive rational prime in this orbit is p. The number of positive integers in this orbit not exceeding X is

$$\sum_{j=1}^{r-1} \left\lfloor (X/p^j)^{1/r} \right\rfloor \le (r-1)(X/p)^{1/r}.$$

Proof. Put $t = p^{1/r}$, and choose the generator τ of $\mathrm{Gal}(E/K)$ with $\tau(t) = \zeta_r t$. If $\alpha^r = a$ generates E/K, then $\tau(\alpha)/\alpha = \zeta_r^j$ for some $1 \le j < r$. The eigenspaces of τ, in the basis $1, t, \dots, t^{r-1}$, are Kt^j. Thus $\alpha = bt^j$, where $b \in K^\times$, and $z := a/p^j = b^r \in \mathbb{Q}^\times$. Taking norms to $\mathbb{Q}$ gives $(N_{K/\mathbb{Q}} b)^r = z^{r-1}$. Consequently $c = z/N_{K/\mathbb{Q}} b$ is rational and $c^r = z$. The reverse implication follows because j is coprime to r. If a is a positive integer, its valuations force c to be a positive integer: outside p this is immediate, and at p the inequality $j + r v_p(c) \ge 0$, with $0 < j < r$, implies $v_p(c) \ge 0$. The representations have distinct p-adic valuations modulo r, which proves the count and the assertion about primes. □

This proposition prevents the prime exceptional-set estimate from being applied to multiple prime radicands defining this same extension. It makes no assertion about transformations that change the extension and provide some additional, separately proved transport map.

Proposition 25.2. *Let $q \notin \{p, r\}$ be prime and $q \equiv 1 \pmod r$. If $v_q(N_{F/\mathbb{Q}}\alpha)$ is not divisible by r, for some $\alpha \in F^\times$, then p is an r-th power modulo q. In particular, a prime $q = |N_{F/\mathbb{Q}}\alpha|$ in this progression is not an auxiliary prime.*

Proof. If p is not an r-th power modulo q, Frobenius acts transitively on the r roots of $X^r - p$ over $\mathbb{F}_q$: it acts by a nonidentity translation in the group of r-th roots of unity. Hence $X^r - p$ is irreducible modulo q. Since $q \nmid rp$, the prime is unramified and inert in F, with residue degree r. The valuation formula for norms then gives $v_q(N\alpha) = r v_{\mathfrak{q}}(\alpha)$, also for fractional α. This proves the contrapositive. □

25.2 The coefficient cost of an additive shift

The sparse multiplicative character expansion used to count exceptional characteristics has coefficient norm 1. A nonzero additive shift does not preserve that property.

Proposition 25.3. *Let ℓ be prime, $r \mid \ell - 1$, and $\ell \ge 2r + 1$. Put $H = (\mathbb{F}_\ell^\times)^r$, fix $c \in \mathbb{F}_\ell^\times$, and define $f_c(u) = 1_H(u + c)$ for $u \in \mathbb{F}_\ell^\times$. For a multiplicative character ψ, set*

$$\widehat{f_c}(\psi) = \frac{1}{\ell - 1} \sum_{u \in \mathbb{F}_\ell^\times} f_c(u)\overline{\psi(u)}, \qquad \mu = \frac{1}{r} - \frac{1_H(c)}{\ell - 1}.$$

Then

$$\begin{aligned} \widehat{f_c}(1) &= \mu, \\ |\widehat{f_c}(\psi)| &\le \frac{1 + (r-1)\sqrt{\ell}}{r(\ell - 1)} \quad (\psi \ne 1), \\ \frac{\sqrt{\ell}}{8r} &\le \sum_{\psi} |\widehat{f_c}(\psi)| \le \sqrt{(\ell - 1)/r}. \end{aligned}$$

For a squarefree product $M = \prod_{\ell\in\mathcal{L}} \ell$ of such primes, the coefficient norm of

$$u \longmapsto \prod_{\ell\in\mathcal{L}} 1_{(\mathbb{F}_\ell^\times)^r}(u+c) \quad \text{on } (\mathbb{Z}/M\mathbb{Z})^\times$$

is the product of the local coefficient norms. In particular it is at least

$$\prod_{\substack{\ell\in\mathcal{L}\\ \ell\mid c}} \frac{\sqrt{\ell}}{8r}.$$

An unshifted local factor has coefficient norm exactly 1.

Proof. Let χ have order r, and extend every character, including the principal character, by zero at zero. Then

$$1_H(v) = \frac{1}{r}\sum_{a=0}^{r-1} \chi^a(v).$$

There are $(\ell-1)/r - 1_H(c)$ allowed values of u, giving the formula for μ. If $\psi \neq 1$, the term $a = 0$ in the Fourier coefficient has numerator of absolute value 1. For $a \neq 0$, substituting $u = -cv$ gives the numerator

$$\chi^a(c)\overline{\psi(-c)}J(\overline{\psi}, \chi^a).$$

Its absolute value is 1 when $\overline{\psi}\chi^a = 1$, and $\sqrt{\ell}$ otherwise. These Jacobi-sum identities follow from $|G(\eta)|^2 = \ell$ for nonprincipal η, $G(\alpha)G(\beta) = J(\alpha,\beta)G(\alpha\beta)$, and $J(\alpha,\overline{\alpha}) = -\alpha(-1)$. They give the stated maximum coefficient bound.

Parseval and $f_c^2 = f_c$ give

$$\sum_{\psi\neq 1} |\widehat{f_c}(\psi)|^2 = \mu(1-\mu).$$

Here $\mu \geq 1/(2r)$ and $1-\mu \geq 1/2$, while the maximum nonprincipal coefficient is at most $\sqrt{\ell}/(\ell-1)$. The inequality $\sum|z_i| \geq \sum|z_i|^2/\max|z_i|$ now gives the lower bound. Cauchy–Schwarz and Parseval give the upper bound. Finally, the Chinese remainder theorem makes both the Fourier transform and its coefficient norm factor over ℓ. When $c = 0$ locally, the original expansion has r coefficients equal to $1/r$. □

Thus an estimate obtained by applying individual Dirichlet bounds and the triangle inequality to a shifted exceptional-set indicator must include this coefficient norm. For example, each changed local factor with $\ell \geq 256r^2$ contributes a factor at least 2. This is a statement about the bound furnished by that method, not a lower bound for the prime-counting error.

25.3 A large family of twists with finite-height zero-free strips

There is a positive zero-density consequence for characters whose modulus depends on p. It still requires an untwisting argument to say anything about the specified untwisted function.

Proposition 25.4. *Suppose $d = \mathrm{ord}_r(p) > 1$ and $v_r(p^d-1) \geq 2$. Let $\rho = \mathrm{Ind}_K^{\mathbb{Q}}\chi$, where χ is a nontrivial character of $\mathrm{Gal}(E/K)$. There is an absolute effective constant $c_5 \geq 1$ such that, with $\delta = 1/(4c_5 n)$, all but*

$$O\big((pT)^{1/4}/\sqrt{n}\big)$$

of the $p-1$ Dirichlet characters $\psi \bmod p$ have $L(s, \rho\otimes\psi)$ free of zeros in

$$\Re s > 1-\delta, \qquad |\Im s| \leq T.$$

This holds unconditionally for every $T \geq 1$. In particular, for $T \leq p$ the exceptional number is $O(\sqrt{p}/\sqrt{n})$.

Proof. Put $B = \mathbb{Q}(\zeta_p)$. Since $p \neq r$, the field $KB = \mathbb{Q}(\zeta_{pr})$ has degree $(r-1)(p-1)$ over $\mathbb{Q}$, so $[KB : K] = p - 1$ and $\mathrm{Gal}(KB/K) \simeq \mathbb{F}_p^\times$. Since $r \nmid p-1$, the extensions E/K and KB/K have coprime degrees r and $p-1$. Thus $\mathrm{Gal}(EB/K) = C_r \times \mathbb{F}_p^\times$. The hard-case congruence implies $p^{r-1} \equiv 1 \pmod{r^2}$, so p is an r-th power in $\mathbb{Q}_r^\times$. Hence E/K is unramified at r; its only ramification is tame and total at the primes above p. Each χ^j, $j \neq 0$, therefore has conductor $p\mathcal{O}_K$.

The characters associated with EB/K are $\chi^j(\psi \circ N_{K/\mathbb{Q}})$. At a prime above p, the two factors have respective orders dividing r and $p-1$. Their nontrivial inertia characters cannot cancel when $j \neq 0$. All these conductors divide $p\mathcal{O}_K$, and the maximum conductor norm is p^n. The n characters with a fixed ψ and $j \neq 0$ have the same L-function, namely $L(s, \rho \otimes \psi)$: the Galois action on $K/\mathbb{Q}$ permutes them transitively and fixes the norm pullback.

The log-free density estimate of Thorner and Zaman (Thorner and Zaman, 2019, Theorem 3.2, printed p. 1051; equation (3-1), p. 1050) gives

$$\sum_{\eta} N(1-\delta, T, \eta) \ll \left(r^{r-2} p^n n^n T^n\right)^{c_5 \delta}.$$

The optional exceptional-zero factor in that estimate is at most 1 and has been discarded. A bad ψ contributes at least n to the left-hand side. Substituting the specified δ and using $r = n + 1$ proves

$$\#\{\text{bad } \psi\} \ll \frac{r^{(r-2)/(4n)} n^{1/4} (pT)^{1/4}}{n} \ll \frac{(pT)^{1/4}}{\sqrt{n}}.$$

□

The conclusion does not identify a good character in deterministic time polynomial in r and $\log p$, and it need not include the principal character. The next calculation concerns the additional problem of transferring bounds from such twists to the untwisted sum.

25.4 An absolute-error untwisting calculation

Lemma 25.5. *Assume rational Dirichlet GRH. There is an absolute constant C such that, for every character $\xi \bmod q$ and $x \geq 2$,*

$$\left| \sum_{m \leq x} \Lambda(m) \xi(m) \log(x/m) - 1_{\xi \text{ principal}} x \right| \leq C \sqrt{x} \log(2q).$$

This is an absolute complex estimate, not only a bound on its real part.

Proof. For primitive nonprincipal ξ, use the Hadamard product and the full complex identity in Lamzouri et al. (2015, equation (2.3) and Lemma 2.2, corrected author p. 5). Writing B_ξ for the logarithmic derivative at zero of the completed function, logarithmic differentiation at 2 gives

$$B_\xi = \frac{\Lambda_\xi'}{\Lambda_\xi}(2) - \sum_{\varrho} \left(\frac{1}{2-\varrho} + \frac{1}{\varrho} \right).$$

The local zero count

$$\#\{\varrho : |\Im \varrho - t| \leq 1\} \ll \log(2q) + \log(|t| + 3)$$

is the $K = \mathbb{Q}$ case of Thorner and Zaman (2019, Lemma 2.5, printed p. 1050). Under GRH it implies $\sum_\varrho (1 + |\Im \varrho|^2)^{-1} \ll \log(2q)$. Since

$$\frac{1}{2-\varrho} + \frac{1}{\varrho} = \frac{2}{\varrho(2-\varrho)} \ll \frac{1}{1 + |\Im \varrho|^2},$$

and the completed logarithmic derivative at 2 is $O(\log(2q))$, we have $|B_\xi| \ll \log(2q)$. Lemma 2.2 cited above now bounds the smoothed sum by

$$O(\sqrt{x}\log(2q) + \log x \log(2q) + (\log x)^2 + 1),$$

which has the claimed size. The principal primitive case follows from the zeta-function identity in their Lemma 2.1. Passing from a primitive character to a character modulo q removes Euler factors at primes dividing q; their total contribution is at most $\omega(q)(\log x)^2 \ll \sqrt{x}\log(2q)$. □

Proposition 25.6. *Assume rational Dirichlet GRH, and let $2 \le x < p$. Define nonnegative weights*

$$w_m = n\Lambda(m)1_{m\equiv 1 \bmod r}\log(x/m), \qquad S_0 = \sum_{m\le x} w_m.$$

There is an absolute $C_A > 0$ such that

$$S_0 = x + O(n\sqrt{x}\log(2r)), \qquad \left|\sum_{m\le x} w_m\psi(m)\right| \le C_A n\sqrt{x}\log(2pr) \quad (\psi \bmod p,\ \psi \ne 1).$$

Suppose a linear combination $f(m) = \sum_{\psi\ne 1} c_\psi\psi(m)$ satisfies

$$\sum_{m\le x} w_m|1 - f(m)| \le \varepsilon S_0, \qquad 0 \le \varepsilon < 1.$$

Then its coefficient norm $U = \sum_{\psi\ne 1}|c_\psi|$ obeys

$$U \ge \frac{(1-\varepsilon)S_0}{C_A n\sqrt{x}\log(2pr)}. \tag{68}$$

If the auxiliary-prime window up to x is empty, these same sums are exactly the twists of the ρ prime-power sum, with the Euler factors at rp removed.

Proof. The identity $n1_{m\equiv 1 \bmod r} = \sum_{\eta \bmod r}\eta(m)$ and Lemma 25.5 give the two estimates. For $\psi \ne 1$, every product $\eta\psi \bmod rp$ is nonprincipal by the Chinese remainder theorem. The approximation assumption gives

$$(1-\varepsilon)S_0 \le \left|\sum_{m\le x} w_m f(m)\right| \le UC_A n\sqrt{x}\log(2pr),$$

which proves (68).

For the final assertion, the character of ρ is n at the identity, -1 on $N\setminus\{1\}$, and zero outside $N = \mathrm{Gal}(E/K)$. An element of $N \rtimes \mathbb{F}_r^\times$ outside N has order equal to the order of its multiplier. Its powers therefore either remain outside N or equal the identity. For a prime $\ell \equiv 1 \bmod r$ below x, the empty-window assumption forces Frobenius to be the identity. Consequently, for all prime powers $m \le x$ coprime to rp, the coefficient of $-L'/L(s,\rho)$ is $n\Lambda(m)1_{m\equiv 1 \bmod r}$. Multiplication by $\psi(m)$ proves the asserted equality. □

For example, once x is sufficiently larger than $n^2\log^2(2pr)$, a convex combination of nonprincipal characters cannot satisfy the approximation hypothesis with any fixed $\varepsilon < 1$. Allowing complex coefficients removes that obstruction only by increasing their norm.

Here is an exact limitation of one common certificate based on such coefficients. Suppose one also has, for all selected twists, a bound of the usual zero-free-strip form

$$|T_{\rho\psi}(x)| \le C_\delta n\log(2pr)\, x^{1-\delta}, \qquad 0 < \delta < \tfrac{1}{2},$$

where $T_{\rho\psi}$ denotes the same sum with the Euler factors at rp removed. Combining these bounds only by the triangle inequality would certify a contradiction to an empty window only if

$$UC_\delta n \log(2pr) x^{1-\delta} < (1-\varepsilon) S_0.$$

Together with (68), this requires

$$x^{1/2-\delta} < C_A/C_\delta. \tag{69}$$

Thus this absolute-error transfer, with a fixed $\delta < 1/2$ and fixed constants, cannot give arbitrarily large windows, even if the selected twists satisfy those bounds. For a variable δ the exact restriction is (69); no uniform claim about the dependence of C_δ is needed. The conclusion concerns this particular use of individual bounds. It gives neither a lower bound for the twisted sums nor an obstruction to a new correlation estimate retaining cancellation among the twists.

25.5 What a bilinear character calculation adds

The issue is not that interpolation by good twists is algebraically impossible. A sufficiently large family always interpolates a short window. This can be proved without any analytic assumption.

Proposition 25.7. *Let $\mathcal{G}$ be any set of characters modulo p, and let $b = (p-1) - |\mathcal{G}|$. For a set $\mathcal{S} \subseteq \mathbb{F}_p^\times$ of size h, suppose $bh < p-1$. Then there are coefficients c_ψ, $\psi \in \mathcal{G}$, such that*

$$\sum_{\psi\in\mathcal{G}} c_\psi \psi(u) = 1 \quad (u \in \mathcal{S}), \qquad \sum_{\psi\in\mathcal{G}} |c_\psi| \le \sqrt{\frac{(p-1)h}{p-1-bh}}.$$

Thus the good twists in Proposition 25.4, after also excluding the principal character, permit exact interpolation whenever the displayed size condition holds. This last assertion is existence of coefficients; it does not give a polynomial-time procedure to identify the good twists.

Proof. If $h = 0$ all coefficients can be taken to be zero. For $h > 0$, form the h-by-$|\mathcal{G}|$ matrix V with $V_{u,\psi} = \psi(u)$. Orthogonality for the complete character family gives

$$VV^* = (p-1)I_h - WW^*,$$

where W is the matrix of the b omitted columns. Since WW^* is positive semidefinite and its trace is bh, the least eigenvalue of VV^* is at least $p-1-bh > 0$. For $\mathbf{1} \in \mathbb{C}^h$, choose $c = V^*(VV^*)^{-1}\mathbf{1}$. Then $Vc = \mathbf{1}$, and

$$\|c\|_2^2 = \mathbf{1}^*(VV^*)^{-1}\mathbf{1} \le \frac{h}{p-1-bh}.$$

Cauchy–Schwarz and $|\mathcal{G}| \le p-1$ finish the proof. □

For comparison, a small selected family is nearly orthogonal to the constant function in the prime weights under A. This is a bilinear estimate in distinct rational characters, not an additional tensor-power identity for ρ.

Proposition 25.8. *Assume rational Dirichlet GRH and $2 \le x < p$. Use the weights of Proposition 25.6, and put $B = C_A n\sqrt{x}\log(2pr)$. For k distinct nonprincipal characters $\psi_1, \dots, \psi_k \bmod p$, suppose $S_0 > (k-1)B$. Then every $c \in \mathbb{C}^k$ satisfies*

$$\sum_{m \le x} w_m \left| 1 - \sum_{i=1}^{k} c_i \psi_i(m) \right|^2 \ge S_0 - \frac{kB^2}{S_0 - (k-1)B}.$$

Proof. The weighted Gram matrix $G_{ij} = \sum_m w_m \overline{\psi_i(m)}\psi_j(m)$ has diagonal entries S_0, and its off-diagonal entries have absolute value at most B, because $\overline{\psi_i}\psi_j$ is nonprincipal. Consequently $G \ge (S_0 - (k-1)B)I$. The vector $v_i = \sum_m w_m \overline{\psi_i(m)}$ satisfies $\|v\|_2^2 \le kB^2$. Minimizing the quadratic expression $S_0 - 2\Re(c^*v) + c^*Gc$ gives $S_0 - v^*G^{-1}v$, which is bounded as stated. □

For $kB \le S_0/2$, the lower bound is $S_0 - 2kB^2/S_0$. The large-family interpolation proposition and this small-family bound are consistent: they quantify different sizes of the available character family. Neither yields the new cancellation estimate for nonabelian twisted sums that would be needed to advance the uniform auxiliary-prime problem.

26 Coefficient searches and character-product lower bounds

This section quantifies what two explicit cubic searches and the classical character-product construction do and do not guarantee at a prescribed characteristic.

26.1 A coefficient family and the cost of Dirichlet-character transfer

There is a particularly simple family with a positive density of successful coefficients. Its multiplicative Fourier expansion also quantifies the loss in applying rational Dirichlet estimates separately to all the resulting characters.

Proposition 26.1 (Density and multiplicative Fourier norm)**.** *Let $p \equiv 2 \pmod 3$ be an odd prime, and define*

$$F(m) = \mathbf{1}_{\{X^3+X+m \text{ is irreducible over } \mathbb{F}_p\}} \qquad (m \in \mathbb{F}_p).$$

For a multiplicative character χ of $\mathbb{F}_p^\times$, put

$$\widehat{F}(\chi) = \frac{1}{p-1} \sum_{m \in \mathbb{F}_p^\times} F(m)\overline{\chi(m)}, \qquad \mu = \frac{p+1}{3(p-1)}.$$

Then

$$\sum_{m \in \mathbb{F}_p} F(m) = \frac{p+1}{3}, \qquad \widehat{F}(1) = \mu.$$

Every nontrivial character satisfies

$$|\widehat{F}(\chi)| \le \frac{4\sqrt{p}}{3(p-1)},$$

and $\widehat{F}(\chi) = 0$ when $\chi(-1) = -1$. Moreover,

$$\frac{(p+1)(p-2)}{6(p-1)\sqrt{p}} \le \sum_{\chi \ne 1} |\widehat{F}(\chi)| \le \sqrt{(p-2)\mu(1-\mu)}.$$

In particular,

$$\frac{p-1}{6\sqrt{p}} \le \sum_{\chi \ne 1} |\widehat{F}(\chi)| \le \frac{\sqrt{p-2}}{2}.$$

Thus the nonprincipal Fourier ℓ^1 norm has order $\sqrt{p}$.

Proof. Let λ be the quadratic character, extended by zero at zero, and let N_m be the number of roots of $X^3 + X + m$ in $\mathbb{F}_p$. Its discriminant is

$$D_m = -4 - 27m^2.$$

The element -3 is a nonsquare: a square root would give a nontrivial cube root of unity in $\mathbb{F}_p$, whereas $3 \nmid p - 1$. Consequently D_m never vanishes. For a separable cubic, the sign of its Frobenius permutation is the quadratic character of its discriminant. Indeed Frobenius acts on the Vandermonde product of its three roots by this sign, and the square of that product is the discriminant. The three possible root counts $0, 1, 3$ therefore give the identity

$$F(m) = \frac{2 + \lambda(D_m) - N_m}{3}. \tag{70}$$

Each $x \in \mathbb{F}_p$ is a root for exactly one parameter, $m = -x^3 - x$, so $\sum_m N_m = p$. For nonzero $c \in \mathbb{F}_p$ one has

$$\sum_{x \in \mathbb{F}_p} \lambda(x^2 - c) = -1 :$$

the equation $y^2 = x^2 - c$ has $p - 1$ solutions, by the bijection $(x, y) \mapsto (x - y, x + y)$ with pairs of product c. It follows that

$$\sum_m \lambda(-4 - 27m^2) = -\lambda(-27) = 1.$$

Summing (70) proves the asserted count. Also $F(0) = 0$, since $X^3 + X$ has the root zero, so this count gives the stated mean on $\mathbb{F}_p^\times$.

We include the character-sum facts used for the coefficient bound. Extend all multiplicative characters by zero at zero, including the trivial character. For characters A, B, set

$$J(A, B) = \sum_{z \in \mathbb{F}_p} A(z)B(1 - z).$$

If A, B, AB are nontrivial, then $|J(A, B)| = \sqrt{p}$. For completeness, let $e_p(z) = \exp(2\pi i z/p)$ and $G(A) = \sum_z A(z)e_p(z)$. Multiplicative orthogonality gives

$$|G(A)|^2 = \sum_{u \neq 0} A(u) \sum_{y \neq 0} e_p((u - 1)y) = p$$

for nontrivial A. Sorting the terms of $G(A)G(B)$ by $x + y$ gives

$$G(A)G(B) = J(A, B)G(AB).$$

The terms with $x + y = 0$ sum to zero because AB is nontrivial. The claimed magnitude follows.

The function F is even, since replacing X by $-X$ identifies the roots at m and $-m$. Hence all its odd-character coefficients vanish. Fix a nontrivial even character χ, and write $\psi = \overline{\chi}$. The character group is cyclic, and $\psi(-1) = 1$, so there is a character η with $\eta^2 = \psi$. Neither η nor $\eta\lambda$ is trivial.

Put $A_0 = -4$ and $B_0 = -27$. Changing variables $y = m^2$ gives

$$\begin{aligned} S_D &:= \sum_{m \neq 0} \psi(m)\lambda(A_0 + B_0 m^2) \\ &= \sum_y \eta(y)\lambda(A_0 + B_0 y) + \sum_y (\eta\lambda)(y)\lambda(A_0 + B_0 y). \end{aligned}$$

The two sums are, up to factors of modulus one, $J(\eta,\lambda)$ and $J(\eta\lambda,\lambda)$. Their two characters and their products are nontrivial, so

$$|S_D| \le 2\sqrt{p}.$$

Similarly,

$$\begin{aligned} S_N &:= \sum_{m\neq 0} \psi(m) N_m = \sum_x \psi\big(-x(x^2+1)\big) \\ &= \sum_y \eta(y)\psi(1+y) + \sum_y (\eta\lambda)(y)\psi(1+y). \end{aligned}$$

Here the two Jacobi sums, again up to factors of modulus one, are $J(\eta,\psi)$ and $J(\eta\lambda,\psi)$. Their products are η^3 and $\eta^3\lambda$. Since $3 \nmid p-1$, cubing is an automorphism of the character group. Thus $\eta^3 = 1$ would force $\eta = 1$, and $\eta^3 = \lambda$ would force $\eta = \lambda$; both are excluded. All the required characters are therefore nontrivial, and

$$|S_N| \le 2\sqrt{p}.$$

The constant term in (70) has zero sum against ψ. Consequently

$$|\widehat{F}(\chi)| = \frac{|S_D - S_N|}{3(p-1)} \le \frac{4\sqrt{p}}{3(p-1)}.$$

Parseval's identity and $F^2 = F$ give

$$\sum_{\chi\neq 1} |\widehat{F}(\chi)|^2 = \mu - \mu^2 = \frac{2(p+1)(p-2)}{9(p-1)^2}.$$

Dividing by the upper bound for the largest summand's modulus proves the lower bound for the ℓ^1 norm. The further simplification uses $(p+1)(p-2) \ge (p-1)^2$, valid for $p \ge 3$. Cauchy–Schwarz with $p-2$ nonprincipal characters proves the upper bound; finally $\mu(1-\mu) \le 1/4$. □

To quantify a direct transfer from rational Dirichlet estimates, use the fixed nonnegative smooth function w supported in $(1,2)$ from Lemma 18.3, with $c_w = \int_1^2 w(t)\,dt > 0$. For a Dirichlet character modulo p, put

$$\Psi_{\chi,w}(X) = \sum_{n\ge 1} \Lambda(n)\chi(n) w(n/X).$$

Under A, uniformly in these characters and $X \ge 2$,

$$\Psi_{\chi,w}(X) = \mathbf{1}_{\chi=1} c_w X + O_w\Big(\sqrt{X}\log(2p)\Big). \tag{71}$$

Here the error is an absolute bound on a complex number, as needed for the possibly complex Fourier coefficients. To verify it, apply Mellin inversion to $-L'/L$ for the primitive character inducing χ and shift the contour past its pole and zeros. The zero term is $-\sum_\rho X^\rho \widehat{w}(\rho)$. Three integrations by parts give $|\widehat{w}(\sigma+it)| \ll_w (1+|t|)^{-3}$ on $0 \le \sigma \le 1$. The unit-height zero count is $O(\log(2p) + \log(|t|+3))$, so the absolute sum of these transforms is $O_w(\log(2p))$. Under A, $|X^\rho| = \sqrt{X}$. The pole at one gives the stated main term; the gamma factor and trivial zeros contribute $O_w(\log(2p))$, exactly as in the proof of Lemma 18.3. For the principal character modulo p, the omitted Euler factor contributes only prime powers of p in $(X, 2X)$; there is at most one, and its contribution is $O_w(\log p)$. This proves (71).

Fourier inversion and a termwise triangle inequality now give

$$\sum_{n\ge 1} \Lambda(n) F(n \bmod p) w(n/X) = \mu c_w X + O_w\Big(\sqrt{pX}\log(2p)\Big). \tag{72}$$

Proper prime powers contribute $O_w(\sqrt{X})$ by Chebyshev's bound. Thus this estimate guarantees a prime parameter in $(X, 2X)$ when $X = C_w p \log^2(2p)$ and C_w is sufficiently large. This still depends on the numerical characteristic p. Proposition 26.1 shows that the factor of order $\sqrt{p}$ is present in the ℓ^1 norm used by this particular termwise transfer. It does not give a lower bound for the error in (72), and it does not exclude cancellation between characters or a different family.

A smaller explicit prime-interval menu. The lower comparison below uses a different family: one parameter for each split prime in a prescribed interval. Its size is asymptotic to $K/(2\log K)$, while Corollary 23.3 uses large prime divisors of the full product and gives linear size. The following statement fixes the smaller menu exactly.

Lemma 26.2. *Let $K \geq 4$ be an integer and put $N(k) = k^2 + 3k + 9$. For each prime*

$$\ell \in (K+3, 2K+3], \qquad \ell \equiv 1 \pmod 3,$$

let k_ℓ be the smaller least nonnegative representative of a root of $N(k) \equiv 0 \pmod \ell$. Then $1 \leq k_\ell \leq K$, $k_\ell \neq 3$, and $v_\ell(N(k_\ell)) = 1$. No other prime in the displayed interval divides $N(k_\ell)$.

For each such parameter there is a cubic Dirichlet character θ_ℓ of conductor dividing $N(k_\ell)$ such that, for every prime $p \equiv -1 \pmod 9$ not dividing $N(k_\ell)$,

$$h_{k_\ell}(Y) = Y^3 - (k_\ell + 3)Y^2 + k_\ell Y + 1 \text{ is irreducible over } \mathbb{F}_p \quad \Longleftrightarrow \quad \theta_\ell(p) \neq 1.$$

These characters are linearly independent over $\mathbb{F}_3$: every product $\prod_\ell \theta_\ell^{e_\ell}$ with some $e_\ell \not\equiv 0 \pmod 3$ is nonprincipal. More strongly, it has a nontrivial component at a prime $\ell \neq 3$.

Proof. The discriminant of N is -27, so it has two distinct roots modulo every prime $\ell \equiv 1 \pmod 3$. Here $\ell > 7$. Neither $\ell - 1$ nor $\ell - 2$ is a root, since $N(-1) = N(-2) = 7$. Thus the sum of the two least representatives is $\ell - 3$, and the smaller is at most $(\ell-3)/2 \leq K$. It cannot be zero or three because $N(0) = 9$ and $N(3) = 27$. Since

$$N(k_\ell) \leq K^2 + 3K + 9 < (K+3)^2 < \ell^2,$$

its ℓ-valuation is one. The product of any two primes larger than $K+3$ exceeds the same upper bound, proving the other divisibility assertion. The exclusion of $\ell = 7$ is necessary: its two least roots are five and six.

We spell out the reciprocity character to control its conductor. In the Eisenstein integers factor $\alpha = k_\ell - 3\zeta_3$ into units, the ramified prime, inert rational primes, and primary split prime factors. For a hard prime p, the units and ramified prime are cubes modulo p: ζ_3 is a cube because $9 \mid p + 1$, and $(1 - \zeta_3)^2 = -3\zeta_3$ is a cube. Inert rational factors are scalars, hence also cubes. Cubic reciprocity for the remaining primary factors, in the split/inert case of (Friesen et al., 1986, p. 395, equation (1.1), and Section 2, pp. 396–397), therefore expresses the cubic residue symbol of α modulo p as a product of rational cubic Dirichlet characters, each of conductor a split prime dividing $N(k_\ell)$. At conjugate prime factors their characters are inverses. This defines θ_ℓ with conductor dividing $N(k_\ell)$. On the progression $p \equiv -1 \pmod 9$, equation (73) proves the displayed irreducibility criterion. Lemma 23.2 gives the character of the cubic field at every unramified prime.

The designated prime ℓ divides the Eisenstein norm exactly once, so precisely one of its two prime factors divides α, with exponent one. Consequently θ_ℓ has a nontrivial local component at ℓ. No other selected character has a component there, by the norm-divisibility assertion above. Hence a product with $e_\ell \neq 0$ retains that nontrivial component. This proves independence. A character may have additional conductor primes; only the designated prime is private to it. □

26.2 Classical character-product lower comparisons

The character product also produces exceptional primes at controlled scales. This is the classical mechanism for least-quadratic-nonresidue lower bounds: (Montgomery and Vaughan, 2007, p. 432, Exercise 9(a)–(e)) explicitly attributes that argument to Montgomery (1971, p. 122). We record the specializations to the present searches.

Lemma 26.3. *Assume A. Let $Q \geq 2$, let $H \leq (\mathbb{Z}/Q\mathbb{Z})^*$ have index I, and let $J \geq I$. For an absolute sufficiently large C, the coset $(-1)H$ contains a prime in $(x/4, x]$, where*

$$x = CJ^2 \log^2(2Q).$$

Proof. Extended by zero on nonunits, the coset indicator is

$$1_{(-1)H}(n) = I^{-1} \sum_{\chi \in H^\perp} \chi(-1)\chi(n).$$

Its coefficients are real, equal to $1/I$ or $-1/I$, and have absolute-value sum one. Lemma 22.1 gives

$$S(y) := \sum_{n \leq y} \Lambda(n) 1_{(-1)H}(n) \log(y/n) = y/I + O(\sqrt{y}\log(2Q)).$$

The nonnegative logarithmic tent $S(x) - 2S(x/2) + S(x/4)$, supported on $(x/4, x]$ and bounded by $\log 2$, has value

$$x/(4I) + O(\sqrt{x}\log(2Q)).$$

Proper prime powers contribute $O(\sqrt{x})$ by Chebyshev's bound. At the stated x, the main term is at least $CJ\log^2(2Q)/4$, while the error is $O(\sqrt{C}J\log^2(2Q))$. An absolute choice of C makes the prime contribution positive. □

Proposition 26.4. *Assume A and fix an odd prime r. There are arbitrarily large primes $p \equiv -1 \pmod{r^2}$ with no good auxiliary prime at most B, where*

$$B = \left(\frac{r-1}{2\log r} + o(1)\right) \log p \log\log p.$$

These are hard inputs: $\mathrm{ord}_r(p) = 2$ and $v_r(p^2 - 1) \geq 2$. The $o(1)$ is along the constructed sequence with r fixed.

Proof. Let $B \to \infty$. Take all primes $\ell \leq B$ with $\ell \equiv 1 \pmod{r}$, let m be their number, and put $M = \prod \ell$, $Q = r^2 M$. The conditions $n \equiv -1 \pmod{r^2}$ and that n be an rth power at every selected ℓ describe a compatible coset $(-1)H$: every odd-order residue character takes value one at -1. The Chinese remainder theorem gives the exact index

$$I = \varphi(r^2) r^m.$$

Lemma 26.3, with $J = I$, supplies a prime $p \in (x/4, x]$ satisfying these conditions, for $x = CI^2\log^2(2Q)$. The fixed-progression prime number theorem and Chebyshev's bound give

$$m = (1 + o(1))\frac{B}{(r-1)\log B}, \qquad \log Q \ll_r B.$$

Consequently

$$\log p = 2m\log r + O_r(\log B) = \left(\frac{2\log r}{r-1} + o(1)\right)\frac{B}{\log B}.$$

Inverting this relation gives the asserted scale and shows $p \to \infty$. The hard-case assertions follow from the chosen congruence and oddness of r. □

Proposition 26.5. *Assume A. For every sufficiently large integer K, the private-prime menu of Lemma 26.2 has a hard prime $p \equiv -1 \pmod 9$ at which all its cubics split, with*

$$\log p = (\log 3 + o(1))K/\log K, \qquad K = (1/\log 3 + o(1)) \log p \log\log p.$$

Proof. Let m be the menu size and Q nine times the least common multiple of its character conductors. The private conductor primes give the bad coset exact index $I = 6 \cdot 3^m$. It is compatible with $p \equiv -1 \pmod 9$ because cubic characters take value one at -1. Apply Lemma 26.3 with $J = I$. Since $m = (1/2 + o(1))K/\log K$ and $\log Q \ll K$, its prime satisfies $\log p = 2m \log 3 + O(\log K)$, giving the formulas. For large K, $p > K^2 + 3K + 9$, so the cubics are unramified. Their residue characters are trivial. A squarefree cubic with square discriminant has only the split and irreducible possibilities here; the nontrivial-character alternative is precisely the latter. Thus all menu cubics split. □

The private menu is not nested as K varies. This proposition concerns the displayed pairs (p, K); it does not by itself refute a different prescribed rule $K = K(p)$ whose moving prime interval is not the same one.

Proposition 26.6. *Assume A. There are arbitrarily large hard primes p such that every h_k with $1 \le k \le K$ splits over $\mathbb{F}_p$, where*

$$K = \left(\frac{1}{2\log 3} + o(1)\right) \log p.$$

Thus scanning only $1 \le k \le c \log p$ fails for infinitely many hard inputs whenever $c < 1/(2\log 3)$. This does not exclude larger polynomial windows in $\log p$.

Proof. For each $1 \le k \le K$, let ϑ_k be the prime-to-three cubic reciprocity character of $k - 3\zeta_3$, allowing the principal character. It gives the required test on $p \equiv -1 \pmod 9$. For example, ϑ_3 is principal, whereas the character of the cyclic cubic field for $k = 3$ has conductor nine. Its conductor divides $N(k)$ by the same factorization argument used above. Common triviality and $p \equiv -1 \pmod 9$ form a compatible coset $(-1)H$. Its index satisfies $I \le 6 \cdot 3^K$, with no independence assumption. For Q nine times the least common multiple of the conductors,

$$\log Q \le \log 9 + \sum_{k \le K} \log N(k) \ll K \log(2K).$$

Choose the explicit upper index $J = 6 \cdot 3^K$ in Lemma 26.3, whether or not $I = J$. The resulting prime satisfies $\log p = 2K \log 3 + O(\log K)$, and all these cubics split by the unramified residue-character criterion. The parameter intervals are nested, so the final scanning assertion follows. When interpreted via elliptic curves, the singular parameter $k = 3$ can be omitted; it causes no difficulty for the polynomial statement. □

The first proposition compares directly with the almost-all upper coefficient $(r-1)/(2\log r)$. The existence of exceptional primes is consistent with that upper theorem and with a possible uniform bound having a larger constant or a higher logarithmic power. These specializations do not establish independence of the desired all-input construction from A.

27 What the abelian quotient does not determine

This section isolates an algebraic obstruction to a direct transfer of prime-counting information from the cyclotomic quotient.

Proposition 27.1 (Failure of abelian character reconstruction)**.** *Let $G = \mathbb{F}_r \rtimes \mathbb{F}_r^\times$, r odd, and $N = \mathbb{F}_r$. The indicator of $N \setminus \{0\}$ is not a complex linear combination of one-dimensional characters of G. In particular, the condition on the Frobenius element in Theorem 1.6 is not a class-function condition on the abelian quotient alone.*

Proof. For $b \neq 1$, the commutator of dilation by b and translation by c is translation by $(b-1)c$. As c varies these exhaust N, and G/N is abelian. Thus $[G, G] = N$. Every one-dimensional character is trivial on N, so any linear combination of them has the same value at the identity and at every nonzero translation. The indicated function has values zero and one there, respectively. This is a contradiction. □

This proposition is an obstruction to that specific character reconstruction. It does not prove a logical independence of Hecke GRH from Dirichlet GRH, nor a lower bound for field construction. A new analytic implication, or a different algebraic construction, could bypass it.

27.1 Cyclotomic factor counts and auxiliary primes

This subsection tests whether the factor-count construction of Bhargava, Ivanyos, Mittal, and Saxena can bypass the auxiliary-prime search by using products of complete cyclotomic polynomials. Their Property 1.1 asks for a squarefree polynomial having, for some degree divisible by r, a number of irreducible factors of that degree not divisible by r; their Theorem 1.4 turns such a polynomial into extensions of any r-power degree (Bhargava et al., 2017, p. 2). In the higher-valuation case, restricting this search to cyclotomic products leads back to the same prime condition.

Proposition 27.2. *Let $p \neq r$ be primes, let r be odd, and put $s = \operatorname{ord}_r(p)$ and $a = v_r(p^s - 1) \geq 2$. Let M be a finite set of distinct positive integers prime to p, and put*

$$f(X) = \prod_{m \in M} \Phi_m(X) \in \mathbb{F}_p[X], \qquad D = \deg f.$$

Then f is squarefree. If, for some d divisible by r, the number of irreducible factors of f of degree d is not divisible by r, there exist $m \in M$ and a prime $\ell \mid m$ such that

$$\ell \neq p, r, \qquad \ell \equiv 1 \pmod{r}, \qquad p^{(\ell-1)/r} \not\equiv 1 \pmod{\ell}. \tag{$*$}$$

Moreover, $\ell \leq 2D^2$. Conversely, for every prime ℓ satisfying $(*)$, the polynomial $\Phi_\ell \bmod p$ has the stated factor-count property and degree $\ell - 1$. If the list M is given, a prime satisfying $(*)$ can be recovered in deterministic time polynomial in D and $\log p$.

Proof. For m prime to p, write $m = r^k \prod_{\ell \neq r} \ell^{e_\ell}$, with positive displayed e_ℓ, and set

$$B_0 = \max(k-1, 0), \quad C_0 = \max(k-a, 0), \quad B_\ell = v_r(\ell - 1), \quad C_\ell = v_r(\operatorname{ord}_\ell(p)).$$

The relevant valuations are

$$v_r(\varphi(m)) = B_0 + \sum_{\ell \neq r} B_\ell,$$
$$v_r(\operatorname{ord}_m(p)) = \max\bigl(C_0, \{C_\ell : \ell \mid m,\ \ell \neq r\}\bigr).$$

An empty maximum is zero and $\operatorname{ord}_1(p) = 1$. The first equality follows from Euler's product. For the second, orders over relatively prime moduli combine by least common multiple. When $\ell \neq r$, the quotient $\operatorname{ord}_{\ell^e}(p)/\operatorname{ord}_\ell(p)$ is an ℓ-power, since the kernel of reduction on the unit groups is an ℓ-group. It does not change the r-adic valuation. For $k \geq 1$, the odd-prime lifting-the-exponent identity gives

$$v_r(p^{su} - 1) = a + v_r(u), \qquad \operatorname{ord}_{r^k}(p) = s r^{\max(k-a, 0)}.$$

Indeed, an exponent making p equal to one modulo r^k must be a multiple of s, and the displayed valuation determines the least possible additional power of r.

For $\ell \neq p, r$, condition $(*)$ is equivalent to $C_\ell = B_\ell > 0$: writing $\ell - 1 = r^{B_\ell} h$ with $(r, h) = 1$, if $B_\ell = 0$, this prime cannot satisfy $(*)$. For $B_\ell > 0$, the order of p divides $(\ell - 1)/r$ exactly when $C_\ell < B_\ell$. Suppose no prime divisor of m satisfies $(*)$ and write $C = v_r(\mathrm{ord}_m(p))$. If $C > 0$, take a local factor attaining this maximum. For a factor $\ell \neq r$ its corresponding B_ℓ is at least $C + 1$. If the r-primary factor attains it, then $k > a$ and $B_0 - C_0 = a - 1 \geq 1$. In either case $v_r(\varphi(m)) \geq C + 1$. Therefore

$$r \mid \mathrm{ord}_m(p) \quad \Longrightarrow \quad r \mid \frac{\varphi(m)}{\mathrm{ord}_m(p)}. \tag{$**$}$$

Every primitive m-th root of unity has a Frobenius orbit of length $\mathrm{ord}_m(p)$. Thus Φ_m has $\varphi(m)/\mathrm{ord}_m(p)$ irreducible factors, all of that degree. Different conductors prime to p have disjoint root sets, which also proves squarefreeness. For $r \mid d$, the number of degree-d factors of f is consequently

$$\sum_{\substack{m \in M \\ \mathrm{ord}_m(p) = d}} \frac{\varphi(m)}{d}.$$

If none of the conductors has a prime divisor satisfying $(*)$, every summand is divisible by r, by $(**)$. This proves the first assertion by contraposition.

For the size bound,

$$\frac{\varphi(m)^2}{m} = \prod_{\ell^e \| m} \ell^{e-2}(\ell - 1)^2 \geq \frac{1}{2}.$$

Every factor is at least one except possibly the factor $1/2$ from $\ell = 2, e = 1$. Hence $\ell \leq m \leq 2\varphi(m)^2 \leq 2D^2$. Conversely, if $(*)$ holds, then $d = \mathrm{ord}_\ell(p)$ is divisible by r, and $(\ell - 1)/d$ is not. The factorization formula for Φ_ℓ proves the converse. Finally, trial division of the given conductors and modular exponentiation test all possible prime divisors in time polynomial in D and $\log p$, since $|M| \leq D$ and each $m \leq 2D^2$. □

For example, with $p = 17$ and $r = 3$, one has $s = a = 2$, and for every $k \geq 3$ the polynomial Φ_{3^k} has three irreducible factors, each of degree $2 \cdot 3^{k-2}$. Complete cyclotomic products cannot remove this divisibility unless a conductor has a prime divisor satisfying $(*)$. The prime 7 does satisfy it: $17^2 \equiv 2 \pmod 7$ and $\mathrm{ord}_7(17) = 6$.

Thus a polynomial-size witness within this class is equivalent, up to a polynomial change of size, to a polynomial-size auxiliary prime. The proposition does not address proper selected factors, other polynomial families, or succinct representations of much larger cyclotomic algebras. It restricts this construction and makes no claim that Dirichlet GRH cannot imply the desired field construction by another argument.

27.2 Selected factors and invariant subalgebras

An explicit small factor of a cyclotomic polynomial could give a different construction, but selecting that factor is a substantive step. The following observations distinguish this possibility from compression by cyclotomic symmetries.

Suppose $c \in D = \mathbb{F}_p[T]/\Phi_r(T)$ generates the Teichmüller subgroup of order $R = r^a$. Put $z = c^{R/r}$, an element of order r. The defining equivariance of c gives $\delta_b(z) = z^b$ for $\delta_b(T) = T^b$ and every $b \in (\mathbb{Z}/r\mathbb{Z})^*$, since the Teichmüller lift of b is congruent to b modulo r. Choose a primitive r-th root ζ in an algebraic closure and write $z(\zeta) = \zeta^u$. Evaluating the equivariance gives $z(\zeta^b) = \zeta^{ub}$ for every b. If $u = 0$, every geometric value of z would be one, contradicting its order. Hence $(u, r) = 1$ and, by equality at all geometric points of D, $c^{r^{a-1}} = T^u$. Its

characteristic polynomial $g(X)$ therefore has degree $r-1$ and divides $\Phi_{r^a}(X)$: its geometric roots are distinct, because their r^{a-1}-th powers are distinct, and each has order r^a. It follows that

$$h(X) = g(X^r)$$

is a squarefree degree-$r(r-1)$ divisor of $\Phi_{r^{a+1}}(X)$. Every irreducible factor of h has degree sr, and the number of factors is $(r-1)/s$, which is prime to r. Thus an explicit degree-r field yields a small selected factor by the Teichmüller construction; conversely any explicit divisor with this factor-count property yields the field by (Bhargava et al., 2017, Theorem 1.4, p. 2). If the selected factor is irreducible, its quotient is a degree-sr field, whose degree-r subfield is obtained directly as the kernel of $\mathrm{Frob}_p^r - 1$. More generally, it is enough to select a nonconstant divisor $g \mid \Phi_{r^a}$ for which $r \nmid \deg(g)/s$. Every irreducible factor of g has degree s; the squarefree polynomial $g(X^r)$ divides $\Phi_{r^{a+1}}$, has all its irreducible factors of degree sr, and has exactly $\deg(g)/s$ of them. This proves the same factor-count condition. Together with the characteristic-polynomial construction above, this is a polynomial-time equivalence between constructing the maximal Teichmüller unit and constructing such a polynomial-size selected divisor of Φ_{r^a}.

Proposition 27.3. *Let $p \neq r$ be primes, r odd, put $s = \mathrm{ord}_r(p)$ and $a = v_r(p^s - 1)$, and let $k > a$. Set*

$$A_k = \mathbb{F}_p[T]/\Phi_{r^k}(T), \qquad \Delta = (\mathbb{Z}/r^k\mathbb{Z})^*.$$

The action of $u \in \Delta$ sends T to T^u. For a subgroup $H \leq \Delta$, put $B = A_k^H$ and $b = v_r([\Delta : H])$. Then B has dimension $[\Delta : H]$, and every simple field factor has degree $d = \mathrm{ord}_{\Delta/H}(pH)$, with

$$v_r(d) = \max(b - a + 1, 0).$$

In particular, if $r \mid d$, then $\dim_{\mathbb{F}_p} B \geq r^a$ and the number of its simple field factors has r-adic valuation exactly $a-1$. Thus, for $a \geq 2$, no such invariant algebra has the factor-count property, irrespective of its dimension. Every unital $\mathbb{F}_p$-subalgebra of A_k stable under all of Δ is of the form A_k^H and obeys this conclusion.

Proof. After extension to an algebraic closure, A_k is the algebra of functions on Δ, with its regular translation action. Its H-invariants are functions on Δ/H, and Frobenius acts by multiplication by p. The number of cosets gives the dimension, and their common orbit length gives the degree of every field factor.

The group Δ is cyclic, with r-primary subgroup of order r^{k-1}. The r-primary component of p has order r^{k-a}. Relative to a generator γ it therefore has the form $\gamma^{r^{a-1}v}$ with $(v, r) = 1$. The r-primary quotient of Δ/H has order r^b, so that element's image has order $r^{\max(b-a+1,0)}$. If this order is divisible by r, then $b \geq a$, which proves the dimension bound. The number of factors is $[\Delta : H]/d$, whose r-adic valuation is then $b - (b - a + 1) = a - 1$.

Finally, a unital subalgebra stable under Δ becomes, over an algebraic closure, the functions constant on the blocks of a translation-invariant partition of Δ. The block containing the identity is a subgroup H, and the blocks are its cosets: translation invariance gives closure under products and inverses and identifies every other block as a translate. The subalgebra is therefore A_k^H after scalar extension, and hence before scalar extension as well. □

Although $\Phi_{r^{a+1}}(X) = \sum_{j=0}^{r-1} X^{jr^a}$ has a short sparse description, this description alone does not give arithmetic in a smaller explicit algebra. The proposition limits compression by Δ-stable subalgebras, including fixed-subgroup constructions. It leaves open a method that computes a small selected factor from the sparse input and breaks that symmetry.

27.3 Cyclic fields and the auxiliary prime

The conductor of a cyclic field gives an exact comparison with the classical auxiliary-prime search.

Proposition 27.4 (Cyclic fields and the least auxiliary prime)**.** *Let r be an odd prime, let $p \neq r$ be prime, and suppose*

$$p^{r-1} \equiv 1 \pmod{r^2}.$$

Let ℓ_0 be the least rational prime satisfying

$$\ell_0 \neq p, \qquad \ell_0 \equiv 1 \pmod{r}, \qquad p^{(\ell_0-1)/r} \not\equiv 1 \pmod{\ell_0}.$$

Among cyclic degree-r number fields in which p is unramified and inert, the least absolute discriminant is exactly

$$\ell_0^{r-1}.$$

Equivalently, for every real $B \geq 1$, such a cyclic field of absolute discriminant at most B exists if and only if such an auxiliary prime at most $B^{1/(r-1)}$ exists.

Proof. We use the classical conductor facts recorded, with a proof, in Cornelissen et al. (2026, Lemma 8.3.1); the local character argument needed here is also given explicitly below. By Kronecker–Weber, a cyclic degree-r field L is cut out by a primitive Dirichlet character χ of order r, of conductor f. All the characters χ^j, $1 \leq j \leq r-1$, have the same kernel, and therefore the same conductor f. The conductor–discriminant formula gives

$$|D_L| = f^{r-1}.$$

We describe the local factors of f. At a prime $\ell \neq r$, an order-r character is trivial on the principal units, a pro-ℓ group. Thus a nontrivial local factor has conductor ℓ, and its existence requires $r \mid \ell - 1$. At r, the order-r character is trivial on the factor of order $r-1$ in $\mathbb{Z}_r^\times$. The group $1 + r\mathbb{Z}_r$ is procyclic, and its subgroup of rth powers is $1 + r^2\mathbb{Z}_r$. Hence the conductor's r-part is either 1 or r^2. Consequently

$$f = r^{2\varepsilon} \prod_{\ell \in S} \ell, \qquad \varepsilon \in \{0, 1\}, \qquad \ell \equiv 1 \pmod{r} \quad (\ell \in S).$$

The assumption on p says that its class modulo r^2 has order dividing $r-1$. Every local order-r character at r therefore takes the value one at p. Since p is unramified and inert in L, one has $p \nmid f$ and $\chi(p) \neq 1$. At least one local character at a prime $\ell \in S$ must consequently take a nontrivial value at p. Such a character on $(\mathbb{Z}/\ell\mathbb{Z})^\times$ has kernel exactly the subgroup of rth powers. Thus ℓ is an auxiliary prime of the required kind, and

$$f \geq \ell \geq \ell_0, \qquad |D_L| \geq \ell_0^{r-1}.$$

Conversely, the unique degree-r subfield of $\mathbb{Q}(\zeta_{\ell_0})$ is cut out by an order-r character of conductor ℓ_0. The auxiliary-prime condition says precisely that its Frobenius at p is nonidentity. Every nonidentity element of its cyclic Galois group is an r-cycle, so p is unramified and inert. Its discriminant is ℓ_0^{r-1}, proving equality. □

The least prime in the statement exists unconditionally, for example by qualitative Chebotarev in the Kummer splitting field; the proposition imposes no effective bound on it. In the hard branch $a = v_r(p^s - 1) \geq 2$, with $s = \mathrm{ord}_r(p)$, the displayed congruence on p follows since $s \mid r - 1$. For $r = 3$, the least discriminant in this cyclic family is therefore ℓ_0^2. A small complex cubic field with prescribed inertness can vary its imaginary quadratic resolvent, so its discriminant condition is a different sufficient condition.

28 Elliptic curves, isogenies, and missing selectors

This section identifies the precise elliptic data that suffice for the finite-field construction and the selection steps not supplied by the known curve.

28.1 An elliptic-curve alternative

Elliptic curves provide a distinct way to construct irreducible polynomials; Couveignes and Lercier's algorithm uses isogeny fibres, with a randomized search for a suitable curve (Couveignes and Lercier, 2013, §4). The following sufficient condition uses the factorization pattern of a division polynomial. It replaces the auxiliary-prime search by a curve-construction problem, whose uniform deterministic solution is not established here.

Proposition 28.1. *Let $p > r$ be primes, with r odd, and let $E/\mathbb{F}_{p^2}$ be an explicitly given elliptic curve of Frobenius trace t. Suppose*

$$v_r(t^2 - 4p^2) = 1.$$

Let d be the order of p in $\mathbb{F}_r^*/\{1,-1\}$, and put $N = (r-1)/(2d)$. The monic r-division polynomial of E is squarefree and has exactly N irreducible factors of degree d and N of degree rd, with no other factors. From E and the given presentation of $\mathbb{F}_{p^2}$ one can therefore construct a presentation of $\mathbb{F}_{p^r}$ deterministically in time polynomial in r and $\log p$.

In particular, the trace $t = 2p - r$ satisfies the condition, and an ordinary elliptic curve with that trace exists over $\mathbb{F}_{p^2}$. This existence assertion does not give a polynomial-time construction of the curve.

Proof. Write $q = p^2$. The Frobenius and division-polynomial facts used here are the ones recalled in (Schoof, 1985, §2, pp. 485–486). Frobenius acts on the free rank-two r-adic Tate module of E with characteristic polynomial $X^2 - tX + q$. Its reduction F on $E[r]$ has a repeated eigenvalue $\lambda = t/2 \pmod r$, and $\lambda^2 = p^2 \pmod r$. Thus $\lambda = \pm p$ in $\mathbb{F}_r^*$. The reduction cannot be scalar. If the integral Frobenius matrix were $\lambda_0 I + rA$, for an integer lift λ_0 of λ, then its characteristic discriminant would be

$$t^2 - 4q = r^2\big((\operatorname{tr} A)^2 - 4\det A\big),$$

contrary to the hypothesis. Hence $F = \lambda(I + U)$ in a suitable basis, with $U \neq 0$ and $U^2 = 0$.

The roots of the r-division polynomial correspond to nonzero vectors of $E[r] \simeq \mathbb{F}_r^2$ modulo the identification $v \sim -v$. Since $r \neq p$, these are $(r^2-1)/2$ distinct roots. On the one-dimensional space $\ker U$, every nonzero vector has orbit length d modulo sign. There are $(r-1)/2$ such sign pairs. For $v \notin \ker U$, the vectors v, Uv are linearly independent, and

$$F^j v = \lambda^j (v + jUv).$$

The equality $F^j v = \pm v$ holds exactly when $r \mid j$ and $\lambda^j = \pm 1$. Since d divides $(r-1)/2$, the least such positive j is rd. There are $r(r-1)/2$ sign pairs outside $\ker U$. Dividing the two numbers of sign pairs by their respective orbit lengths gives N factors of degree d and N of degree rd.

The integer N is positive and smaller than r. The degree-rd factor count therefore satisfies the property of (Bhargava et al., 2017, Property 1.1 and Theorem 1.4, p. 2). The division polynomial has degree $(r^2-1)/2$ and is computed by the usual division-polynomial recurrences in polynomially many field operations. The factor-count construction gives a degree-r extension of $\mathbb{F}_{p^2}$, hence an explicitly presented field of order p^{2r}, in polynomial time. Compute the kernel of the $\mathbb{F}_p$-linear map $z \mapsto z^{p^r} - z$ in this field. It is the unique subfield of order p^r. Choose a basis element outside its one-dimensional prime field and compute its minimal polynomial; since r is prime, that polynomial has degree r. These are polynomial-time operations in dimension $2r$ over $\mathbb{F}_p$.

For the proposed trace,

$$(2p - r)^2 - 4p^2 = r(r - 4p),$$

whose r-adic valuation is one because $r \neq p$ and r is odd. Also $0 < 2p - r < 2p$ and $p \nmid 2p - r$. Waterhouse's classification, Theorem 4.1(1), therefore gives an ordinary elliptic curve over $\mathbb{F}_{p^2}$ with that trace (Waterhouse, 1969, p. 536). □

The quadratic base extension is useful here. Over $\mathbb{F}_p$, an element of $\mathrm{GL}_2(\mathbb{F}_r)$ whose order is divisible by r must have a nontrivial unipotent part and a repeated eigenvalue λ; its determinant is λ^2. When p is a quadratic nonresidue modulo r, Frobenius on $E[r]$ cannot have this form for any $E/\mathbb{F}_p$. In particular, merely requiring trace 2 modulo r would not produce a unipotent matrix in the residual case $r \nmid p-1$.

The remaining input in Proposition 28.1 is a deterministic polynomial-time method for finding a curve with $v_r(t^2-4p^2)=1$. The trace test itself can be performed by Schoof's algorithm; an efficient test does not bound the number of curves that must be examined. The fixed trace $2p-r$ has characteristic discriminant $-r(4p-r)$, so a direct CM construction has no established uniform polynomial bound in r and $\log p$ from this choice. One could instead seek any trace with the valuation property. Rational Dirichlet estimates control quadratic splitting conditions, but a CM construction also needs the appropriate class-polynomial root and Frobenius behaviour; those requirements are not removed by a quadratic splitting test.

An explicit curve for $p=17$, $r=3$. Let $u^2=3$ over $\mathbb{F}_{17}$ and set

$$E: \quad y^2 = x^3 + x + (1+6u) \qquad \text{over } \mathbb{F}_{289}.$$

The element 3 is a nonsquare modulo 17. The nonsingularity expression is $4+27(1+6u)^2 = 6+u \neq 0$. For $x = A+Bu$, write

$$R = A^3 + 9AB^2 + A + 1, \qquad I = 3A^2B + 3B^3 + B + 6 \qquad \pmod{17}.$$

Then $x^3+x+(1+6u) = R+Iu$, and this nonzero element is a square in $\mathbb{F}_{289}$ precisely when R^2-3I^2 belongs to

$$\{1,2,4,8,9,13,15,16\} \subset \mathbb{F}_{17}.$$

The following table counts the 17 values of A for each B. The three count columns distinguish zero, nonzero square, and nonsquare right-hand sides.

B	zero	square	nonsquare
0	0	7	10
1	0	10	7
2	0	8	9
3	0	8	9
4	0	11	6
5	0	4	13
6	0	7	10
7	0	7	10
8	0	9	8
9	0	7	10
10	0	6	11
11	0	6	11
12	0	9	8
13	0	13	4
14	0	6	11
15	0	8	9
16	0	3	14
total	0	129	160

Including the point at infinity gives $\#E(\mathbb{F}_{289}) = 1+0+2\cdot 129 = 259$. Thus $t = 289+1-259 = 31$ and $t^2 - 4\cdot 289 = -195$, of 3-adic valuation one.

The normalized 3-division polynomial factors as

$$X^4 + 2X^2 + (4+7u)X + 11 = (X + 5 + 7u) \cdot \big(X^3 + (12+10u)X^2 + (4+2u)X + 10 + 3u\big).$$

The cubic factor has no root in $\mathbb{F}_{289}$, as can be checked on its 289 elements; it is therefore irreducible. Let α be its image in the resulting degree-three extension. The trace to the degree-three subfield over $\mathbb{F}_{17}$ is

$$\beta = \alpha + \alpha^{17^3} = (5+13u) + (5+5u)\alpha + (9+5u)\alpha^2.$$

Direct reduction using the cubic gives

$$\beta^{17^3} = \beta, \qquad \prod_{i=0}^{2} (X - \beta^{17^i}) = X^3 + 7X^2 + 2X + 5.$$

The latter polynomial has values

$$5, 15, 11, 16, 2, 9, 9, 8, 12, 10, 8, 12, 11, 11, 1, 4, 9$$

at $0, 1, \ldots, 16$, respectively, so it has no root over $\mathbb{F}_{17}$ and is irreducible. This finite computation illustrates the conditional construction; it gives no uniform bound for locating such a curve.

The numerical elliptic example is checked by *An elliptic cubic extension at seventeen.*

A related elliptic-data interface. Spaenlehauer (2026, Definition 1.2 and Theorem 1.3, p. 2) gives an unconditional deterministic root-finding algorithm when the input includes an elliptic curve, its quadratic twist, their orders and factorizations, and bases of their rational point groups. Its complexity depends on the largest prime factor of these orders. Constructing such data with suitably small prime factors is posed separately in Problem 1.7, p. 3. This conditional-on-data interface does not construct the required elliptic input from the field characteristic alone.

28.2 Rational three-torsion and cubic extraction

A rational torsion point makes the remaining extraction step explicit.

Proposition 28.2 (The Tate family with rational 3-torsion)**.** *Let k be a field of characteristic different from 3 containing a primitive cube root of unity ζ. Let $a, b \in k$ satisfy*

$$bR \neq 0, \qquad R = a^3 - 27b.$$

The elliptic curve

$$E_{a,b} : y^2 + axy + by = x^3$$

has the point $(0,0)$ of order 3, and its 3-division polynomial is

$$\psi_3(X) = Xg(X), \qquad g(X) = 3X^3 + a^2X^2 + 3abX + 3b^2.$$

The polynomial g is irreducible over k if and only if R is a noncube in $k^\times$. If R is a cube, g splits into three distinct linear factors over k.

Proof. The Weierstrass invariants are $b_2 = a^2$, $b_4 = ab$, $b_6 = b^2$, and $b_8 = 0$. They give the discriminant $b^3(a^3 - 27b)$ and the displayed division polynomial. The line $y = 0$ is the tangent at $(0,0)$ and has intersection multiplicity three there; hence this nonidentity point has order three.

Choose c with $c^3 = R$ in an algebraic closure. The three roots of g are

$$x_c = -\frac{a^2 + ac + c^2}{9}, \qquad c^3 = R.$$

Indeed, if $W = a^2 + ac + c^2$, then $27b = (a - c)W$. Substitution of $X = -W/9$ into g gives zero: after taking out the factor $W^2/243$, the remaining expression is

$$-W + 3a^2 - 3a(a - c) + (a - c)^2 = 0.$$

These roots are distinct. If $c \neq c'$ and $x_c = x_{c'}$, then $a + c + c' = 0$. The remaining cube root of R is consequently a, so $a^3 = R$, contrary to $b \neq 0$.

If R is a cube in k, all three values of c lie in k. If R is a noncube, $X^3 - R$ is irreducible, and its splitting field over k is a cyclic cubic extension. Its Galois group permutes the three values of c transitively, and the injective formula above shows that it also permutes the three roots of g transitively. Thus g is irreducible. □

The root correspondence has the explicit inverse

$$c = a + \frac{3b}{x_c};$$

none of the roots x_c is zero because $g(0) = 3b^2 \neq 0$. Thus selecting one of the other three cyclic 3-isogeny kernels from a curve with the distinguished kernel generated by $(0, 0)$ is exactly a cube-root extraction for R, even when R is known to be a cube. The parent curve is determined by quotienting by the distinguished kernel, so it adds no extra generic parameter to these formulas.

For finite $k = \mathbb{F}_q$ containing μ_3, Frobenius on $E_{a,b}[3]$ fixes $(0, 0)$ and has determinant $q \equiv 1 \pmod 3$. It is therefore either the identity or a nontrivial unipotent matrix. The proposition says that the second case is equivalent to the noncube condition on $a^3 - 27b$: its other three pairs $\{P, -P\}$ form one Frobenius orbit. In particular, over $\mathbb{F}_{p^2}$ with $p \equiv -1 \pmod 9$, selecting such a nontrivial unipotent curve from this Tate family retains the original cubic-nonresidue condition. Choosing both a and b in $\mathbb{F}_p$ always gives a cube when $R \neq 0$, because cubing is a bijection on $\mathbb{F}_p^\times$.

28.3 A cubic isogeny descent and its exact factor counts

The following statement concerns the complete collection of descending isogenies. It does not rule out a method that selects one branch by using additional information.

Proposition 28.3. *Let $p > 3$ be a prime with $p \equiv -1 \pmod 9$, and let $E/\mathbb{F}_p$ be ordinary, with a rational point of order 3. Write $t = p + 1 - \#E(\mathbb{F}_p)$ and $v = v_3(t)$, and suppose $\mathrm{End}_{\overline{\mathbb{F}}_p}(E) = O$ has discriminant $D < -4$. For $k \geq 1$, let $C_k(X)$ be the polynomial whose roots are the j-invariants obtained from E by k successive strictly descending 3-isogenies, each counted once. Then $C_k \in \mathbb{F}_p[X]$ is squarefree, of degree $2 \cdot 3^{k-1}$. Every irreducible factor of C_k has degree*

$$d_k = 2 \cdot 3^{\max(k-v,0)},$$

and the number of these factors is

$$N_k = 3^{\min(k,v)-1}.$$

Consequently, whenever $3 \mid d_k$, one has $v_3(N_k) = v - 1$. If $v \geq 2$, no C_k satisfies the degree-3 factor-count property.

Proof. The rational point implies $3 \mid \#E(\mathbb{F}_p)$, and $p \equiv -1 \pmod 3$ therefore gives $3 \mid t$. Ordinarity implies $t \neq 0$. Moreover

$$t^2 - 4p \equiv 1 \pmod 3.$$

Since $\mathbb{Z}[\pi] \subseteq O$ for the Frobenius endomorphism π, the discriminant relation shows that $3 \nmid D$ and that 3 splits in O. Thus E is maximal at 3 and

$$O \otimes \mathbb{Z}_3 \simeq \mathbb{Z}_3 \times \mathbb{Z}_3.$$

For $O_k = \mathbb{Z} + 3^k O$, the conductor exact sequence, with 3 prime to the conductor of O, gives

$$G_k := \ker\big(\mathrm{Cl}(O_k) \longrightarrow \mathrm{Cl}(O)\big) \simeq (O/3^kO)^*/(\mathbb{Z}/3^k\mathbb{Z})^* \simeq (\mathbb{Z}/3^k\mathbb{Z})^*.$$

The unit correction is trivial because $D < -4$. In particular $|G_k| = 2 \cdot 3^{k-1}$; see the exact sequence and class-number formula in (Bröker et al., 2012, §3, (3) and (5)).

Here is also why this kernel counts distinct descendants. In characteristic zero, extension of ideals gives the map from O_k-CM classes to O-CM classes. The fibre over the canonical CM lift of E is a G_k-torsor. Locally at 3,

$$O_k \otimes \mathbb{Z}_3 = \{(x, y) \in \mathbb{Z}_3^2 : x \equiv y \pmod{3^k}\}.$$

Increasing k gives the unique chain of orders reached by strict descent. Equivalently, the standard ideal-lattice description of a 3-isogeny gives two descending neighbours at the split maximal order, and three at each subsequent level, each with a unique ascending neighbour. Thus the above fibre is precisely the set of endpoints of the strictly descending chains, rather than a multiset of paths; see (Bröker et al., 2012, §4.1). The ordinary CM reduction theorem applies because p splits in the CM field and $p \nmid 3D$. The CM values of the same order remain distinct on reduction, by (Cremona and Sutherland, 2024, §2.4, Proposition 4). It follows that the $2 \cdot 3^{k-1}$ descendants remain distinct.

Write $\pi = (\alpha, \beta)$ in the displayed 3-adic product. Then $\alpha + \beta = t$ and $\alpha\beta = p$, and both entries are units. In the quotient G_k the Frobenius is the class of $\rho = \alpha/\beta$, up to inversion. Its image in $\mathrm{Cl}(O)$ is trivial because the Frobenius prime ideal is generated by π. Hence it acts on the fibre by translation by this one element. We have

$$\rho \equiv -1 \pmod 3, \qquad v_3(\rho + 1) = v_3(t/\beta) = v.$$

The lifting-the-exponent identity therefore gives

$$\mathrm{ord}_{(\mathbb{Z}/3^k\mathbb{Z})^*}(\rho) = 2 \cdot 3^{\max(k-v,0)}.$$

All translation orbits have this length. Dividing $|G_k|$ by this length gives N_k, proves Frobenius stability of the root set, and proves all the asserted factor counts. □

The same obstruction survives subgroup compression of this whole fibre. Indeed, for $H \leq G_k$, put $b = v_3([G_k : H])$. In G_k/H the 3-part of the Frobenius order is $3^{\max(b-v+1,0)}$. If this order is divisible by 3, the number of Frobenius orbits has valuation exactly $v - 1$. This follows directly by writing G_k as its cyclic 2-part times its cyclic 3-part and observing that the Frobenius exponent in the latter has valuation $v - 1$. It applies to the invariant algebra $(\mathbb{F}_p[X]/C_k)^H$ and does not presume that an arbitrary selected factor is such an invariant algebra.

When $v = 1$, there is an effective short descent. The two first descendants form an irreducible quadratic factor of $\Phi_3(j(E), X)$, obtainable by distinct-degree gcds; horizontal roots can be removed without choosing them. Use this quadratic as a presentation of $\mathbb{F}_{p^2}$, and denote its class of X by j_1. At this first level,

$$\Phi_3(j_1, X)/(X - j(E))$$

is irreducible of degree 3 over $\mathbb{F}_{p^2}$. Its quotient presents $\mathbb{F}_{p^6}$, and the kernel of $\mathrm{Frob}_p^3 - 1$ gives its degree-3 subfield by linear algebra. All these operations take polynomial time in $\log p$. If $v \geq 2$, the same quotient at the first level is a split cubic, and selecting one of its roots is additional work not performed by the complete-fibre construction.

28.4 What a known parent supplies in cubic descent

The next calculation separates a known isogeny from a choice of its continuation. Suppose $p \equiv -1 \pmod 9$, and an ordinary curve $E_0/\mathbb{F}_p$ has rational three-torsion and nonzero trace divisible by 9. After the first descending isogeny $\phi : E_0 \longrightarrow E_1$ over $k = \mathbb{F}_{p^2}$, assume the three-torsion of E_1 is rational, as in the intermediate levels of Proposition 28.3. Its known ascending kernel can be normalized to $(0, 0)$ in $y^2 + axy + by = x^3$. The normalization uses translation and its tangent line, and requires no cubic root.

By Proposition 28.2, selecting one of the three other kernels is equivalent, by constant-size algebraic formulas, to selecting a root of $C^3 = a^3 - 27b$. The parent is recovered by quotienting by the known kernel, so these generic parent data do not add an independent parameter choosing a root. Over $k(a, b)$ the radicand has valuation one at the divisor $a^3 - 27b = 0$ and is not a cube. Consequently there is no rational function in the generic parameters that makes this choice. This assertion does not exclude algorithms using finite-field Frobenius or the special descent sequence.

There is an explicit Frobenius action, but its effect is a permutation of the choices. Write π for the p-Frobenius of E_0 and $C_0 = \ker \phi$. At the first descent, $\pi(C_0) \neq C_0$. On $E_1[3]$, the computable endomorphism $\alpha = \phi\pi\widehat{\phi}$ is nonzero of rank one: the image of $\widehat{\phi}$ is C_0, and ϕ is nonzero on $\pi(C_0)$. Moreover $\alpha^2 = [3]\phi\pi^2\widehat{\phi} = 0$ there. Thus its kernel and image are the ascending line, and $1 + \alpha$ cyclically permutes the three other lines. The coordinate algebra already has this symmetry in the form

$$k[X]/(3X^3 + a^2X^2 + 3abX + 3b^2) \simeq k[C]/(C^3 - a^3 + 27b), \qquad C = a + 3b/X,$$

where $C \mapsto \zeta_3 C$ rotates the roots. Mapping the known points of $E_0[3]$ through ϕ yields only $\ker \widehat{\phi}$, because $\widehat{\phi}\phi = [3]$. A nonparent continuation instead corresponds to a cyclic subgroup of order nine above C_0; its preimage cannot be $E_0[3]$, whose image is exactly the excluded ascending line.

Theorem 28.4 (Selection on an ordinary descending chain)**.** *Let $p > 3$ be prime with $p \equiv -1 \pmod 9$, and present $k = \mathbb{F}_p[\zeta]$, where $\zeta^2 + \zeta + 1 = 0$. Suppose an ordinary elliptic curve $E_0/\mathbb{F}_p$ and a specified point of order three in $E_0(\mathbb{F}_p)$ are given. Consider the following restricted oracle: on a descending 3-isogeny chain from E_0, after the first descent and strictly above the floor, choose one of the three cyclic kernels other than the known ascending dual kernel. The current curve, incoming isogeny, and a generator of the dual kernel are given.*

Relative to this given ordinary starting curve, this oracle and extraction of a cube root of an arbitrary promised cube in k are polynomial-time oracle equivalent. More precisely, if t_0 is the trace of E_0 and $v = v_3(t_0)$, at most $v - 1$ restricted oracle calls produce a maximal 3-primary unit of $k^\times$, and hence solve every promised cube-root instance. All remaining work is deterministic and polynomial in $\log p$.

Proof. The rational point gives $3 \mid p + 1 - t_0$, hence $3 \mid t_0$. Since E_0 is ordinary and $p > 3$, $t_0 \neq 0$. The Frobenius polynomial $X^2 - t_0X + p$ has discriminant congruent to 1 modulo three. It therefore has distinct roots λ_+, λ_- in $\mathbb{Z}_3$, reducing to $1, -1$; its local endomorphism order is maximal and split. Moreover

$$v_3(\lambda_+^2 - \lambda_-^2) = v_3((\lambda_+ - \lambda_-)t_0) = v.$$

Here is the local lattice calculation for the descent, including the termination assertion. Up to independent scaling in the two eigenlines, a lattice of local conductor 3^i is

$$L_i = \mathbb{Z}_3(1,1) + \mathbb{Z}_3(0,3^i).$$

For completeness, normalize both coordinate projections of any lattice to $\mathbb{Z}_3$. Its two coordinate kernels have the same index 3^i, and the induced identification modulo 3^i is multiplication by a unit. Rescaling one coordinate gives exactly the displayed lattice. In its displayed basis, the p^2-Frobenius matrix is

$$\begin{pmatrix} \lambda_+^2 & 0 \\ (\lambda_-^2 - \lambda_+^2)/3^i & \lambda_-^2 \end{pmatrix}.$$

The reduction is the identity for $i < v$, and is nontrivial unipotent for $i = v$.

At $i > 0$, put $u = (1,1)$ and $w = (0,3^i)$. Among the four index-three superlattices, adjoining $w/3$ gives L_{i-1}. The other three adjoin $(u + cw)/3$, $c = 0,1,2$; normalizing their two projections as above gives conductor 3^{i+1}. Thus there is exactly one ascending kernel and every other kernel descends. At $i = 0$ the two Frobenius eigenlines are the horizontal kernels. The two remaining kernels are conjugate over $\mathbb{F}_p$ and rational over k, and either gives the first descent. These statements concern kernels, so coincident j-invariants caused by automorphisms do not change the argument.

The first descent is explicit. Move the given point and its tangent to $(0,0)$ and $y = 0$, obtaining $y^2 + axy + by = x^3$. Since cubing is bijective on $\mathbb{F}_p$, scale by the uniquely computable cube root of b to obtain $E_K : y^2 + Kxy + y = x^3$. With $c^3 = K^3 - 27$ in $\mathbb{F}_p$ and $d = \zeta c$, the point

$$Q = \left(-\frac{3}{K-d}, \frac{K-c}{K-d} \right)$$

generates a descending kernel; the identities are verified in Proposition 28.5 below. Its x-coordinate is not in $\mathbb{F}_p$, so it is not a horizontal kernel. Quotienting by this kernel is explicit. The image of the original rational point generates the dual kernel: its line is different from the selected kernel, and the image of $E[3]$ under a 3-isogeny is its dual kernel. The same argument tracks an explicit dual generator at every later step. Degree-three quotient maps can be computed from the kernel x-coordinate alone.

At every intermediate depth all four kernel lines are rational over k. At depth v, Frobenius fixes only the ascending line and cyclically permutes the other three. Hence the cubic for their x-coordinates is irreducible. No trace computation is needed: at each stage normalize at the known dual generator to $y^2 + axy + by = x^3$ and test whether $R = a^3 - 27b$ is a cube. The elementary Tate formula

$$x_z = -\frac{a^2 + az + z^2}{9}, \qquad z^3 = R, \qquad z = a + \frac{3b}{x_z}$$

identifies its three remaining kernel coordinates with the cube roots of R. Thus a noncube is detected exactly at the floor. Before the floor the restricted oracle makes the next descent. There are at most $v - 1$ such calls, and Hasse's bound gives $v \le \log_3(2\sqrt{p})$.

Write $p^2 - 1 = 3^a m$, with $3 \nmid m$. At the floor the noncube R gives $\eta = R^m$ of exact order 3^a. To extract a root of a nonzero promised cube γ, choose u with $3u \equiv 1 \pmod{m}$ and write $3u - 1 = jm$; for $m = 1$ take $u = 0$. Then $x_0 = \gamma^u$ has error

$$\gamma / x_0^3 = \gamma^{-jm} \in \langle \eta \rangle.$$

Its discrete logarithm to base η is divisible by three. Successive base-three digits compute this logarithm in $O(a^2)$ group operations. If it is e with $0 \le e < 3^a$, then $x_0 \eta^{e/3}$ is a cube root of γ. The zero input is immediate. Conversely, a general promised cube-root oracle applied

to R in the displayed Tate formula selects a nonparent kernel at any intermediate stage. All formulas have bounded degree, and exponentiation and field arithmetic require polynomial time in $\log p$. □

The ordinary starting curve is an additional input in this theorem. The theorem does not construct such an input from p alone and does not establish a smaller sufficient analytic hypothesis for finite-field construction.

Proposition 28.5 (The first continuation instance)**.** *Under the same conditions on p and k, let $E_K : y^2 + Kxy + y = x^3$, where $K \in \mathbb{F}_p \setminus \{3\}$. Put $\delta = \zeta - \zeta^2$, let $c \in \mathbb{F}_p$ be the unique cube root of $K^3 - 27$, and put $d = \zeta c$. The point Q in the preceding proof has order three. Moving Q and its tangent to the origin gives Tate coefficients*

$$a_0 = \frac{K + 2d}{\delta}, \qquad b_0 = \frac{\delta d}{K - d}, \qquad a_0^3 - 27b_0 = \left(\frac{K - d}{\delta}\right)^3.$$

For $t = K/c$, the resulting first continuation problem is cube-root extraction for

$$b_0(t) = \frac{\delta\zeta}{t - \zeta}, \qquad t \in \mathbb{F}_p \setminus \{1\}.$$

Every such t occurs. Every nonzero promised cube in k can be transformed, by deterministic polynomial-time field operations and a known cube multiplier, to a scalar multiple of one of these values. If its associated E_K is ordinary, one valid restricted selector call extracts its cube root. Supersingular associated curves are not covered by that oracle reduction.

Proof. Since $K^3 - d^3 = 27$, the denominators are nonzero, and

$$-\frac{3}{K - d} = -\frac{K^2 + Kd + d^2}{9}.$$

The latter is the Tate 3-torsion formula. Substituting $y = (K - c)/(K - d)$ verifies the curve equation. Its tangent slope m satisfies $K + 2m = (K + 2d)/\delta$, while $2y + Kx + 1 = \delta d/(K - d)$, giving a_0, b_0. Using $\delta^2 = -3$ and $(K - d)(K^2 + Kd + d^2) = 27$ proves the remaining cube identity. Conversely, for $t \neq 1$ take

$$c = \left(\frac{27}{t^3 - 1}\right)^{1/3} \in \mathbb{F}_p, \qquad K = tc.$$

The curve's j-invariant, useful for identifying the exceptional supersingular cases, is

$$j(E_K) = \frac{27t^3(t^3 + 8)^3}{(t^3 - 1)^3}.$$

For any nonsingular Tate curve $E : y^2 + axy + by = x^3$, the quotient by $(0, 0)$ has equation

$$E' : y^2 + axy + by = x^3 - 5abx - (a^3b + 7b^2)$$

and division polynomial

$$\psi_{3,E'}(X) = (3X + a^2)(X^3 - 9abX - b(a^3 + 27b)).$$

The linear factor is the dual kernel. The other roots correspond bijectively to cube roots of b through

$$X = az + 3z^2, \quad z^3 = b, \qquad z = \frac{aX + 9b}{a^2 + 3X}.$$

Substitution verifies both identities; the denominator cannot vanish because the cubic at $-a^2/3$ is $-(a^3-27b)^2/27 \neq 0$. Applying this with a_0, b_0 proves the assertion about the first continuation.

For the arbitrary-input reduction write $\gamma = A+B\zeta$. If $B = 0$, its root is uniquely computable in $\mathbb{F}_p$. Otherwise, when $A \neq 2B$ and $A \neq B$, put

$$t = \frac{B-2A}{A-2B}, \qquad s = \frac{A^2-AB+B^2}{A-2B}.$$

Then $s \in \mathbb{F}_p^\times$, $t \neq 1$, and direct multiplication gives $\gamma = sb_0(t)$. All scalars are cubes in k, so $b_0(t)$ is a promised cube, and multiplication by the prime-field cube root of s transfers roots back. The two excluded projective classes are those of $\delta\zeta = -2-\zeta$ and ζ^2. Multiplication by the known cube $(1-\zeta)^3 = -3\delta$ sends them to the classes of ζ and $1-\zeta$, respectively, neither excluded. After extracting a root, divide by $1-\zeta$.

If E_K is ordinary, a cubic $b_0(t)$ means that the first descending curve is above its floor, so the corresponding selector call satisfies the oracle's promise. Without ordinarity this inference is unavailable. The $j = 0$ cases $t = 0, -2$ themselves have explicit roots: the latter is scalar, and for the former γ is a scalar multiple of δ, whose root is $\delta/\sqrt[3]{-3}$ in view of $\delta^3 = -3\delta$. Other supersingular cases remain outside the proved reduction. □

The coordinate, tangent, quotient-root, and arbitrary-input identities in the proposition are checked by exact arithmetic, for every admissible K and every nonzero cube at $p = 17, 53, 71, 107$, by the dependency-free program *Exact identities for restricted cubic isogeny descent*. These finite checks supplement the identities; they do not supply an asymptotic ordinary-start construction.

28.4.1 An ordinary curve after extending the base

The ordinary-start condition can be solved over a quartic extension, although the cubic selector remains a separate operation.

Proposition 28.6 (An ordinary Tate curve over a quartic extension)**.** *Let $p > 3$ be prime with $p \equiv 2 \pmod 3$. There is an unconditional deterministic algorithm, polynomial in $\log p$, that constructs $L = \mathbb{F}_{p^4}$ and an ordinary curve E/L with a specified L-rational point of order three. After constructing a quadratic presentation*

$$L = k[U], \qquad k = \mathbb{F}_{p^2}, \qquad U^2 = \eta, \quad \eta \in k^\times \setminus (k^\times)^2,$$

one of the two curves

$$E_A : y^2 + Axy + y = x^3, \qquad A = U \text{ or } 2U,$$

has $j(E_A) \notin k$ and therefore is ordinary. In both cases the point $(0,0)$ has order three.

Proof. The element $d = -3$ is a quadratic nonresidue modulo p, so $k = \mathbb{F}_p[\delta]$, $\delta^2 = d$, is explicit. Tonelli–Shanks with this known nonresidue extracts every promised square root in $\mathbb{F}_p$ in polynomial time. Promised square roots in k reduce to the same operation. Indeed, for $z = A + B\delta$ with $B \neq 0$, take a square root h in $\mathbb{F}_p$ of $A^2 - dB^2$. The two elements $(A+h)/2, (A-h)/2$ have product $dB^2/4$, a nonsquare, so exactly one is a square. If $x^2 = (A+h)/2$ is the square choice, then

$$x + \frac{B}{2x}\delta$$

is a square root of z. Changing the sign of h selects the other choice when necessary. For $B = 0$, use a prime-field square root of A if it exists, and otherwise use $\delta\sqrt{A/d}$.

Starting from $-1 \in k$, repeatedly take a square root while the current element is a square in k. Its order, initially two, doubles at each step. After at most $v_2(p^2-1)-1$ steps the result η is a nonsquare. This constructs L and supplies $U^{p^2} = -U$ without any unproved search bound.

For a nonzero A with $A^{p^2} = -A$, neither $A^3 - 27$ nor $A^3 + 27$ vanishes. The displayed Tate curve is nonsingular, has the flex $(0, 0)$, and has invariant

$$j(A) = \frac{A^3(A^3-24)^3}{A^3-27}.$$

Coefficient conjugation sends $j(A)$ to $j(-A)$. Writing $u = A^3$, direct expansion gives

$$j(A) - j(-A) = \frac{2u(u^4 - 216u^2 - 373248)}{u^2 - 729}.$$

Thus $j(A) \in k$ if and only if

$$(A^6)^2 - 216A^6 - 373248 = 0.$$

Put $x = U^6 = \eta^3$. This is a nonsquare of k and hence does not belong to $\mathbb{F}_p$, since every nonzero element of $\mathbb{F}_p$ is a square in k. If both trial curves had j in k, then x and $64x$ would both be roots of $X^2 - 216X - 373248$. They are distinct: $p \neq 3, 7$, because $p \equiv 2 \pmod 3$. Their sum would give

$$65x = 216.$$

If $p \mid 65$, this is impossible since $p > 3$; otherwise it would put x in $\mathbb{F}_p$, again impossible. At least one trial therefore has $j \notin k$.

Finally, a supersingular elliptic curve in characteristic p has purely inseparable multiplication-by-p of degree p^2. It factors as relative Frobenius squared followed by an isomorphism, so the curve and its p^2-Frobenius conjugate have the same j-invariant. Every supersingular j therefore lies in $\mathbb{F}_{p^2}$. The selected curve is ordinary. All operations above are field arithmetic, exponentiation, or promised quadratic extraction in fields of bounded degree. □

This construction supplies an ordinary starting curve over $\mathbb{F}_{p^4}$, with explicit three-torsion. It does not produce a prime-field ordinary parent for Theorem 28.4. Nor does ordinary status imply that a given nonparent cubic is irreducible: the scalar-Frobenius intermediate case can still require a split cubic selector. A noncube found over $\mathbb{F}_{p^4}$ would descend to one over $\mathbb{F}_{p^2}$ by the relative norm, but the construction does not assert that either trial already provides such a noncube.

Example 28.7. Ordinarity of the two trial curves does not ensure a noncubic radicand. Take $p = 53$, write $k = \mathbb{F}_{53}(\zeta)$ with $\zeta^2 + \zeta + 1 = 0$, and put $\eta = 1 + 17\zeta$. Then $\eta^{1404} = -1$, so $L = k(U)$, $U^2 = \eta$, is quadratic. For $A = U, 2U$ and $B = 1$, respectively, the values of A^6 are $19 + 32\zeta$ and $50 + 34\zeta$. Substitution into the ordinarity polynomial of Proposition 28.6 gives

$$(A^6)^2 - 216A^6 - 373248 = 34 + 11\zeta, \quad 9 + 41\zeta.$$

Both are nonzero, so both j-invariants lie outside k and both curves are ordinary. On the other hand,

$$N_{L/k}(U^3 - 27) = 21 + 21\zeta, \qquad N_{L/k}(8U^3 - 27) = 43 + 19\zeta.$$

Each of these elements has 936th power one in k, where $936 = (53^2 - 1)/3$. The norm map induces an isomorphism of the cubic-class groups of L and k: for a generator of L^* its norm is a generator of k^*. Thus both Tate radicands are cubes in L. All displayed equalities follow by multiplication using $\zeta^2 = -\zeta - 1$ and repeated squaring; the program *Two ordinary Tate trials* records these powers and the quadratic-field extraction checks.

The fact that supersingular invariants lie in the quadratic field is classical (Silverman, 2009, Theorem V.3.1(a)(iii), pp. 145–146). For ordinary curves over the prime field, Hashimoto and Nuida (Hashimoto and Nuida, 2026, Theorem 3.10 and Corollary 3.11, p. 12) identify the surface position and first descending Hessian three-isogeny. Their supersingularity test takes a cubic nonresidue in $\mathbb{F}_{p^2}$ as auxiliary input (Hashimoto and Nuida, 2026, Section 4, pp. 13–14); Algorithm 2, step 4(c), subsequently extracts cubic roots there. These statements explain the geometric descent but do not construct the missing nonresidue. The prescribed-radicand construction below isolates the remaining root-selection input on an ordinary curve over a larger field.

Proposition 28.8 (Infinitely many failures of the two noncube trials)**.** *Fix $a \geq 2$. There are infinitely many primes p satisfying*

$$p \equiv 5 \pmod{48}, \qquad v_3(p+1) = a,$$

such that, for every primitive sixteenth root of unity $U \in \mathbb{F}_{p^4}$ and every $t \in \{1, 2\}$,

$$(tU)^3 - 27 \in (\mathbb{F}_{p^4}^{\times})^3.$$

For these primes U^2 is a nonsquare in $\mathbb{F}_{p^2}$, and at least one of the two curves $y^2+tUxy+y = x^3$ is ordinary by Proposition 28.6. Thus neither trial supplies a noncube, regardless of the square-root choices in the two-primary tower construction.

Proof. Put $L = \mathbb{Q}(\zeta_{48})$ and let τ be the automorphism $\zeta_{48} \mapsto \zeta_{48}^5$. It has order four. Let $\mathcal{A}$ consist of all numbers

$$t^3\xi^3 - 27, \qquad t \in \{1, 2\}, \quad \xi \text{ primitive of order } 16.$$

Every number in $\mathcal{A}$ is nonzero and integral. The set is stable under $\mathrm{Gal}(L/\mathbb{Q})$. Consequently

$$M = L(\alpha^{1/3} : \alpha \in \mathcal{A})$$

is Galois over $\mathbb{Q}$, and $\mathrm{Gal}(M/L)$ has exponent dividing three.

Choose any lift h of τ to M. Since h^4 belongs to that exponent-three kernel, $h^{12} = 1$. The element $g = h^{-3}$ still restricts to τ, and $g^4 = 1$. In particular g has order four.

Let $C = \mathbb{Q}(\zeta_{3^{a+1}})$. The extension LC/L is cyclic of degree 3^a. Its intersection with M has Galois group of exponent at most three, hence

$$M \cap LC \subseteq L\mathbb{Q}(\zeta_9).$$

Intersecting with C gives $M \cap C \subseteq \mathbb{Q}(\zeta_9)$; the reverse inclusion of $\mathbb{Q}(\zeta_3)$ is automatic. On this intersection, g acts by complex conjugation. Indeed its action has order dividing four and induces conjugation on $\mathbb{Q}(\zeta_3)$, while the only such lift to $\mathbb{Q}(\zeta_9)$ has order two and is conjugation.

The automorphism σ of C corresponding to

$$-1 + 3^a \pmod{3^{a+1}}$$

also restricts to conjugation on $\mathbb{Q}(\zeta_9)$, since $a \geq 2$. Thus g and σ agree on $M \cap C$ and extend to an automorphism γ of MC. The qualitative Chebotarev theorem supplies a positive-density set of rational primes with Frobenius conjugate to γ; see (Serre, 1981, Section 2.1, Theorem 1, p. 131). Exclude the finitely many ramified primes and primes dividing the norms of numbers in $\mathcal{A}$.

For each remaining prime the cyclotomic restrictions give $p \equiv 5 \pmod{48}$ and $p \equiv -1 + 3^a \pmod{3^{a+1}}$, hence $v_3(p+1) = a$. The residue degree of a prime of L over p is four, and its Frobenius in M/L is the fourth power of a conjugate of g, hence is trivial. Thus every element

of $\mathcal{A}$ has a cube root in the residue field $\mathbb{F}_{p^4}$. Reduction sends the primitive sixteenth roots in L bijectively to all primitive sixteenth roots in that field, proving the claim.

Finally $p \equiv 5 \pmod 8$ gives $v_2(p^2-1) = 3$. For every such U, the element U^2 has order eight in $\mathbb{F}_{p^2}^\times$ and is therefore nonsquare. Repeated square roots of -1 in that quadratic field produce a primitive eighth root, and either square root of it in the quartic field is primitive of order sixteen. All possible choices were included in $\mathcal{A}$. □

This proposition concerns a fixed algebraic family and every fixed exact valuation $a \geq 2$. It gives no lower bound for a number of trials growing with $\log p$, and it does not rule out an adaptive construction using additional information about p.

Proposition 28.9 (An ordinary Tate curve with prescribed radicand)**.** *Let $p > 3$ be prime with $p \equiv 2 \pmod 3$, let $k = \mathbb{F}_{p^2}$, and present $L = k[U]$ with U^2 a nonsquare in k. Given any $\gamma \in k^\times$, for each $A \in \{U, 2U\}$ put*

$$B = \frac{A^3 - \gamma}{27}, \qquad E_A : y^2 + Axy + By = x^3.$$

Both curves are nonsingular, and at least one has $j(E_A) \notin k$, hence is ordinary. Its Tate radicand is exactly $A^3 - 27B = \gamma$.

Consequently a single oracle call selecting a cyclic order-three kernel other than the distinguished kernel generated by $(0,0)$ on a presented ordinary Tate curve over L, under the promise that its remaining cubic splits, suffices to extract a cube root of any promised cube in k. All preprocessing is unconditional and polynomial in $\log p$.

Proof. Write a bar for the nontrivial automorphism of L/k. Then $\bar{A} = -A$. Since $\gamma \in k^\times$, neither $A^3 - \gamma$ nor $A^3 + \gamma$ vanishes. Thus $B \neq 0$ and the discriminant $B^3\gamma$ is nonzero. The point $(0,0)$ has order three.

Substitution in the Tate invariants, with $u = A^3/\gamma$, gives

$$j(E_A) = J(u) := \frac{27u(u+8)^3}{(u-1)^3}, \qquad \overline{J(u)} = J(-u).$$

Expanding the difference gives the identity

$$J(u) - J(-u) = \frac{54u(u^2+8)(u^4+259u^2+64)}{(u^2-1)^3}.$$

Here $u \neq 0$ and $u^2 \neq 1$. Put

$$x = \frac{U^6}{\gamma^2} \in k^\times.$$

Since U^2 is nonsquare in k, its cube divided by a square is nonsquare. In particular $x \notin \mathbb{F}_p$, and the same holds for $64x$. The factor $u^2 + 8$ cannot vanish at either trial, because $-8 \in \mathbb{F}_p$.

If both trial invariants were in k, then x and $64x$ would both be roots of

$$X^2 + 259X + 64.$$

They are distinct since $p \nmid 63$. Their sum would imply $65x = -259$. If $p \mid 65$, this contradicts $p \nmid 259$; otherwise it puts x in $\mathbb{F}_p$. Both alternatives are impossible. Thus one invariant lies outside k, and the supersingular-invariant theorem makes that curve ordinary.

For a promised cube γ in k, the remaining Tate division cubic on this selected curve splits over L. For any root coordinate X selected by the oracle, the explicit Tate inverse gives

$$z = A + \frac{3B}{X}, \qquad z^3 = \gamma.$$

The denominator is nonzero because the constant term of the remaining division cubic is $3B^2 \neq 0$. All three cube roots of γ already lie in k, since $\mu_3 \subset k$. Therefore the returned z belongs to the embedded copy of k; extracting its coordinates in that subfield is immediate. The field construction, invariant tests, and formulas have the claimed deterministic complexity by Proposition 28.6. □

The oracle here allows ordinary inputs over $\mathbb{F}_{p^4}$ and is not required to recognize an incoming edge as ascending. It is therefore a different domain from the restricted prime-field parent oracle in Theorem 28.4. The proposition removes the supersingular exception for this ordinary-input root-selection formulation; it does not construct a selector and does not establish a uniform Dirichlet-GRH-only field algorithm.

28.5 An ascending selector over the quartic field

The explicit ordinary curve constructed above removes the starting-curve input when the selector is allowed over the quartic field. The precise promise and its use at the volcano floor matter.

Proposition 28.10 (The ascending-parent selector over the quartic field)**.** *Let $p > 3$ be prime with $p \equiv -1 \pmod 9$, and put $k = \mathbb{F}_{p^2}$ and $L = \mathbb{F}_{p^4}$. Suppose a deterministic algorithm selects a cyclic order-three kernel other than a specified ascending kernel whenever its input is an ordinary elliptic curve over L, strictly above its 3-isogeny volcano floor, with all its three-torsion rational over L. The input includes a rational generator of the ascending kernel, equivalently the dual kernel of a known descending incoming isogeny. Assume a stated polynomial time bound on these promised inputs.*

Then one can deterministically construct an irreducible cubic over $\mathbb{F}_p$, a maximal three-primary unit in k, and cube roots of arbitrary promised cubes in L, in time polynomial in $\log p$ using that algorithm. No ordinary starting curve is an additional input. Conversely, promised cube extraction in L implements the stated selector.

The reduction runs a polynomial-time total completion of the selector: each invocation is clocked at its promised time bound, and every returned kernel is checked. No behavior off the geometric promise is assumed.

Proof. Use Proposition 28.9 with $\gamma = 1$. This constructs an ordinary curve

$$E : y^2 + Axy + By = x^3 \quad \text{over } L, \qquad A^3 - 27B = 1.$$

Its four three-isogeny kernels are explicit: one is generated by $(0, 0)$, and the other three have abscissas

$$-\frac{A^2 + Az + z^2}{9}, \qquad z \in \{1, \zeta, \zeta^2\}.$$

All kernel lines are rational. Since $\mu_3 \subset L$ and $(0, 0)$ is a rational point, Frobenius acts as the identity on $E[3]$, so their generators are rational as well. Their ordinates require only promised square extraction in the already constructed quadratic tower. That extraction is deterministic by the same norm-and-two-square-roots formulas used to construct the tower.

Start all four outgoing isogeny branches. At least one first edge descends: at a maximal-order vertex there are at most two horizontal kernels, and below it there is one ascending kernel and three descending ones. On a branch whose first edge descends, the dual of the incoming edge is the unique ascending parent thereafter. Every nonparent choice continues down. These general ordinary-volcano facts hold in the split, inert, and ramified local cases; exact sources are (Sutherland, 2013, Section 2.9, p. 514; Lemma 6, p. 515; Theorem 7 and Remark 8, pp. 516–517). We count cyclic kernel lines, retaining distinct edges whose target j-invariants coincide. The initial invariant lies outside $\mathbb{F}_{p^2}$ and is therefore neither 0 nor 1728. After a descent the endomorphism order has conductor divisible by three, so no vertex on that descending branch

has either exceptional invariant. The multiplicities at exceptional surface vertices described in the cited Remark 8 therefore do not affect this branch. Parallel branching to find a descending path is the classical FindShortestPathToFloor mechanism (Sutherland, 2013, Section 3.1, pp. 519–520); here the prescribed-radicand preprocessing gives the starting curve and the stated selector supplies the root selections.

After every quotient, keep a rational generator of the dual kernel. At the initial step it is the image of a known independent three-torsion point. At later steps it is the image of the previous parent generator, since the chosen kernel is different. A linear coordinate change moving this point and its tangent to the origin gives a nonsingular Tate equation $y^2 + axy + by = x^3$ with the parent generated by $(0, 0)$.

Test $R = a^3 - 27b$ before calling the selector. If R is a noncube in L, keep it and stop. If it is a cube, the remaining division cubic splits. Invoke the clocked selector, check that its output is a three-kernel distinct from the parent, and continue. Discard a branch upon timeout or an invalid output. On the truly descending branch all calls strictly before the floor satisfy the promise, so that branch is never discarded.

The height is $O(\log p)$. More explicitly, if t is the L-trace, the conductor of $\mathbb{Z}[\pi]$ has square dividing $t^2 - 4p^4$, whose absolute value is at most $4p^4$. Thus the height is at most $\frac{1}{2}\log_3(4p^4)$. A fixed larger integer bound is used as the clock for every branch. At a floor reached by descent there is only one rational kernel (Sutherland, 2013, Proposition 10, p. 518). Frobenius fixes the known rational parent point and has determinant $p^4 \equiv 1 \pmod 3$ on three-torsion. Both eigenvalues are therefore one. It cannot be the identity, which would make all four kernels rational. Thus it is nontrivial unipotent, fixing the ascending line and cycling the other three. The corresponding remaining cubic is irreducible and R is a noncube. Thus some branch stops successfully within the bound, independently of every off-promise answer on the other branches. A successful noncube test is itself a complete algebraic certificate for the subsequent output.

Put $u = N_{L/k}(R)$. Then

$$u^{(p^2-1)/3} = R^{(p^4-1)/3} \neq 1,$$

so u is a noncube in k. If $a = v_3(p+1)$, then $u^{(p^2-1)/3^a}$ has exact order 3^a. Its image in L has the same order, which is maximal because $v_3(p^4-1) = a$. Successive base-three digit extraction therefore computes roots of every promised cube in L.

For a direct cubic output, let

$$n = N_{k/\mathbb{F}_p}(u), \qquad v = n^{1/3} \in \mathbb{F}_p, \qquad f(X) = X^3 - 3vX - \mathrm{Tr}_{k/\mathbb{F}_p}(u).$$

The prime-field cube root is unique and computable by exponentiation. If $c^3 = u$, the three roots of f are $\zeta^j c + v/(\zeta^j c)$, for $j = 0, 1, 2$. They are distinct: equality for two roots would imply $u^2 = n$, hence $u = \overline{u}$, which is impossible for a noncube in k because every nonzero prime-field scalar is a cube there. The cyclic extension $k(c)/k$ acts transitively on these three roots. Thus f is irreducible over k and consequently over $\mathbb{F}_p$.

Conversely, on any promised selector input move the given parent point to Tate form. The split remaining cubic has the displayed Tate roots with $z^3 = a^3 - 27b$. One promised cube extraction supplies such a z, and hence a kernel different from the parent. All transformations and verification have polynomial complexity. □

This proposition concerns ordinary inputs over $\mathbb{F}_{p^4}$ and permits the explicitly specified clocked total completion. It does not apply the original prime-field-parent oracle to a new domain without justification. It proves an equivalence for the enlarged selector problem and leaves the original uniform construction question unchanged.

28.6 Complementary torsion and the local endomorphism action

The Weil pairing identifies exactly what a second high-order torsion point would add to the known rational eigenspace.

Proposition 28.11 (Complementary rational high torsion)**.** *Let $p > 3$ be prime with $p \equiv 2 \pmod 3$, let $a = v_3(p+1)$, $n = 3^a$, and $k = \mathbb{F}_p(\sqrt{-3})$. Suppose an elliptic curve $E/\mathbb{F}_p$ is given with a point P of exact order n in $E(\mathbb{F}_p)$. Let $E^{(-3)}$ be its quadratic twist by -3. Constructing a point of exact order n in $E^{(-3)}(\mathbb{F}_p)$ is deterministically polynomial-time equivalent to constructing an element of exact order n in $k^\times$. The polynomial bound is in $\log p$ and the presented curve data.*

Proof. Write t for the trace of E. The given point implies $p+1-t \equiv 0 \pmod n$. Since $n \mid p+1$, we have $t \equiv 0 \pmod n$, so the twist also has group order $p+1+t$ divisible by n. The rational three-primary subgroup of either curve is cyclic, because two independent rational three-torsion points would give $\mu_3 \subset \mathbb{F}_p$ by the Weil pairing. Thus a point of the required order exists on the twist.

Suppose first that such a twist point has been constructed, and let $Q \in E(k)$ be its image under the standard quadratic-twist isomorphism. Prime-field Frobenius fixes P and sends Q to $-Q$. If a multiple of P equals a multiple of Q, applying Frobenius shows that this common point is killed by two. It is also killed by the odd integer n, and is therefore zero. The two points form a basis of $E[n]$, so $e_n(P,Q)$ has exact order n.

The pairing requires no auxiliary point search. In the shifted Miller formula take $S = 2P+Q$:

$$e_n(P,Q) = \frac{f_P(Q+S)f_Q(-S)}{f_P(S)f_Q(P-S)}.$$

The two evaluations of f_P are outside $\langle P \rangle$, because their Q-coefficients are one and two. The two evaluations of f_Q are outside $\langle Q \rangle$, because their P-coefficients are minus two and minus one. These coefficients are units modulo n. Hence none of the intermediate Miller line evaluations meets a zero or pole in the corresponding cyclic subgroup. The pairing is computable in polynomial time by the standard algorithm (Silverman, 2009, Theorem XI.8.1 and Remark XI.8.2, pp. 393–397).

Conversely suppose a maximal three-primary unit in k is given. It gives roots of every promised cube in k by successive base-three digits. In particular it permits deterministic factorization of a split separable cubic over $\mathbb{F}_p$. Write the depressed cubic as $X^3 + uX + v$. Since cubing permutes $\mathbb{F}_p$, its three distinct roots force $u \neq 0$. Put $D = (v/2)^2 + (u/3)^3$. Its discriminant $-108D$ is a square in $\mathbb{F}_p$. A square root of D can be computed in k using $\sqrt{-3}$ and a square root of the discriminant in $\mathbb{F}_p$, where -3 is a known nonsquare. Both nonzero values $-v/2 \pm \sqrt{D}$ are cubes in k. Indeed, for the three roots r_0, r_1, r_2 and a primitive cube root ζ, the two resolvents $(r_0 + \zeta r_1 + \zeta^2 r_2)/3$ and $(r_0 + \zeta^2 r_1 + \zeta r_2)/3$ have product $-u/3$ and the sum of their cubes is $-v$. Extract $c^3 = -v/2 + \sqrt{D}$ in k. The three roots are $\zeta^j c - u/(3\zeta^j c)$ for $j = 0, 1, 2$.

We first obtain a rational point of order three on the twist. Put $P_0 = [n/3]P$. Frobenius on $E[3]$ has distinct eigenvalues $1, -1$. Its action on the four kernel lines has two fixed lines and a transposed pair. Therefore

$$\gcd(\psi_{3,E}(X), X^p - X)$$

has degree two; removing the known line of P_0 leaves the other linear factor. Its nonzero points have Frobenius -1, so their images on the quadratic twist are rational. Their coordinates use only square extraction with the known nonsquare. This gives a rational twist point Q_1 of order three.

We describe a constant-degree lifting step on the twist, now denoted C. Suppose $Q_i \in C(\mathbb{F}_p)$ has order 3^i, where $i < a$. Cyclicity and divisibility of the rational group order by 3^a imply that Q_i is three-divisible. Choose the known rational subgroup $G = \langle [3^{i-1}]Q_i \rangle$, and compute the degree-three isogeny

$$\phi : C \longrightarrow C' = C/G$$

and its dual. These are constant-degree rational maps computable from the given kernel by the standard explicit three-isogeny formulas. The dual kernel is the Cartier dual of the constant group G, hence has Frobenius -1 and no nonzero prime-field point. The dual is therefore a bijection on rational points, since the two isogenous curves have the same rational group order.

Its inverse image R_i of Q_i is obtained by a constant-degree abscissa equation and a linear gcd with $X^p - X$. There is exactly one prime-field abscissa: an additional rational abscissa would give either another rational preimage or a pair interchanged by negation; the latter would force $Q_i = -Q_i$, impossible for a nonzero point of odd order. Recover the ordinate by square extraction and verify the map.

Because Q_i is three-divisible and $\widehat{\phi}\phi = [3]$, the point R_i lies in $\phi(C(\mathbb{F}_p))$. Its three preimages under ϕ are rational and differ by G. Their abscissas are distinct: equality of two abscissas would make the two points negatives, forcing a preimage of order 3^{i+1} to have a double in G, which is impossible. The abscissa equation is therefore a split separable cubic over $\mathbb{F}_p$. Use the preceding cubic extraction procedure, recover an ordinate, and verify the image. This constructs Q_{i+1} with $[3]Q_{i+1} = Q_i$ and exact order 3^{i+1}. There are $a - 1 = O(\log p)$ such steps, each of constant algebraic degree and polynomial arithmetic complexity. □

For the exact hard branch $v_3(p+1) = 2$, a given rational nine-torsion point lies in one Frobenius eigenspace; obtaining a complementary rational nine-torsion point on the twist is already equivalent to the original maximal-unit gate. If only a point of order nine is given when $a > 2$, the forward pairing yields a ninth root, not automatically a maximal unit. The proposition does not assert that a single given rational high-torsion point alone solves or is equivalent to the gate.

Proposition 28.12 (The full local-order action on three-torsion)**.** *Let $\mathcal{O}$ be the maximal order in the quadratic $\mathbb{Q}_3$-algebra of an ordinary elliptic curve, and choose a $\mathbb{Z}_3$-basis $(1, \omega)$ of $\mathcal{O}$. Thus $\omega^2 - T\omega + N = 0$ for $T, N \in \mathbb{Z}_3$. Suppose its three-adic Tate lattice is explicitly identified, as a lattice with the multiplication action, with $I_k = \mathbb{Z}_3 + 3^k\mathbb{Z}_3\omega$ for $k \geq 1$. On $I_k/3I_k$, the image of the whole local order I_k is $\{aI + bJ : a, b \in \mathbb{F}_3\}$, where $J = \left(\begin{smallmatrix} 0 & 0 \\ 1 & 0 \end{smallmatrix}\right)$. Its unique common eigenline is the ascending line. There exists a unit cyclically permuting the three remaining lines.*

Proof. Put $\vartheta = 3^k\omega$. In the basis $(1, \vartheta)$, multiplication by ϑ is

$$\begin{pmatrix} 0 & -3^{2k}N \\ 1 & 3^kT \end{pmatrix},$$

which reduces to J because T, N are integral and $k \geq 1$. Thus every local-order action has the asserted form. The common eigenline is $\ker J = \operatorname{im} J$, generated by ϑ. Enlarging the lattice by $\vartheta/3$ gives I_{k-1}, identifying the ascending line in this lattice model. For comparison, if one took $\omega = (0, 1/3)$ in $\mathbb{Q}_3 \times \mathbb{Q}_3$ and $k = 1$, the displayed lattice would be the maximal split order. Its multiplication algebra modulo three is $\mathbb{F}_3 \times \mathbb{F}_3$. This explains the integral basis hypothesis. The other lines have representatives $(1, c)$, $c \in \mathbb{F}_3$; the action of $1 + \vartheta$ sends c to $c + 1$. □

This is a statement about the explicit lattice model, rather than a freeness assertion for an unspecified Tate module. A unit in its reduction modulo three may represent an isogeny of degree prime to three; it need not represent an automorphism of the elliptic curve. Dividing Frobenius by further powers of three, whenever the quotient is an endomorphism, remains in the same algebra and gives no second eigenline. On the split coordinate algebra of the three descending kernels, a nontrivial cyclic eigenvector has values $(c, \zeta c, \zeta^2 c)$ and cube c^3. A Fourier projector is a linear operator on that algebra, rather than a primitive algebra idempotent. Obtaining an evaluation of the eigenvector still asks for a cube root. This calculation does not exclude algorithms that use coordinates or other data to choose a kernel.

28.7 The first successful Tate model is a simplest-cubic test

The success condition $v = 1$ has an exact elementary parametrization, so point counting is not needed to test it. The statement below also includes supersingular curves; it does not rely on the ordinary volcano description.

Proposition 28.13. *Let $p > 3$, $p \equiv 2 \pmod 3$, and $k \in \mathbb{F}_p$ with $k \neq 3$. Set*

$$E_k : \quad y^2 + kxy + y = x^3, \qquad h_k(X) = X^3 - (k+3)X^2 + kX + 1.$$

Then h_k is squarefree and is either split or irreducible over $\mathbb{F}_p$. It is irreducible exactly when

$$v_3(\#E_k(\mathbb{F}_p)) = 1.$$

If additionally $p \equiv -1 \pmod 9$, this is equivalent to $v_3(t(E_k)) = 1$. Every nonsingular Tate form $y^2 + axy + by = x^3$ over $\mathbb{F}_p$ is isomorphic over $\mathbb{F}_p$ to one of these models, with $k = a/c$ and $c^3 = b$; this cube root is unique and computable by exponentiation.

Proof. The discriminant of E_k is $k^3 - 27$, which is nonzero precisely when $k \neq 3$, because cubing is bijective in $\mathbb{F}_p$. The point $P = (0, 0)$ has order 3. The quotient isogeny $\phi : E_k \to E'_k = E_k/\langle P \rangle$ has the formulas

$$E'_k : \; y^2 + kxy + y = x^3 - 5kx - (k^3 + 7), \qquad \phi_x(x) = x + k/x + 1/x^2.$$

The corresponding second coordinate is

$$\phi_y(x, y) = y - \frac{(kx + 2)y + (kx + 1)^2}{x^3}.$$

For example the x-formula follows by adding P and $-P$ in the group law and summing their changes in x; substitution gives the displayed quotient equation. Its 3-division polynomial is

$$\psi_{3,E'_k}(X) = (3X + k^2)(X^3 - 9kX - (k^3 + 27)).$$

The point $Q' = (k + 3, 4)$ lies on E'_k and on its 3-division polynomial, so has order 3. The equation $\phi_x(x) = k + 3$ is exactly $h_k(x) = 0$.

A root $x \in \mathbb{F}_p$ of h_k is nonzero. Choose a point above it in the fibre of Q' or $-Q'$. Frobenius fixes its x-coordinate, so fixes the point or negates it. Negation is impossible because its image is a rational point of order 3. Hence the root gives a rational preimage of Q' after a possible sign change. Conversely a rational preimage gives a root. Thus h_k has a rational root exactly when $Q' \in \phi(E_k(\mathbb{F}_p))$.

Both rational groups have the same order and have cyclic 3-primary part: full rational 3-torsion would, by the Weil pairing, imply $\mu_3 \subset \mathbb{F}_p$, contrary to $p \equiv 2 \pmod 3$. The isogeny has rational kernel of order 3, so its image has index 3 in $E'_k(\mathbb{F}_p)$. If the common group order has valuation $m \geq 1$, that image contains the unique subgroup of order 3 exactly when $m \geq 2$. Since Q' generates it, the asserted equivalence follows. Direct calculation gives

$$\operatorname{disc}(h_k) = (k^2 + 3k + 9)^2.$$

The quadratic factor $k^2 + 3k + 9$ never vanishes in $\mathbb{F}_p$, since its discriminant -27 is a nonsquare. The square discriminant restricts the Frobenius permutation to A_3, giving the split-or-irreducible alternative. Finally $p + 1 \equiv 0 \pmod 9$ gives the trace equivalence, and the change $x = c^2X$, $y = c^3Y$ gives the normalization. □

The same cubic is an exact one-parameter Eisenstein residue test. If $\zeta^2+\zeta+1=0$ in $\mathbb{F}_{p^2}$, put $z=k-3\zeta$. With $w=(Y+\zeta)/(Y+\zeta^2)$, direct expansion gives

$$(w-1)^3 h_k\left(\frac{\zeta-w\zeta^2}{w-1}\right)=(k-3\zeta)w^3-(k-3\zeta^2). \tag{73}$$

Neither excluded denominator value is a root, because $k-3\zeta$ and $k-3\zeta^2$ are nonzero. A cubic over $\mathbb{F}_p$ is irreducible if and only if it remains irreducible over its quadratic extension. The displayed transformation and Kummer theory therefore give

$$h_k \text{ irreducible over } \mathbb{F}_p \quad\Longleftrightarrow\quad k-3\zeta \text{ is a noncube in } \mathbb{F}_{p^2}.$$

Indeed the Kummer radicand is $z^p/z=N(z)/z^2$, and $N(z)\in\mathbb{F}_p^*$ is a cube because $p\equiv 2 \pmod 3$. Here $N(z)=k^2+3k+9$.

For integral k, the identity

$$-h_k(-X)=X^3+(k+3)X^2+kX-1$$

identifies the family with the simplest cubic polynomial of parameter $-k-3$; see (Hashimoto and Aoki, 2024, §1, (1.1)–(1.3)). Its irreducibility over $\mathbb{Q}$ follows directly from the rational-root test, and the displayed square discriminant then gives a cyclic cubic field L_k. If m_k is its conductor, the conductor-discriminant formula and the polynomial-index formula give

$$\mathrm{Disc}(L_k)=m_k^2, \qquad m_k \mid k^2+3k+9.$$

Suppose now $p\equiv -1 \pmod 9$ and $h_k \bmod p$ is irreducible. A primitive cubic Dirichlet character of conductor m_k is nontrivial at p. Its local conductor at 3 divides 9, and that local component is trivial at $p\equiv -1 \pmod 9$. At every other prime its nontrivial local component has conductor a prime $\ell\equiv 1 \pmod 3$ and kernel the cubes modulo ℓ. Consequently some prime

$$\ell\le m_k\le k^2+3k+9 \quad\text{satisfies}\quad \ell\equiv 1 \pmod 3, \qquad p^{(\ell-1)/3}\not\equiv 1 \pmod \ell.$$

Thus a uniformly polylogarithmic bound on an *integer* parameter k successful in this normalized family would give the earlier good-auxiliary-prime bound. Normalization can turn small original coefficients a,b into a large residue $k=a/b^{1/3}$, so this height implication is specifically about the normalized integer parameter.

28.8 A supersingular point and an affine-line noncube

A fixed supersingular curve gives a second geometric description of the missing cubic witness. Its point-to-unit map uses the classical Weil pairing and the order-three automorphism; its reverse map uses explicit three-isogeny descent. The proof below identifies both operations and the precise search that remains.

Proposition 28.14 (A supersingular point formulation)**.** *Let $p>3$ be prime with $p\equiv 2 \pmod 3$, put $a=v_3(p+1)$, and present $k=\mathbb{F}_p[\delta]$, where $\delta^2=-3$. For*

$$E/\mathbb{F}_p : y^2=x^3-3,$$

constructing a point of exact order 3^a in $E(\mathbb{F}_p)$ and constructing an element of exact order 3^a in $k^\times$ are deterministically polynomial-time equivalent. Moreover, for every affine point $P=(X,Y)\in E(\mathbb{F}_p)$,

$$P\in 3E(\mathbb{F}_p) \quad\Longleftrightarrow\quad Y+\delta\in (k^\times)^3.$$

Thus a nondivisible point can be sought by scanning the single prime-field parameter Y, with $X=(Y^2+3)^{1/3}$ uniquely computable in $\mathbb{F}_p$.

Proof. Cubing permutes $\mathbb{F}_p$, so each Y gives exactly one affine point and $\#E(\mathbb{F}_p) = p + 1$. Its 3-primary subgroup is cyclic: two independent rational points of order three would give a primitive cube root in $\mathbb{F}_p$ by the Weil pairing. In particular it has order 3^a.

First suppose P has that exact order. With $\zeta = (-1 + \delta)/2$, the automorphism

$$\iota : E_k \longrightarrow E_k, \qquad (x, y) \longmapsto (\zeta x, y)$$

has order three. Write $P_0 = [3^{a-1}]P$. The line generated by P_0 cannot be invariant under ι: an order-three action on a one-dimensional $\mathbb{F}_3$-space would be trivial, whereas the nonzero fixed points of ι have $x = 0$ and $y^2 = -3$, and hence are not rational over $\mathbb{F}_p$. Thus $P_0, \iota(P_0)$ are independent, so $P, \iota(P)$ form a basis of $E[3^a]$. Their Weil pairing has exact order 3^a in $k^\times$.

This pairing is computable without a point-search subroutine. Use the explicit point of order two

$$S = (3^{1/3}, 0) \in E(\mathbb{F}_p).$$

It is outside the odd-order subgroup generated by P and $Q = \iota(P)$. Miller's algorithm and the shifted evaluation formula compute

$$e_{3^a}(P, Q) = \frac{f_P(Q + S) f_Q(-S)}{f_P(S) f_Q(P - S)}, \qquad \operatorname{div}(f_P) = 3^a(P) - 3^a(O),$$

in polynomial time. All evaluation points have a nonzero two-primary component, so the Miller line functions encounter neither zeros nor poles from the three-primary subgroups. See (Silverman, 2009, Chapter III, Section 8; Theorem XI.8.1 and Remark XI.8.2, pp. 393–397).

We next prove the divisibility criterion and the constructive reverse implication. Put $E' : y^2 = x^3 + 81$. The explicit degree-three isogeny

$$\phi : E \longrightarrow E', \qquad (x, y) \longmapsto \left(\frac{x^3 - 12}{x^2}, \frac{y(x^3 + 24)}{x^3} \right)$$

has no nonidentity rational kernel point because -3 is a nonsquare. It is therefore a bijection on prime-field points, as both curves have $p + 1$ points. Its dual is

$$\widehat{\phi} : E' \longrightarrow E, \qquad (z, w) \longmapsto \left(\frac{z^3 + 324}{9z^2}, \frac{w(z^3 - 648)}{27z^3} \right), \qquad \widehat{\phi}\phi = [3].$$

These identities follow by substitution; equivalently they are the standard $j = 0$ three-isogeny and its dual, with the target scaling made explicit. The associated map $Y + \delta$ is a specialization, up to conjugation, of the classical explicit three-isogeny descent map; see (Cohen and Pazuki, 2009, Definition 1.3 and Proposition 1.4, p. 372). We verify the required finite-field case directly below. Consequently $\widehat{\phi}(E'(\mathbb{F}_p)) = 3E(\mathbb{F}_p)$.

For $P = (X, Y)$ with $Y \neq 0$, a preimage under the dual exists over $\mathbb{F}_p$ exactly when

$$f_P(z) = z^3 - 9Xz^2 + 324$$

has a root there. The ordinate is then $27Yz^3/(z^3 - 648)$. Its denominator cannot vanish for such a root when $Y \neq 0$. The cubic

$$h_P(w) = w^3 - 3Xw - 2Y$$

has discriminant 324 and has roots corresponding to those of f_P through

$$z = 3(w^2 - X), \qquad w = \frac{6Y}{z - 6X}.$$

Substitution proves the identities, and the denominators are nonzero when $Y \neq 0$.

Set $\gamma = Y + \delta$, whose norm is X^3. If $c^3 = \gamma$ in k, then the norm of c is X, because cubing is bijective in $\mathbb{F}_p$. Hence $w = c + X/c$ belongs to $\mathbb{F}_p$ and is a root of h_P. Conversely, if γ is a noncube, the cyclic cubic extension $k(c)/k$ acts transitively on the three distinct values $c + X/c$, so h_P is irreducible over k and over $\mathbb{F}_p$. For distinct cube roots equality of these values would force their product to be X, and cubing would give $\gamma = \overline{\gamma}$, impossible since $\delta \neq 0$. This proves the criterion for $Y \neq 0$. For $Y = 0$, the point has order two and is three-divisible; δ is a cube because $\delta^2 = -3 \in \mathbb{F}_p^\times$ is a cube and the cube-class group has order three.

Finally, a known maximal three-primary unit gives promised cube extraction in k by successive base-three digits. Start with the explicit point

$$P_1 = (12^{1/3}, 3) \in E(\mathbb{F}_p)[3] \setminus \{O\}.$$

For a point $P_i = (X, Y)$ of order 3^i, $i < a$, cyclicity guarantees a rational preimage under $[3]$. The proved criterion therefore says that $Y + \delta$ is a promised cube. Extract $c^3 = Y + \delta$ using the maximal unit, and form $w = c + X/c$ and $z = 3(w^2 - X)$. These lie in $\mathbb{F}_p$ and the displayed formulas give a point of E' mapping to P_i under the dual. The resulting E'-point has a unique prime-field preimage under ϕ; its abscissa is the unique root over $\mathbb{F}_p$ of

$$Z^3 - X'Z^2 - 12,$$

found as a linear gcd with $Z^p - Z$. Indeed the lifted point of E' has odd order, since its image has odd order and the dual kernel has order three; its ordinate is therefore nonzero. For any prime-field root z of the fiber cubic, $z \neq 0$ and $z^3 + 24 \neq 0$: the latter equality would force the target ordinate to vanish. The ordinate $Y'z^3/(z^3 + 24)$ is then rational and gives a preimage. The bijectivity of ϕ on rational points therefore makes this root unique. Recover its ordinate by that formula. This constructs P_{i+1} with $[3]P_{i+1} = P_i$. After $a - 1 = O(\log p)$ iterations its order is 3^a. All quadratic extraction uses the known nonsquare -3. □

For the hard branch $p \equiv -1 \pmod 9$, if a small integer Y yields a nondivisible point, the Eisenstein integer $Y + \sqrt{-3}$ has norm $Y^2 + 3$. The established cubic reciprocity argument then gives a good auxiliary prime $\ell \leq Y^2 + 3$. Group operations applied only to points already in $3E(\mathbb{F}_p)$ cannot leave that subgroup. Thus this supersingular formulation makes the remaining noncube search explicit; it is an equivalence, not a strictly weaker sufficient condition.

Proposition 28.15 (Scalar Frobenius in the supersingular walk)**.** *Let $p > 3$, $p \equiv 2 \pmod 3$, and let $E_0/\mathbb{F}_p$ be supersingular. Every curve reached by isogenies defined over $k = \mathbb{F}_{p^2}$ has k-Frobenius $[-p]$. Consequently every such curve has all its three-torsion rational over k. Whenever it is put in Tate form using a rational point of order three, its remaining cubic splits and its Tate radicand is a cube. Such a walk never reaches an ordinary floor producing a noncube.*

Proof. For a supersingular prime-field curve, Frobenius and its dual are purely inseparable. Their sum $[t_0]$ has zero differential, so $p \mid t_0$. Hasse's bound and $p > 3$ give $t_0 = 0$, hence prime-field Frobenius satisfies $\pi^2 = [-p]$. If $\phi : E_0 \to E_1$ is defined over k, then $\pi_{k,E_1}\phi = \phi\pi_{k,E_0} = [-p]\phi$. Surjectivity of ϕ gives $\pi_{k,E_1} = [-p]$. Induction handles the whole walk. Since $-p \equiv 1 \pmod 3$, Frobenius fixes all three-torsion points. The explicit Tate root formula then gives the last assertions. □

The program *Supersingular points and cubic classes* checks the isogeny maps, divisibility criterion, and maximal-order lift on the stated finite examples.

28.9 A genus-two family does not separate three trace lifts

Frobenius polynomials can reveal factors when their values differ between components, so a possible alternative to direct root extraction is to attach a curve to each unknown trace. Kayal's idea and Poonen's generalization make this criterion precise: Poonen's Hypothesis Z asks for a fixed abelian family whose specializations have distinct zeta functions, and his Theorem 5.1 gives deterministic polynomial factoring under that hypothesis (Poonen, 2019, author pp. 2–5). Booher, Howe, Sutherland and Voloch propose an explicit fifteen-curve version; their separation assertion is Conjecture 10.1 and their factoring conclusion, Theorem 10.4, is conditional on it (Booher et al., 2026, pp. 38–41). Their family and conjecture are different from the family considered here.

The following example tests one direct genus-two construction on a three-element cyclotomic trace fibre. Elliptic quotients of dihedral genus-two curves are classical; see Cardona and Quer (2002, Section 5, Proposition 5.1, author pp. 19–20) for the characteristic-zero description. We give the finite-field quotients and the point-count identity explicitly, so the calculation does not require a transfer of that characteristic-zero statement.

Proposition 28.16. *Over $\mathbb{F}_{107}$, put $C_3(T) = T^3 - 3T$. Then*

$$C_3(47) = -1, \qquad \{t : C_3(t) = 47\} = \{17, 98, 99\}.$$

Each of these three values is the trace of an order-27 element whose cube is the same chosen order-9 element of $\mathbb{F}_{107^2}$. The smooth projective curves

$$\mathcal{C}_t : \quad y^2 = x^6 - tx^3 + 1 \qquad (t = 17, 98, 99)$$

all have Frobenius characteristic polynomial

$$X^4 + 133X^2 + 11449.$$

Consequently this polynomial, even together with the lower trace 47 and characteristic 107, does not distinguish the three trace branches.

Proof. All arithmetic in the explicit tables is modulo 107. Multiplication gives

$$(T - 17)(T - 98)(T - 99) = T^3 - 3T - 47, \qquad 47^3 - 3 \cdot 47 = -1.$$

Since $107 \equiv 3 \pmod 4$, write $\mathbb{F}_{107^2} = \mathbb{F}_{107}[i]/(i^2 + 1)$. For $\alpha = 77 + 8i$ and $\zeta = 53 + 9i$, direct multiplication gives

$$N(\alpha) = 1, \quad \alpha + \alpha^{-1} = 47, \quad \alpha^3 = \zeta, \quad \zeta^3 = 1, \quad \zeta \neq 1.$$

Thus α has exact order nine. Its three cube roots are

t	17	98	99
v_t	$62 + 3i$	$49 + 75i$	$103 + 29i$

and each satisfies $N(v_t) = 1$, $v_t^3 = \alpha$, and $v_t + v_t^{-1} = t$. Multiplication by ζ cycles the three entries. Their order is therefore 27. The identity $(v + v^{-1})^3 - 3(v + v^{-1}) = v^3 + v^{-3}$ explains the trace fibre. Fixing α also fixes the orientation: of v_t, v_t^{-1}, only the first cubes to α.

The sextic is squarefree when $t^2 \neq 4$: a common zero with its derivative would satisfy $x \neq 0$, $2x^3 = t$, and then $1 - t^2/4 = 0$. The three parameters avoid these values, so the curves have genus two. Their two elliptic quotients are the smooth projective models of

$$E_t^+ : v^2 = (u + 2)(u^3 - 3u - t), \qquad E_t^- : w^2 = (u - 2)(u^3 - 3u - t).$$

Both quartics are squarefree for $t^2 \neq 4$, and their monic leading terms give rational points at infinity. The quotient maps on $x \neq 0$ are

$$u = x + x^{-1}, \qquad v = \frac{y(x+1)}{x^2}, \qquad w = \frac{y(x-1)}{x^2}.$$

These degree-two maps extend to the smooth projective models.

Here is a direct proof of the zeta-function factorization, including the fibres omitted by these affine formulas. Over any finite extension $k/\mathbb{F}_{107}$, let χ be the quadratic character with $\chi(0) = 0$. For every $u \in k$,

$$\sum_{\substack{x \in k^* \\ x + x^{-1} = u}} \chi(x) = \chi(u+2) + \chi(u-2).$$

If $u^2 - 4$ is nonsquare, both sides are zero. If it is a nonzero square, the roots are x, x^{-1}, and $(x \pm 1)^2/x = u \pm 2$ proves the identity. At $u = 2$ and $u = -2$, the values are respectively 1 and $\chi(-1)$. For $x \neq 0$, the sextic is $x^3(u^3 - 3u - t)$; its value at zero is one. Summing the identity and including the two points at infinity on each monic even-degree model gives, with $q = |k|$,

$$\#\mathcal{C}_t(k) = \#E_t^+(k) + \#E_t^-(k) - q - 1.$$

This holds over every extension. The exponential definition of zeta functions therefore makes the genus-two Frobenius polynomial the product of the two elliptic Frobenius polynomials.

It remains to compute six small elliptic counts. For each $f_{t,e}(u) = (u+e)(u^3 - 3u - t)$, $e = 2, -2$, evaluate $u = 0, \ldots, 106$. A nonzero value is a square exactly when $f_{t,e}(u)^{53} = 1$. The complete calculation is specified by the following integer-arithmetic loop; its three counts are ordered as nonsquare, zero, square.

```
for t in (17, 98, 99):
    for e in (2, -2):
        counts = [0, 0, 0]
        for u in range(107):
            z = ((u + e)*(u**3 - 3*u - t)) % 107
            if z == 0:
                j = 1
            else:
                j = 2 if pow(z, 53, 107) == 1 else 0
            counts[j] += 1
        print(t, e, counts)
```

It gives

	$e = 2$			$e = -2$		
t	nonsquare	zero	square	nonsquare	zero	square
17	49	1	57	58	1	48
98	49	1	57	58	1	48
99	49	1	57	58	1	48

Thus the character sums are 8 and -10. A monic quartic model has $107 + 2 + \sum \chi(f(u))$ points, so $\#E_t^+ = 117$ and $\#E_t^- = 99$. Their traces are -9 and 9, giving

$$\begin{aligned} P_{\mathcal{C}_t}(X) &= (X^2 + 9X + 107)(X^2 - 9X + 107) \\ &= (X^2 + 107)^2 - 81X^2 = X^4 + 133X^2 + 11449. \end{aligned}$$

The map from the three trace values to this polynomial is constant, which proves the stated failure of separation. □

This example does not refute eventual separation in another family, an adaptive collection of families, or Poonen’s Hypothesis Z. It also does not show that intermediate operations in a point-counting algorithm cannot reveal a factor despite identical final polynomials. In this finite example any of the three displayed traces is a valid lift; the conclusion concerns only the distinguishing information carried by this one family’s Frobenius polynomial.

The collision above has a structural feature that can be explained uniformly. The two elliptic quotients displayed below at a single parameter are related by a three-isogeny and a quadratic twist. This explains the vanishing odd coefficients of the genus-two Frobenius polynomial, but does not force that polynomial to be constant on a trace-lifting fibre. The dihedral elliptic-quotient phenomenon is classical (Cardona and Quer, 2002, Proposition 5.1, author pp. 19–20); the following calculation fixes the models and the twist sign used here.

Proposition 28.17. *Let k be a field of characteristic different from 2 and 3, and let $t \in k \setminus \{2, -2\}$. The smooth projective curve*

$$C_t : \quad y^2 = x^6 - tx^3 + 1$$

has genus two. Two complementary degree-two elliptic quotients are

$$E_{\pm}(t) : \quad v^2 = (s \pm 2)(s^3 - 3s - t), \qquad \phi_{\pm}(x, y) = \left(x + x^{-1},\ \frac{y(x \pm 1)}{x^2} \right).$$

Use $(-2, 0)$ as the identity on $E_+(t)$, and $(2, 0)$ on $E_-(t)$. Their product induces a k-isogeny

$$\mathrm{Jac}(C_t) \longrightarrow E_+(t) \times E_-(t)$$

of degree four. Moreover $E_+(t)$ has a rational subgroup of order three whose quotient is k-isomorphic to $E_-(t)^{(-3)}$. Here the quadratic D-twist of $V^2 = U^3 + AU^2 + BU + C$ means

$$V^2 = U^3 + DAU^2 + D^2BU + D^3C.$$

If $k = \mathbb{F}_p$, $p > 3$, and $a_t = p + 1 - \#E_+(t)(\mathbb{F}_p)$, then

$$a(E_-(t)) = \left(\frac{-3}{p} \right) a_t.$$

In particular, for $p \equiv -1 \pmod 3$, the characteristic polynomial of Frobenius on $\mathrm{Jac}(C_t)$ is

$$P_t(T) = T^4 + (2p - a_t^2)T^2 + p^2.$$

Proof. The derivative of $x^6 - tx^3 + 1$ is $3x^2(2x^3 - t)$. A common zero would be nonzero and satisfy $x^3 = t/2$, forcing $1 - t^2/4 = 0$. Thus the sextic is squarefree and its smooth hyperelliptic model has genus two.

Put $s = x + x^{-1}$ and $v_{\pm} = y(x \pm 1)/x^2$. The identities

$$x^3 + x^{-3} = s^3 - 3s, \qquad (x \pm 1)^2/x = s \pm 2$$

prove the displayed quotient equations. The maps are invariant under

$$\sigma(x, y) = (x^{-1}, y/x^3), \qquad \iota\sigma(x, y) = (x^{-1}, -y/x^3),$$

respectively, where ι is the hyperelliptic involution. Since $x^2 - sx + 1 = 0$, these are the corresponding degree-two function-field inclusions. The rational maps extend to morphisms of the smooth projective curves. The cubic factor in each quartic has discriminant $27(4 - t^2)$; its value at -2, respectively 2, is $-(t + 2)$, respectively $2 - t$. Thus both quartics are smooth genus-one models with the indicated rational identities.

Their regular differentials pull back as

$$\phi_+^*(ds/v) = (x-1)\,dx/y, \qquad \phi_-^*(ds/v) = (x+1)\,dx/y.$$

These are independent, so the induced map F from the Jacobian to the product is an isogeny. If G is the sum of the two pullbacks, the quotient identities give

$$GF = (1+\sigma^*) + (1+(\iota\sigma)^*) = [2].$$

Here $\iota^* = -1$ on the Jacobian. The maps F and G are dual for the principal polarizations, so have equal degree. Since $\deg[2] = 16$ in dimension two, $\deg F = 4$.

The change $s \mapsto -s$, with v unchanged, identifies $E_-(t)$ with $E_+(-t)$, preserving the indicated origins. Write $a = t+2 \neq 0$. On $E_+(t)$ put

$$X = -\frac{a}{s+2}, \qquad V = \frac{av}{(s+2)^2}, \qquad Y = V - 3X - a.$$

These substitutions give, successively,

$$V^2 = X^3 + 9X^2 + 6aX + a^2, \qquad \mathcal{E}_t:\ Y^2 + 6XY + 2aY = X^3.$$

The inverse on $X \neq 0$ is $s = -a/X - 2$, $v = aV/X^2$, $V = Y + 3X + a$. The maps therefore extend to an isomorphism of the smooth projective curves and send $(-2, 0)$ to the point at infinity. On $\mathcal{E}_t$, the line $Y = 0$ is the tangent at $(0,0)$ and meets the cubic there with multiplicity three. Consequently $P = (0,0)$ has exact order three; its subgroup is $\{O, (0,0), (0,-2a)\}$.

For the equation $Y^2 + bXY + cY = X^3$, the quotient by this subgroup has equation

$$Y'^2 + bX'Y' + cY' = X'^3 - 5bcX' - c(b^3 + 7c).$$

For clarity, this is the direct specialization of Vélu (1971, translated pp. 3–4, equations (8)–(11)): the representative $(0,0)$ has $g^x = 0$, $g^y = -c$, $t_Q = bc$, $u_Q = c^2$, and the totals in those formulas are $t_{\mathrm{V}} = bc$, $w_{\mathrm{V}} = c^2$. The rational kernel makes both the quotient and its isogeny defined over k. One may also verify the quotient directly using

$$X' = X + \frac{bc}{X} + \frac{c^2}{X^2},$$

$$Y' = Y - \frac{c^2(2Y + bX + c)}{X^3} - \frac{bc(bX + Y)}{X^2} - \frac{bc^2}{X^2}.$$

Set $b = 6$, $c = 2a$, and complete the square by $V' = Y' + 3X' + a$. The quotient becomes

$$V'^2 = X'^3 + 9X'^2 - 54(t+2)X' - 27(t+2)(t+18).$$

The -3-twist of the completed-square model of $E_+(-t)$ is

$$W^2 = U^3 - 27U^2 + 54(2-t)U - 27(t-2)^2.$$

Substitution $U = X' + 12$, $W = V'$, gives exactly the preceding quotient. This proves the asserted twist identity, including its sign.

Over $\mathbb{F}_p$, rational isogenies preserve elliptic Frobenius traces, and a quadratic D-twist multiplies the trace by $\left(\frac{D}{p}\right)$. Hence $a(E_-(t)) = \left(\frac{-3}{p}\right) a_t$. For $p \equiv -1 \pmod 3$, the polynomial $Z^2 + Z + 1$ has no root in $\mathbb{F}_p$, so its discriminant -3 is a nonsquare. The Jacobian isogeny therefore gives

$$P_t(T) = (T^2 - a_t T + p)(T^2 + a_t T + p) = T^4 + (2p - a_t^2)T^2 + p^2.$$

□

At special parameters there can be other degree-two quotients. Over $\mathbb{Q}$ at $t = 0$, the map $(x, y) \mapsto (x^2, y)$ gives $v^2 = u^3 + 1$, with $j = 0$. The displayed pair has $j = 54000$, as follows from its Weierstrass models. Thus the complementary pair exhibiting the splitting need not exhaust the elliptic quotients.

Proposition 28.18. *The Frobenius polynomial in Proposition 28.17 is not universally constant on the three-point fibres used in the trace-lifting step. Specifically, in $\mathbb{F}_{53}$, put $u = 4$ and $C_3(Z) = Z^3 - 3Z$. Then $C_3(u) = -1$, and the three roots of $C_3(t) = u$ are $5, 7, 41$. Their Frobenius polynomials are, respectively,*

$$T^4 + 97T^2 + 2809, \qquad T^4 + 25T^2 + 2809, \qquad T^4 + 25T^2 + 2809.$$

There is also no generic isogeny between the following two branches. Let z be transcendental over $\overline{\mathbb{Q}}$, let ζ be a primitive cube root of unity, and put

$$t = z + z^{-1}, \qquad t' = \zeta z + \zeta^{-1} z^{-1}.$$

Although $C_3(t) = C_3(t')$, the Jacobians of C_t and $C_{t'}$ are not isogenous over $\overline{\overline{\mathbb{Q}}(z)}$.

Proof. The finite-field identities are obtained by direct substitution. The completed-square elliptic model is

$$V^2 = f_t(X), \qquad f_t(X) = X^3 + 9X^2 + 6(t+2)X + (t+2)^2.$$

Let n_-, n_0, n_+ count the $X \in \mathbb{F}_{53}$ for which the quadratic character of $f_t(X)$ is $-1, 0, 1$. The exact counts are

t	n_-	n_0	n_+	$\#E_+(t)(\mathbb{F}_{53})$	a_t
5	28	0	25	51	3
7	31	0	22	45	9
41	22	0	31	63	−9

For an independently reproducible exhaustive check, the following integer-arithmetic calculation gives the first three count columns:

```
for t in (5, 7, 41):
    counts = [0, 0, 0]       # nonsquare, zero, square
    for x in range(53):
        f = (x**3 + 9*x*x + 6*(t+2)*x + (t+2)**2) % 53
        symbol = pow(f, 26, 53)
        index = 1 if symbol == 0 else 2 if symbol == 1 else 0
        counts[index] += 1
    print(t, counts)
```

Since $\#E_+ = 1 + n_0 + 2n_+$, the traces in the table follow. The formula in Proposition 28.17 then gives the three stated polynomials. In particular this fibre has unequal polynomials, while Proposition 28.16 gives an equal fibre at $p = 107$.

For the function-field assertion, the displayed parameters satisfy

$$C_3(t) = z^3 + z^{-3} = C_3(t').$$

The normalized elliptic curve has

$$c_4 = 144(5 - 2t), \quad \Delta = 432(2 - t)(t + 2)^3, \quad j(t) = \frac{6912(5 - 2t)^3}{(2 - t)(t + 2)^3}.$$

At $z = 1$, we have $t - 2 = (z - 1)^2/z$, so $j(t)$ has a pole. By contrast $t'(1) = \zeta + \zeta^{-1} = -1$, so $j(t')$ is integral at this valuation and its displayed model has unit discriminant.

Isogenous elliptic curves in characteristic zero have the same support of poles of their j-invariants, including after a finite extension of the function field. Indeed, after a finite extension one can factor an isogeny into cyclic prime-degree isogenies. At each step the two j-invariants satisfy the classical modular polynomial, monic with integral coefficients in each

variable. Integrality of one invariant in a valuation ring therefore forces integrality of the other, and conversely. Pole or nonpole status persists under extension of a valuation. This precise elementary pole-support argument is also recorded in Griffon and Pazuki (2022, Remark 5.5, author V2 p. 13). Thus $E_+(t)$ and $E_+(t')$ cannot become isogenous over any finite algebraic extension.

Over the algebraic closure of the function field, the -3 twist is trivial. Proposition 28.17 gives

$$\mathrm{Jac}(C_t) \sim E_+(t)^2, \qquad \mathrm{Jac}(C_{t'}) \sim E_+(t')^2.$$

An isogeny between the Jacobians would yield a nonzero homomorphism between these two products. Some elliptic matrix entry would be nonzero, hence would be an isogeny $E_+(t) \to E_+(t')$, a contradiction. □

These assertions concern the displayed family and, in the second case, the specified generic pair. They do not exclude isogenies after particular finite-field specializations, nor do they supply a uniformly separating family or a general obstruction to adaptive Frobenius constructions. The explicit model transformations have constant degree, but they do not produce a maximal three-primary point or an input-free construction of the desired finite field.

29 Quadratic class fields and witness selection

This section gives class-field formulations of the missing witness and tracks the discriminant budget they require.

29.1 A ramified ring-class alternative

A whole CM class polynomial avoids choosing one CM j-invariant. The resulting sufficient condition still contains a cyclic r-extension splitting condition, which is distinct from quadratic splitting in the CM field.

Proposition 29.1. *Let r be an odd prime and $D < -4$ a negative fundamental discriminant with $r \mid D$ and $r \nmid h(D)$. Put $K = \mathbb{Q}(\sqrt{D})$, $O = O_K$, and $O_r = \mathbb{Z} + rO$. Let $p \nmid rD$ be an odd prime splitting as $pO = \mathfrak{p}\overline{\mathfrak{p}}$. Let e be the order of $[\mathfrak{p}]$ in $\mathrm{Cl}(O)$, and choose $\pi = (t + v\sqrt{D})/2$ with $\mathfrak{p}^e = (\pi)$. The reduction of the class polynomial H_{r^2D} modulo p is squarefree. If $r \nmid v$, all its irreducible factors have degree re, and their number is $h(D)/e$. If $r \mid v$, all its factors have degree e, and their number is $rh(D)/e$. Thus it satisfies the degree-r factor-count property exactly when $r \nmid v$, equivalently when*

$$\pi \bmod rO \notin \mathbb{F}_r^* \quad \Longleftrightarrow \quad v_r(t^2 - 4p^e) = 1.$$

In the positive case, a presentation of $\mathbb{F}_{p^r}$ can be constructed deterministically in time polynomial in r, $|D|$, and $\log p$.

Proof. Since $D < -4$, both orders have unit group $\{1, -1\}$. The conductor exact sequence is

$$1 \longrightarrow (O/rO)^*/\mathbb{F}_r^* \longrightarrow \mathrm{Cl}(O_r) \longrightarrow \mathrm{Cl}(O) \longrightarrow 1;$$

see (Bröker et al., 2012, §3, (3)). Here its conductor-coprimality hypothesis holds because O is maximal. Since r ramifies and is odd, $O/rO \simeq \mathbb{F}_r[\epsilon]/(\epsilon^2)$. Modulo scalar units, its units are $1 + b\epsilon$, with multiplication given by addition of b. The kernel therefore has order r, and $h(r^2D) = rh(D)$.

Let g be the class of $\mathfrak{p} \cap O_r$. Its image has order e, which is prime to r. Consequently $\mathrm{ord}(g)$ is either e or re. Its e-th power in the kernel is the class of $\pi \bmod rO$, up to inversion according

to the Artin-map convention. It is trivial exactly when that residue is scalar. The image of $\sqrt{D}$ is a nonzero nilpotent generator of O/rO, so this is equivalent to $r \mid v$. The identity

$$t^2 - 4p^e = Dv^2$$

and $v_r(D) = 1$ give the valuation criterion.

By CM reciprocity, the roots of H_{r^2D} form a regular $\mathrm{Cl}(O_r)$-torsor. At p split in K and prime to the conductor and discriminant, ordinary reduction preserves this torsor and its distinct roots; Frobenius acts by translation by g or g^{-1}. The factorization statement is the Frobenius computation in (Cremona and Sutherland, 2024, §2.3); its Proposition 4 establishes squarefreeness for every split prime under consideration. Every orbit therefore has length $\mathrm{ord}(g)$. Dividing $rh(D)$ by the two possible orbit lengths gives the stated factor counts. In the positive case $h(D)/e$ is prime to r, so the factor-count construction applies.

For the complexity assertion, the elementary complex-evaluation construction computes H_Δ deterministically in time polynomial in $|\Delta|$; a coarse bound suffices here. Enumerate its primitive reduced quadratic forms, evaluate the corresponding j-values with certified precision, and multiply the linear factors. The form bounds give polynomially many roots, $|q| \le e^{-\pi\sqrt{3}}$ at every reduced CM point, and $|j| \le e^{O(\sqrt{|\Delta|})}$. The exponentially convergent j-series therefore gives polynomial-size coefficient intervals that can be narrowed below $1/2$ and rounded using CM integrality. Taking $\Delta = r^2 D$ proves the required polynomial bound for constructing and reducing H_{r^2D}. The remaining algebra has polynomial cost in its degree and $\log p$. □

The units hypothesis is material: for $D = -3, r = 3$ the additional units remove the proposed order-r kernel. The ramified-prime calculation above uses the exact sequence directly, not the unramified-prime hypothesis of (Bröker et al., 2012, Lemma 3.1).

Let L/K be the Hilbert class field and L_r/K the ring class field of conductor r. The positive condition in Proposition 29.1 says that a prime of L above $\mathfrak{p}$ has nontrivial Frobenius in $\mathrm{Gal}(L_r/L) \simeq C_r$. Rational Dirichlet GRH can control the quadratic splitting test for K, but no polynomial bound for a suitable D with this additional condition is established here. The fact that L_r/K is abelian does not identify its characters with rational Dirichlet characters. A uniform existence bound $|D| \le \mathrm{poly}(r, \log p)$ for a positive case would suffice: enumerate the discriminants, count their reduced forms, and apply deterministic distinct-degree decomposition to $H_{r^2D} \bmod p$ to test the factor-count property. This search requires neither a generator π nor a chosen CM root.

Example 29.2. Take $D = -15$ and $r = 3$. Here $h(-15) = 2$ and $h(-135) = 6$. The primes 17 and 107 are both 17 (mod 45), but

$$4 \cdot 17^2 = 14^2 + 15 \cdot 8^2, \qquad 4 \cdot 107^2 = 106^2 + 15 \cdot 48^2.$$

Both prime ideals have class order $e = 2$. The first coefficient v is prime to 3, and the second is not. Accordingly,

$$H_{-135} \bmod 17 = X^6 + X^5 + 10X^4 + 16X^3 + 15X^2 + 9X + 3$$

is irreducible, whereas

$$H_{-135} \bmod 107 = X^6 + 46X^5 + 88X^4 + 9X^3 + 70X^2 + 80X + 22$$

has exactly three distinct irreducible quadratic factors. For a root α of the degree-six polynomial over $\mathbb{F}_{17}$, the element $\beta = \alpha + \alpha^{17^3}$ has minimal polynomial

$$X^3 + X^2 + 4X + 8 \in \mathbb{F}_{17}[X].$$

The program *CM class-polynomial and subfield certificate* derives H_{-15} and Φ_3 from exact modular q-series and certified CM intervals, computes the integer resultant, and verifies both factor patterns and this subfield polynomial by Frobenius arithmetic and gcds. The independently derived input polynomials agree with the discriminant--15 singular moduli displayed in (Duke, 2006, p. 1 of the author manuscript) and the numerical Φ_3 coefficients in the official PARI source PARI Group (2026).

29.2 A class-group period formulation

The preceding construction can be expressed more generally by taking periods of CM values over class-group cosets. Decompositions of this kind are classical; see (Sutherland, 2012, §2.5), which develops the constructions of Enge–Morain and Hanrot–Morain. The proposition below states the particular finite-field selection condition that is needed here. It does not assert a bound for a discriminant satisfying that condition.

Proposition 29.3. *Let $\Delta < -4$ be an imaginary quadratic order discriminant, $A = \mathrm{Cl}(O_\Delta)$, and $h = |A|$. Let $p \nmid \Delta$ be a prime split in $K = \mathbb{Q}(\sqrt{\Delta})$, and let $g \in A$ be the class of a prime above p. Let r be an odd prime. Write $rA = \{a^r : a \in A\}$ in multiplicative notation. If $g \notin rA$, then a presentation of $\mathbb{F}_{p^r}$ can be constructed deterministically in time polynomial in $|\Delta|$ and $\log p$.*

More precisely, index the CM j-values by A as $(j_b)_{b\in A}$. For an index-r subgroup $H \le A$ and $1 \le k \le h-1$, define

$$\theta_{aH,k} = \sum_{b\in aH} j_b^k, \qquad P_{H,k}(X) = \prod_{aH\in A/H} (X - \theta_{aH,k}).$$

Every $P_{H,k}$ belongs to $\mathbb{Z}[X]$ and has degree r. If $g \notin rA$, at least one such polynomial is irreducible modulo p. All these polynomials can be enumerated and computed in time polynomial in $|\Delta|$; the class g need not be computed. Conversely, if $g \in rA$, every $P_{H,k} \bmod p$ is a product of linear factors, with repetitions allowed.

Proof. Let L be the ring class field. The CM values form a regular $A = \mathrm{Gal}(L/K)$-torsor. Choose the indexing so that complex conjugation sends j_b to $j_{b^{-1}}$; equivalently, take the principal CM value as the initial value. Translations by A and inversion both permute the cosets of H. They therefore permute the $\theta_{aH,k}$. The coefficients of $P_{H,k}$ are algebraic integers fixed by $\mathrm{Gal}(L/\mathbb{Q}) = A \rtimes \langle c\rangle$, where c acts by inversion. Hence they are rational integers. This also explains why passing from coefficients over K to coefficients over $\mathbb{Q}$ is valid in this CM setting.

At a prime above p, ordinary CM reduction gives h distinct residues $\bar{j}_b$; this is the injectivity statement in (Cremona and Sutherland, 2024, §§2.3–2.4, Proposition 4). It is stronger than the assertion that p is unramified in L. Frobenius acts on these residues by $b \mapsto gb$, or by its inverse depending on the Artin-map convention.

The group A/rA is a vector space over $\mathbb{F}_r$. If $g \notin rA$, choose a linear functional nonzero on its image, and let H be the inverse image of its kernel. Then $[A : H] = r$ and gH generates A/H. Suppose that for every $0 \le k \le h-1$ the two sums

$$\sum_{b\in H} \bar{j}_b^k, \qquad \sum_{b\in gH} \bar{j}_b^k$$

were equal. The difference of the two coset-indicator vectors would then be annihilated by the Vandermonde matrix $V = (\bar{j}_b^k)_{0\le k<h,\, b\in A}$. Its determinant is a nonzero product of differences of the distinct $\bar{j}_b$, so V is invertible. The indicator difference is nonzero, because H and gH are disjoint. This is a contradiction in every characteristic, including when p divides $|H|$. The sums for $k = 0$ are equal, so some $1 \le k \le h-1$ separates the two cosets.

For this k, the Frobenius orbit of $\bar{\theta}_{H,k}$ is nonconstant and its length divides the prime r. Its length is therefore r. These are exactly the r roots of $P_{H,k} \bmod p$, proving irreducibility. If instead $g \in rA$, its image is trivial in every index-r quotient of A. Frobenius then fixes every coset sum, so every root of every $P_{H,k} \bmod p$ lies in $\mathbb{F}_p$.

It remains to justify the deterministic complexity assertion. Put $M = |\Delta|$. The reduced positive primitive forms have $1 \le a \le \sqrt{M/3}$, $|b| \le a$, and $c = (b^2 + M)/(4a) \ge a$. Enumerating these integer pairs and applying the usual boundary conventions gives all $h = O(M)$ classes. Composition and reduction give the entire class-group table using polynomially many operations on polynomial-size integers. From this table one constructs A/rA by exhaustive group operations and elementary linear algebra. If its dimension is d, the number of index-r subgroups is $(r^d - 1)/(r - 1) \le h$. All of them can thus be enumerated in polynomial time. This uses an explicitly enumerated group, rather than a discrete logarithm oracle.

The CM values and the period polynomials can be computed by certified complex evaluation and integer rounding, with a coarse polynomial bound sufficient here. At a reduced CM point, $q = \exp(2\pi i\tau)$ satisfies

$$e^{-\pi\sqrt{M}} \le |q| \le e^{-\pi\sqrt{3}} < 1/200.$$

The absolutely convergent identity

$$j(q) = \frac{\left(1 + 240\sum_{n\ge1}\sigma_3(n)q^n\right)^3}{q\prod_{n\ge1}(1-q^n)^{24}}$$

gives $|j(q)| \le 2^{O(\sqrt{M})}$. For example, $\sigma_3(n) \le n^4$ bounds the series tails geometrically after a polynomial factor; the omitted logarithm of the product is also geometrically bounded. Thus N certified absolute bits of each j require only $O(N+\sqrt{M}+\log(N+2))$ terms, with polynomial-size rational interval arithmetic. For $k < h$, each period has size at most $h\,2^{O(h\sqrt{M})}$, and a coefficient of its degree-r period polynomial has logarithmic size $O(rh\sqrt{M} + r\log(h+1))$. The elementary estimates $|z^k - w^k| \le k\max(|z|,|w|)^{k-1}|z-w|$ and the product telescoping identity show that a polynomial number of additional guard bits makes every final coefficient interval have radius less than $1/4$. Its unique integer is then recovered exactly. This establishes a coarse rigorous polynomial bound; it does not invoke the heuristic near-optimal floating-point error estimate discussed in (Enge, 2009, §1, Corollary 1.3).

Finally, reduce the resulting integer polynomials modulo p and test them for irreducibility. For a monic polynomial P of prime degree r, the conditions

$$X^{p^r} \equiv X \pmod{P}, \qquad \gcd(P, X^p - X) = 1$$

are equivalent to irreducibility. Modular powering and gcd computation have polynomial cost in r and $\log p$, and $r \le h$ whenever an index-r subgroup exists. At least one test succeeds in the positive case. No chosen CM root and no computation of g are required. □

This gives another precise sufficient input for the original field construction problem: a discriminant

$$|\Delta| \le \mathrm{poly}(r, \log p), \qquad \left(\frac{\Delta}{p}\right) = 1, \qquad [\mathfrak{p}] \notin r\,\mathrm{Cl}(O_\Delta).$$

A uniform existence bound of this form would suffice, since one could enumerate all discriminants in the range and test the period polynomials. No such bound is established from rational Dirichlet GRH here. In the ramified-conductor situation of Proposition 29.1, the r-Sylow subgroup of $\mathrm{Cl}(O_r)$ has order r. Consequently its positive case is exactly the condition $g \notin r\,\mathrm{Cl}(O_r)$ in the present proposition.

29.3 An equivalent cubic-field gate when $r = 3$

For $r = 3$, the class-group formulation has an exact interpretation in terms of cubic fields. This is a specialization of the classical ring-class correspondence, not a new source of a bounded witness.

Proposition 29.4. *Let p be a prime and $B \geq 1$ a real number. The following conditions are equivalent.*

1. *There is an imaginary quadratic order discriminant $\Delta < -4$ with $|\Delta| \leq B$ and $p \nmid \Delta$, such that p splits in $K = \mathbb{Q}(\sqrt{\Delta})$ and the class g of a prime above p does not belong to $3\,\mathrm{Cl}(O_\Delta)$.*
2. *There is a complex cubic number field M with $|\mathrm{Disc}(M)| \leq B$ and $p \nmid \mathrm{Disc}(M)$ in which p is inert.*

In the forward direction, M can be chosen with $\Delta = c^2\,\mathrm{Disc}(M)$ for an integer $c \geq 1$. In the reverse direction one may take $\Delta = \mathrm{Disc}(M)$. Consequently, a uniform bound $B \leq \mathrm{poly}(\log p)$ for either existence condition would suffice for deterministic polynomial-time construction of $\mathbb{F}_{p^3}$.

Proof. Write $A = \mathrm{Cl}(O_\Delta)$. The class-field correspondence in (Bhargava and Varma, 2016, §5.1, Lemma 34 and Proposition 35, pp. 256–258) identifies index-three subgroups $H \leq A$ with isomorphism classes of cubic fields M satisfying $\Delta = c^2\,\mathrm{Disc}(M)$. In this correspondence, the cyclic cubic extension of K cut out by H is the normal closure of M over $\mathbb{Q}$. The discriminant is negative, so M has one real and two complex embeddings, and its normal closure has group $C_3 \rtimes C_2 \simeq S_3$. The involution acts on C_3 by inversion. The ray-class minus-part in the cited Lemma 34 is essential: the entire ray class group is not being identified with the ring class group.

Suppose the first condition holds. Since $A/3A$ is a vector space over $\mathbb{F}_3$, there is an index-three subgroup H not containing g. In the corresponding S_3 extension, Frobenius at p lies in the normal subgroup C_3, because p splits in K. Its image is nontrivial because $g \notin H$. It is therefore a three-cycle on the three embeddings of M, which says precisely that p is inert in M. The discriminant relation gives both $|\mathrm{Disc}(M)| \leq B$ and unramifiedness at p.

Conversely, let M satisfy the second condition, and put $\Delta = \mathrm{Disc}(M)$. Its quadratic resolvent ring is the imaginary quadratic order of discriminant Δ. The same correspondence supplies an index-three subgroup of $A = \mathrm{Cl}(O_\Delta)$ cutting out the normal closure of M over its quadratic resolvent field. Inertness means that Frobenius is a three-cycle. It is even, so p splits in the quadratic resolvent field, and it is nontrivial in the quotient A/H. Thus $g \notin H$, and in particular $g \notin 3A$. The correspondence includes nonfundamental discriminants and ramification at 3, so neither direction discards these cases.

For the algorithmic consequence, enumerate all negative order discriminants in the asserted range and apply Proposition 29.3. The work is polynomial in B and $\log p$. The existence formulation does not require a given cubic field, its defining polynomial, or the ideal class g. □

What a uniform counting theorem currently supplies. Let $N_-(B; p)$ count complex cubic fields of absolute discriminant less than B in which p is unramified and inert. Theorem 1.3 and Table 1 of (Bhargava et al., 2024, pp. 2–3 of the checked author manuscript) give, for every fixed $\epsilon > 0$, uniformly in the prime p,

$$N_-(B; p) = \frac{p^2}{12\zeta(3)(p^2 + p + 1)} B + \kappa(p) B^{5/6} + O_\epsilon\left(p^{2/3} B^{2/3+\epsilon}\right),$$

$$\kappa(p) = \frac{4\sqrt{3}\,\zeta(1/3)}{15\Gamma(2/3)^3\zeta(5/3)} \frac{1 - p^{-1/3}}{1 - p^{-5/3}}.$$

The constants in the error term depend on ϵ, not on p. The leading coefficient is bounded below by a positive absolute constant, and $\kappa(p)$ is bounded. For every fixed $\delta > 0$, this estimate therefore proves $N_-(B;p) > 0$ once $B \geq C_\delta p^{2+\delta}$. For example, take $\epsilon = \delta/(6(2+\delta))$ and compare the error with B. At $B = (\log p)^C$ for a fixed C, this error estimate does not certify positivity. This is a limitation of the displayed estimate; it does not rule out a stronger estimate, an argument using Assumption A, or another construction.

The trace still to be controlled. For any finite family $\mathcal{C}$ of complex cubic fields unramified at p, let D_M be the fundamental discriminant of the quadratic resolvent of M. Let $a_M(p)$ be the trace of the standard two-dimensional representation of S_3 at Frobenius. The three splitting types $(111), (12), (3)$ give trace values $2, 0, -1$, respectively, while the quadratic character $\chi_{D_M}(p)$ takes values $1, -1, 1$. Consequently the exact identity

$$3\#\{M \in \mathcal{C} : p \text{ is inert in } M\} = |\mathcal{C}| + \sum_{M \in \mathcal{C}} \chi_{D_M}(p) - \sum_{M \in \mathcal{C}} a_M(p)$$

isolates the remaining term. Class-number and reflection formulas enumerate the available cubic quotients; they do not, on their own, bound this last Frobenius-trace sum for a given prime p. No estimate removing the numerical $p^{2/3}$ dependence in the polylogarithmic discriminant range is established here from Assumption A.

29.4 A quadratic-unit form of the cubic CM gate

The quadratic-unit formulation makes a residue test and its output polynomial explicit. The bounded-parameter selection condition remains separate.

Lemma 29.5. *Let $p \equiv 2 \pmod 3$ be prime, let $u \in \mathbb{F}_{p^2}^*$, and suppose $u^{p+1} = n \in \{1, -1\}$. Put $t = u + u^p$. Then*

$$u \notin (\mathbb{F}_{p^2}^*)^3 \quad \Longleftrightarrow \quad X^3 - 3nX - t \text{ is irreducible over } \mathbb{F}_p.$$

The assertion includes scalar u and repeated-discriminant specializations.

Proof. If $u = w^3$ with $w \in \mathbb{F}_{p^2}^*$, then $N(w)^3 = n$. Cubing is bijective on $\mathbb{F}_p^*$, so $N(w) = n$. Thus $x = w + n/w = w + w^p$ belongs to $\mathbb{F}_p$ and satisfies $x^3 - 3nx = t$. Conversely, if $x \in \mathbb{F}_p$ satisfies this equation, choose a root $w \in \mathbb{F}_{p^2}^*$ of $Z^2 - xZ + n$. Expanding its third power shows

$$w^3 + n^3/w^3 = t.$$

Hence w^3 is a root of $Z^2 - tZ + n$, whose roots are u, u^p. In either case u is a cube in $\mathbb{F}_{p^2}$. A cubic over a field is reducible exactly when it has a root there. □

Proposition 29.6. *Let $d > 1$ be squarefree, $3 \nmid d$, and put $C = \mathbb{Q}(\sqrt{d})$ and $F = \mathbb{Q}(\sqrt{-3d})$. Write D_F for the fundamental discriminant of F and assume $3 \nmid h(F)$. Let $\epsilon > 1$ be a fundamental unit of C, and set*

$$n = N_{C/\mathbb{Q}}(\epsilon) \in \{1, -1\}, \qquad t = \mathrm{Tr}_{C/\mathbb{Q}}(\epsilon).$$

Let $p \nmid 6d$ be prime, $p \equiv -1 \pmod 9$, and suppose p splits in F. Then p is inert in C. With $u = \epsilon \bmod p \in \mathbb{F}_{p^2}^*$, the following conditions are equivalent:

1. *u is a noncube in $\mathbb{F}_{p^2}$;*
2. *$X^3 - 3nX - t$ is irreducible over $\mathbb{F}_p$;*
3. *p is inert in the complex cubic field $M = \mathbb{Q}(\epsilon^{1/3} + n\epsilon^{-1/3})$;*

4. *the prime class in the conductor-3 ring class group of F has nontrivial image in its cyclic cubic quotient.*

This field M has discriminant $9D_F$. Given such a d, the test in (1) and the output polynomial in (2) are computable deterministically in time polynomial in the numerical d and $\log p$.

Proof. Theorems 1.1 and 1.3 of Evans et al. (2025) give the discriminant assertion. To specify the dependency, put $v = \epsilon^{1/3}$ and use their ring class field for $\mathbb{Z}[\sqrt{-3d}]$. Inclusion of v there would give a cubic field of discriminant D_F, forcing $3 \mid h(F)$ by Hasse's correspondence. Thus that inclusion is excluded, and their Theorem 1.3 gives $\mathrm{Disc}(M) = 9D_F$.

Since 3 ramifies in F, $h(9D_F) = 3h(D_F)$. As $3 \nmid h(D_F)$, the conductor-3 ring class group has a unique subgroup of index three. Its fixed field in the ring class field is cyclic cubic over F and is the S_3 normal closure of M over $\mathbb{Q}$. A subgroup generated by complex conjugation fixes a cubic field isomorphic to M. For a prime p split in F, a nonidentity element in the cyclic quotient acts as a three-cycle on the embeddings of M. This proves (3)⇔(4).

We have $(-3/p) = -1$ and $(D_F/p) = 1$, so $(d/p) = -1$. This proves inertness in C and gives $N(u) = n$. The preceding lemma proves (1)⇔(2). For the remaining Frobenius equivalence, let N be the normal closure of M and put $B = FC = C(\zeta_3)$. Then N has quadratic resolvent F and

$$NB = C(\zeta_3, \epsilon^{1/3}).$$

The residue field of B at the chosen prime above p is $\mathbb{F}_{p^2}$. Its cubic Kummer Frobenius is nontrivial exactly when u is a noncube. This Frobenius restricts to the square of the Frobenius in $\mathrm{Gal}(N/F)$; squaring preserves nontriviality in a group of order 3. Since p splits in F, nontriviality is precisely the inert splitting type in M. This proves (1)⇔(3), without making an unjustified assumption that p is prime to the index of the displayed generator of M.

Continued fractions compute the fundamental unit in time polynomial in numerical d, with coefficient bit lengths polynomial in d. Arithmetic in $\mathbb{F}_p[Y]/(Y^2 - d)$ then tests $u^{(p^2-1)/3} \neq 1$ by exponentiation. The integers n and t give the stated output polynomial. □

This gives a sufficient all-input condition: for each hard prime p, find a squarefree $d \leq \mathrm{poly}(\log p)$ with $3 \nmid d$, $3 \nmid h(\mathbb{Q}(\sqrt{-3d}))$, p split in that imaginary quadratic field, and the fundamental unit noncubic modulo its inert prime in $\mathbb{Q}(\sqrt{d})$. Quadratic splitting, class-number indivisibility, and this particular Frobenius condition are distinct requirements.

Beckwith, Raum, and Richter prove that for every prime $\ell \geq 2$ and finite set S_+ of odd primes, some imaginary quadratic field has class number prime to ℓ and splits at every prime in S_+ (Beckwith et al., 2024, Theorem A). Their result includes prescribed split primes $p \equiv -1 \pmod{\ell}$. The stated theorem and its proof do not bound the least discriminant in terms of S_+, do not prescribe ramification at 3, and do not test the nontrivial cubic Frobenius required above. Accordingly that qualitative existence theorem does not supply the bounded-d selector used here. Likewise the reciprocity formulas of (Evans et al., 2025, Theorem 1.6) and (Evans and Van Veen, 2026, Theorem 1.1) concern their specified $p \equiv 1 \pmod 3$ setting and are not used as an all-input selection bound in the present inert setting.

30 Small searches, fixed candidates, and Jacobi sums

This section compares concrete nonresidue searches with the auxiliary-prime gate and proves limitations of fixed candidate families.

30.1 Small cubic nonresidues recover the auxiliary prime

The following reduction identifies a precise limitation of a search through small coefficients in the first unresolved prime-degree case. It is a consequence of classical cubic reciprocity.

Proposition 30.1 (Cubic reciprocity and coefficient height)**.** *Let $p \equiv -1 \pmod 9$ be prime, put $\mathcal{O} = \mathbb{Z}[\zeta]$ with $\zeta^2 + \zeta + 1 = 0$, and identify $\mathcal{O}/(p)$ with $\mathbb{F}_{p^2}$. Suppose that the image of $\alpha = A + B\zeta$, where $A, B \in \mathbb{Z}$, is a nonzero noncube in $\mathcal{O}/(p)$. Then there is a rational prime ℓ such that*

$$\ell \equiv 1 \pmod 3, \qquad p^{(\ell-1)/3} \not\equiv 1 \pmod \ell,$$
$$\ell \le A^2 - AB + B^2 \le 3H^2, \quad H = \max\{|A|, |B|\}.$$

Conversely, every rational prime $\ell \equiv 1 \pmod 3$ satisfying $p^{(\ell-1)/3} \not\equiv 1 \pmod \ell$ admits integers A, B such that

$$A^2 - AB + B^2 = \ell, \qquad |A|, |B| \le 2\sqrt{\ell/3},$$

and every element of norm ℓ has noncube image in $\mathcal{O}/(p)$.

Consequently a bound polynomial in $\log p$ for the height of a cubic nonresidue in this quadratic presentation is equivalent, up to a change of polynomial, to such a bound for an auxiliary prime with the displayed property. The corresponding exhaustive searches are deterministic and polynomial in those numerical bounds and $\log p$.

Proof. We fix the normalization explicitly. A prime element $\pi \in \mathcal{O}$ coprime to 3 is called primary here when $\pi \equiv 2 \pmod{3\mathcal{O}}$. Each prime element not above 3 has a unique primary associate. For a prime element π not above 3, and an element γ coprime to π, write

$$\left(\frac{\gamma}{\pi}\right)_3 \equiv \gamma^{(N\pi-1)/3} \pmod \pi, \qquad \left(\frac{\gamma}{\pi}\right)_3 \in \{1, \zeta, \zeta^2\}.$$

This definition uses $N(p) = p^2$ for an inert rational prime p. The cubic reciprocity law in Friesen et al. (1986, equation (1.1), p. 395; the split–inert proof is in Section 2, pp. 396–397) says that

$$\left(\frac{\pi}{\varpi}\right)_3 = \left(\frac{\varpi}{\pi}\right)_3$$

for primary primes of distinct norm.

The ring $\mathcal{O}$ is Euclidean for $N(A + B\zeta) = A^2 - AB + B^2$, hence is a unique factorization domain. Its prime elements, up to associates, are $1 - \zeta$, the rational primes $q \equiv 2 \pmod 3$, and the elements of rational prime norm $\ell \equiv 1 \pmod 3$. We first eliminate all factors except the last type. Every nonzero rational integer modulo p is a cube already in $\mathbb{F}_p$, because $\gcd(3, p-1) = 1$. Also ζ is a cube in $\mathbb{F}_{p^2}$: its cubic character is $\zeta^{(p^2-1)/3} = 1$ since $9 \mid p^2 - 1$. Thus all six units $\{\pm 1, \pm\zeta, \pm\zeta^2\}$ have cube image. Finally

$$(1 - \zeta)^2 = -3\zeta$$

is a cube. The quotient $\mathbb{F}_{p^2}^\times / (\mathbb{F}_{p^2}^\times)^3$ has order 3, so squaring is injective on it; therefore $1 - \zeta$ is itself a cube. These statements cover units, the ramified prime above 3, and every inert rational prime factor other than p. The factor p cannot divide α, since its residue is nonzero.

Factor α in $\mathcal{O}$. Multiplicativity of the cubic character, together with its nontrivial value on α, implies that some prime factor π of norm $\ell \equiv 1 \pmod 3$ has $\left(\frac{\pi}{p}\right)_3 \neq 1$. Multiplying π by a unit makes it primary without changing this value or its norm. The rational prime p is itself primary, since $p \equiv 2 \pmod 3$; it is coprime to π and has norm $p^2 \neq \ell$. Reciprocity therefore gives

$$\left(\frac{p}{\pi}\right)_3 = \left(\frac{\pi}{p}\right)_3 \neq 1.$$

The residue field $\mathcal{O}/(\pi)$ has ℓ elements, so this is exactly $p^{(\ell-1)/3} \not\equiv 1 \pmod \ell$. In particular $\ell \notin \{3, p\}$. Since $\pi \mid \alpha$, its norm divides the positive integer $N\alpha$, proving $\ell \le N\alpha \le 3H^2$.

Conversely, a rational prime $\ell \equiv 1 \pmod 3$ splits in $\mathcal{O}$, and unique factorization gives an element π of norm ℓ. Normalize π to be primary. The hypothesis on p means $\left(\frac{p}{\pi}\right)_3 \neq 1$ in

the residue field $\mathcal{O}/(\pi) \cong \mathbb{F}_\ell$. Reciprocity gives $\left(\frac{\pi}{p}\right)_3 \neq 1$. The same argument applies to the conjugate prime, and multiplying by a unit preserves the cube class modulo p. Every element of norm ℓ is an associate of one of these two primes, so every such element is a noncube modulo p. Writing it as $A + B\zeta$, the identities

$$N\pi = (A - B/2)^2 + 3B^2/4 = (B - A/2)^2 + 3A^2/4$$

prove the claimed coordinate bounds.

For the algorithmic assertion, from a height bound H one may enumerate rational primes at most $3H^2$ and test the displayed congruence by modular powering; the first part proves success. From an auxiliary prime ℓ, enumerate integer pairs in the displayed box until $A^2 - AB + B^2 = \ell$. There are $O(\ell)$ pairs, and the second part proves both existence and the required noncube property. These running times are polynomial in the numerical search bounds and $\log p$; no claim of a running time polynomial in $\log H$ or $\log \ell$ is needed. □

Proposition 30.2 (Exact density on an affine line)**.** *Let p be prime, $r > 1$ an odd divisor of $p + 1$, and $\chi : \mathbb{F}_{p^2}^\times \to \mu_r$ a character of exact order r. For $\alpha \in \mathbb{F}_{p^2} \setminus \mathbb{F}_p$,*

$$\sum_{t \in \mathbb{F}_p} \chi(\alpha + t) = -1.$$

More precisely, the value 1 occurs $(p+1)/r - 1$ times, and each other value in μ_r occurs $(p+1)/r$ times. For $r = 3$, exactly $2(p + 1)/3$ members of this affine line are noncubes.

Proof. The restriction of χ to $\mathbb{F}_p^\times$ is trivial: its order divides both the odd integer $r \mid p + 1$ and $p - 1$, whose greatest common divisor divides 2. Thus χ induces a surjective character on the cyclic quotient $\mathbb{F}_{p^2}^\times/\mathbb{F}_p^\times$, of order $p + 1$. The p elements $\alpha + t$ and the element 1 are one representative of each quotient class. Indeed, an equality $\alpha + t = c(\alpha + u)$ with $c \in \mathbb{F}_p^\times$ forces $c = 1$ and $t = u$, and no $\alpha + t$ is in $\mathbb{F}_p$. Every character value appears $(p + 1)/r$ times on the full quotient. Removing the representative 1 proves both assertions. □

This density statement concerns the entire affine line. It gives no bound polynomial in $\log p$ on the first successful integer parameter t. For $p \equiv -1 \pmod 9$, a successful parameter t in $\zeta + t$ already yields, by Proposition 30.1, a good auxiliary prime $\ell \leq t^2 - t + 1$.

Proposition 30.3 (An explicit cubic candidate family)**.** *Let $p \equiv 2 \pmod 3$ be an odd prime, put $\mathbb{F}_{p^2} = \mathbb{F}_p(\zeta)$, and let $u \in \mathbb{F}_p$. Set*

$$v = u^2 - u + 1, \qquad t_u = 2 - \frac{3}{v}, \qquad f_u(Y) = Y^3 - 3Y - t_u.$$

Then $v \neq 0$, and f_u is irreducible over $\mathbb{F}_p$ if and only if $u + \zeta$ is a noncube in $\mathbb{F}_{p^2}$.

Proof. The quadratic $U^2 - U + 1$ has no root in $\mathbb{F}_p$ when $p \equiv 2 \pmod 3$ and $p \neq 2$. Put $\alpha = u + \zeta$ and $b = \alpha/\alpha^p$. Its norm is 1, and a direct calculation gives

$$b + b^{-1} = \frac{2u^2 - 2u - 1}{u^2 - u + 1} = t_u.$$

For the cubic character χ of $\mathbb{F}_{p^2}$, $\chi(\alpha^p) = \chi(\alpha)^p = \chi(\alpha)^{-1}$. Hence $\chi(b) = \chi(\alpha)^2$, so b is a cube exactly when α is a cube.

If $b = x^3$ for $x \in \mathbb{F}_{p^2}$, then $N(x)^3 = N(b) = 1$. Since the cube map is bijective on $\mathbb{F}_p^\times$, $N(x) = 1$. Thus $y = x + x^{-1}$ is fixed by pth-power Frobenius and belongs to $\mathbb{F}_p$. The identity $y^3 - 3y = x^3 + x^{-3} = t_u$ gives a root of f_u. Conversely, if $y \in \mathbb{F}_p$ is a root of f_u, a root x of $X^2 - yX + 1$ belongs to $\mathbb{F}_{p^2}$, is nonzero, and satisfies $x^3 + x^{-3} = t_u = b + b^{-1}$. Therefore x^3 is one of b, b^{-1}, so b is a cube. A cubic over a field is irreducible exactly when it has no root in that field, completing the proof. □

30.2 Restrictions on all-parameter rational noncube families

Proposition 30.4. *Let $p > 3$ be prime with $p \equiv 2 \pmod 3$, and put $q = p^2$. Let $C = A/B \in \mathbb{F}_q(T)^*$, where A, B are relatively prime polynomials of degrees at most D. Suppose*

$$p > (8D + 4)^2$$

and $C(t)$ is a noncube for every $t \in \mathbb{F}_p$ at which it is finite and nonzero. Then

$$C(T) = c\, b(T)\, g(T)^3, \qquad b \in \mathbb{F}_p(T)^*, \quad g \in \mathbb{F}_q(T)^*,$$

where

$$c = \operatorname{lc}(A)/\operatorname{lc}(B) \in \mathbb{F}_q^* \setminus (\mathbb{F}_q^*)^3.$$

Conversely every function of this form has noncubic value at every base-field argument at which it is finite and nonzero.

Proof. Write a bar for coefficientwise p-Frobenius, and put $R = \overline{C}/C$. Suppose first that R is not a cube over $\overline{\mathbb{F}}_p(T)$. The smooth projective model of

$$W^3 = R(T)$$

is geometrically connected. The semilinear involution that conjugates constants, fixes T, and sends W to W^{-1} descends this curve to $\mathbb{F}_p$. The reduced numerator and denominator of R have degrees at most $2D$; moreover their degrees are equal, so infinity is unramified. There are at most $4D$ branch points. The tame degree-three Riemann–Hurwitz formula therefore gives $g \le 4D - 2$.

For a valid $t \in \mathbb{F}_p$, put $z = C(t)$. The fibre equation is $w^3 = z^{p-1}$. It has a solution in $\mathbb{F}_q$ exactly when z is a cube, since $p - 1 \equiv 1 \pmod 3$. Every such solution has norm one: its norm has cube one, and cubing is bijective on $\mathbb{F}_p^*$. These are precisely the $\mathbb{F}_p$-rational points of the descended fibre. By hypothesis there are none. All rational points must therefore lie above the at most $2D$ finite zeros and poles of C, or above infinity. There are at most three points above each, giving

$$\#X(\mathbb{F}_p) \le 3(2D + 1).$$

On the other hand, Hasse–Weil gives

$$\#X(\mathbb{F}_p) \ge p + 1 - 2g\sqrt{p} \ge p + 1 - (8D - 4)\sqrt{p}.$$

For $\sqrt{p} > 8D + 4$ this exceeds $3(2D + 1)$, a contradiction. The Kummer and Hasse–Weil inputs are standard; see, for example, (Bartoli et al., 2024, Section 1.1, Theorems 1.1–1.2).

It follows that the divisor of R is divisible by three. Factor C into monic irreducibles over $\mathbb{F}_q$:

$$C = c \prod_f f^{e_f}, \qquad c = \operatorname{lc}(A)/\operatorname{lc}(B).$$

The divisor condition says $e_f \equiv e_{\bar f} \pmod 3$. Choose a common representative in $\{0, 1, 2\}$ on each conjugation orbit. The products over these orbits, with those representative exponents, give a monic polynomial $b \in \mathbb{F}_p[T]$. The remaining exponents are divisible by three and give a monic rational function g with $C = cbg^3$.

There is at least one valid base-field argument because $p > 2D$. At any such argument the local orders of b and g can cancel, but their leading units still give $C(t) = c\, b_0 g_0^3$ with $b_0 \in \mathbb{F}_p^*$ and $g_0 \in \mathbb{F}_q^*$. Every element of $\mathbb{F}_p^*$ is a cube in $\mathbb{F}_q^*$. Thus the cubic class of $C(t)$ is exactly that of c. This proves that c is a noncube and also proves the converse, including possible cancelling zeros or poles in the displayed factorization. □

The base-field factor is essential. If a is any noncube, then $C(T) = aT$ has noncubic values at every nonzero base-field argument, although C is not a constant times a rational cube. The proposition therefore factors out both base-field rational functions and cubes. It says that an all-parameter family of the stated small degree already exposes a noncube as its leading-coefficient ratio. It does not bound the least successful parameter in a family taking both cubic and noncubic values.

Proposition 30.5. *Let $p > 3$ be prime with $p \equiv -1 \pmod 9$ and put*

$$h_k(Y) = Y^3 - (k+3)Y^2 + kY + 1.$$

Let $m \geq 1$ be prime to p, and let $(a, m) = 1$. There are infinitely many rational primes $k \equiv a \pmod m$ for which $k \not\equiv 3 \pmod p$ and h_k is irreducible over $\mathbb{F}_p$. There are also infinitely many such prime parameters for which h_k splits completely over $\mathbb{F}_p$.

Proof. Choose $\zeta_3 \in \mathbb{F}_{p^2}$. The established Tate-family identity says that h_k is irreducible exactly when $k - 3\zeta_3$ is a noncube. The lines represented by $k - 3\zeta_3$, for $k \in \mathbb{F}_p$, are precisely the p nonscalar lines of $\mathbb{F}_{p^2}^*/\mathbb{F}_p^*$. This quotient is cyclic of order $p + 1$; its cube subgroup has index three, and the omitted scalar line belongs to it. Hence exactly $(p - 2)/3$ parameters are cubic and $2(p + 1)/3$ are noncubic.

The parameter $k = 0$ is cubic because ζ_3 is a cube when $9 \mid p + 1$. Thus some noncubic parameter is different from both 0 and 3. A nonzero nonsingular split parameter is $k = -3$, since $-3 - 3\zeta_3 = 3\zeta_3^2$ is a cube. Fix one residue of each kind. The Chinese remainder theorem combines it with $k \equiv a \pmod m$ to give a reduced class modulo mp. Dirichlet's theorem gives infinitely many primes in each class. □

Thus a single rational congruence condition of modulus prime to p cannot force the status of all prime parameters in this family. The conclusion has no least-prime bound and does not exclude a bounded-window selector. Cubic reciprocity instead makes a fixed integer parameter into a cubic Dirichlet character of the varying input prime; that orientation is compatible with the preceding almost-all theorem but supplies no pointwise selection theorem.

30.3 Two possible transfers that do not identify the desired witness

The following checks distinguish a rational character at a selected input from a character that detects the required finite-field output.

Proposition 30.6. *A small prime detected by a cubic character associated with $p - 3\zeta$ need not supply an irreducible polynomial h_k modulo p. For $p = 17$, the norm of $17 - 3\zeta$ is the prime 349. The prime 7 is a cubic nonresidue modulo 349, since*

$$7^{116} \equiv 226 \pmod{349}.$$

Nevertheless

$$h_7(Y) \equiv (Y - 2)(Y - 9)(Y - 16) \pmod{17}.$$

Proof. The norm is $17^2 + 3 \cdot 17 + 9 = 349$; trial division by $2, 3, 5, 7, 11, 13, 17$ proves primality. Repeated squaring gives the displayed power, and multiplying the three linear factors gives $Y^3 - 10Y^2 + 7Y + 1$ modulo 17. A nontrivial cubic character modulo 349 therefore detects 7, but its detection does not imply that h_7 is irreducible at the prescribed input. □

Proposition 30.7. *Let a number field L of degree N split completely at a prime p, and suppose $N > p$. For every integral primitive element α of L, the index $[\mathcal{O}_L : \mathbb{Z}[\alpha]]$ is divisible by p.*

Proof. Complete splitting gives $\mathcal{O}_L/p\mathcal{O}_L \cong \mathbb{F}_p^N$. The subalgebra generated by a single element of this product has dimension equal to the number of its distinct coordinates, which is at most p. It cannot be the whole algebra when $N > p$. If the displayed index were prime to p, reduction of $\mathbb{Z}[\alpha] \subset \mathcal{O}_L$ modulo p would be surjective, contradicting this dimension bound. □

Thus a large compositum splitting completely at p is not contradicted by counting roots of an arbitrary primitive-element polynomial. A use of that polynomial must also control its index.

30.4 A finite comparison of cubic witnesses

The following computation illustrates the different sufficient conditions; it provides no bound beyond its stated range. For all 1,599 primes $p \le 100{,}000$ with $p \equiv -1 \pmod 9$, we compared exact modular tests for the six quantities below. The cubic menu consists of monic polynomials $X^3 + aX^2 + bX + c$ with $a \in \{-1, 0, 1\}$, $-12 \le b \le 12$, $1 \le c \le 24$, negative polynomial discriminant, and no rational root. These restrictions leave 1,451 entries, ordered by absolute polynomial discriminant and then lexicographically by coefficients. A successful entry must be unramified and irreducible modulo p.

Quantity minimized for each prime	Largest minimum	First prime
Good auxiliary prime ℓ	43	5,669
Eisenstein height $\max(\lvert A\rvert, \lvert B\rvert)$	5	5,669
Affine parameter $t \ge 0$ with $t + \zeta_3$ noncubic	11	38,501
Absolute polynomial discriminant in the stated menu	204	41,579
Positive m with $X^3 + X + m$ irreducible	23	96,857
Positive m with $X^3 - X + m$ irreducible	26	35,099

Table 2: Exact cubic searches for the stated finite range.

At $p = 41{,}579$, the first successful menu entry is $X^3 + X^2 + X + 3$, of discriminant -204. Its field discriminant divides the polynomial discriminant, but maximal orders and least field discriminants were not computed. The program *Finite comparison of cubic construction witnesses* records every prime and witness. Its irreducibility test uses $\gcd(f, X^p - X)$; three hundred small cases were also compared with direct root enumeration. The fixed-menu result below explains why success of any one finite menu on this interval cannot establish success at every prime.

30.5 Restrictions on fixed candidates and global units

Proposition 30.8 (A finite menu of cyclotomic candidates)**.** *Fix an odd prime r, an integer $a \ge 2$, and finitely many nonzero elements $\alpha_1, \dots, \alpha_m$ of $K = \mathbb{Q}(\zeta_r)$. There is a positive-density set of rational primes p such that*

$$\operatorname{ord}_r(p) = 2, \qquad v_r(p^2 - 1) = a,$$

and every α_j reduces to an rth power in every residue field of K above p. The finitely many primes at which a chosen element has zero, pole, or ramified reduction are omitted.

Proof. Let M be obtained from K by adjoining rth roots of all $\mathbb{Q}$-conjugates of every α_j. Then $M/\mathbb{Q}$ is Galois, and M/K is abelian of exponent dividing r by elementary Kummer theory. Write $C_b = \mathbb{Q}(\zeta_{r^b})$ and $L = MC_a$. The group $\operatorname{Gal}(L/K)$ embeds in

$$\operatorname{Gal}(M/K) \times \operatorname{Gal}(C_a/K);$$

its exponent therefore divides r^{a-1}. Since C_{a+1}/K is cyclic of degree r^a, this implies

$$L \cap C_{a+1} = C_a.$$

Indeed, the intersection contains C_a, and any larger intersection would be all of C_{a+1}, giving a cyclic quotient of exponent r^a.

Let c denote complex conjugation on L. Choose any $t \in \{1, \dots, r-1\}$, and let γ be the automorphism of C_{a+1} with cyclotomic exponent $-1 + tr^a \pmod{r^{a+1}}$. On the intersection C_a, both c and γ act by inversion. They consequently extend to a common automorphism g of LC_{a+1}. By the qualitative Chebotarev density theorem, a positive-density set of unramified rational primes has Frobenius conjugate to g. Every such prime satisfies

$$p \equiv -1 + tr^a \pmod{r^{a+1}},$$

because the cyclotomic Galois group is abelian. Thus $p \equiv -1 \pmod{r}$ and $v_r(p+1) = a$.

The restriction of each Frobenius representative to M is a conjugate of complex conjugation and has square 1. Each prime of K above p has residue degree 2. Its relative Frobenius in M/K is therefore trivial. Consequently its residue field extension in M is trivial: the rth roots adjoined to M reduce into that same $\mathbb{F}_{p^2}$ residue field. This proves the asserted rth-power property at every prime of K above p. □

The preceding proposition treats fixed elements of $\mathbb{Q}(\zeta_r)$. Rational shifts of a maximal r-power root lie at a higher cyclotomic level. The following statement keeps that root maximal in the residue field and makes all the prescribed character tests scalar on a split block of rank r.

Proposition 30.9 (Fixed shifts of maximal cyclotomic roots)**.** *Fix an odd prime r, an integer $a \geq 2$, and a finite set $\mathcal{T} \subset \mathbb{Q}$. Put $R = r^a$. There is a set $\mathcal{P}$ of rational primes of positive natural density among the rational primes such that every $p \in \mathcal{P}$ avoids the denominators of $\mathcal{T}$ and satisfies*

$$p \equiv -1 + r^a \pmod{r^{a+1}}.$$

For every element $\alpha \in \mathbb{F}_{p^2}$ of exact order R, put $\zeta = \alpha^{R/r}$. Then, for every $t \in \mathcal{T}$ and $0 \leq i < r$, the element $t + \alpha\zeta^i$ is nonzero and

$$(t + \alpha\zeta^i)^{(p^2-1)/r} = (t + \alpha)^{(p^2-1)/r}.$$

In particular, $\mathrm{ord}_r(p) = 2$ *and $v_r(p^2-1) = a$, so α has maximal r-power order in $\mathbb{F}_{p^2}^\times$.*

Equivalently, in the split $\mathbb{F}_{p^2}$-algebra

$$B_\alpha = \mathbb{F}_{p^2}[X]/(X^r - \alpha^r), \qquad u = X \bmod (X^r - \alpha^r),$$

each $t + u$ is a unit and $(t+u)^{(p^2-1)/r}$ belongs to the diagonal copy of $\mathbb{F}_{p^2}$.

Proof. Choose a complex primitive Rth root α_0, and write

$$H = \mathbb{Q}(\alpha_0), \qquad K = \mathbb{Q}(\alpha_0^r), \qquad \zeta_0 = \alpha_0^{R/r}.$$

The extension H/K is cyclic of degree r, with automorphisms $\alpha_0 \mapsto \alpha_0\zeta_0^j$ for $0 \leq j < r$. For each $t \in \mathcal{T}$, each $v \in (\mathbb{Z}/R\mathbb{Z})^\times$, and $0 \leq i < r$, set

$$w_{t,v,i} = \frac{t + \alpha_0^v\zeta_0^i}{t + \alpha_0^v} \in H^\times.$$

These elements are defined and nonzero: every $\alpha_0^v\zeta_0^i$ is a primitive Rth root, hence is nonrational because $R \geq 9$. The family contains every $\mathbb{Q}$-conjugate of every member. Indeed the

automorphism with exponent b sends $w_{t,v,i}$ to $w_{t,vb,ib}$, with the last index taken modulo r. Moreover,

$$N_{H/K}(w_{t,v,i}) = 1.$$

To see this, the conjugates of the numerator have the form $t + \alpha_0^v \zeta_0^{vj+i}$, and those of the denominator have the form $t + \alpha_0^v \zeta_0^{vj}$. Since v is prime to r, these are the same r factors.

Let L be obtained from H by adjoining an rth root of every $w_{t,v,i}$. Because $\mu_r \subset H$ and the family is stable under all automorphisms of $H/\mathbb{Q}$, the extension $L/\mathbb{Q}$ is Galois. The action on the chosen radicals embeds $G = \mathrm{Gal}(L/H)$ into a product of copies of μ_r.

The class of α_0 does not lie in the subgroup of $H^\times/H^{\times r}$ generated by the classes of the $w_{t,v,i}$. Otherwise taking H/K norms of an identity

$$\alpha_0 = h^r \prod_{t,v,i} w_{t,v,i}^{e_{t,v,i}}, \qquad h \in H^\times,$$

would give

$$\alpha_0^r = N_{H/K}(h)^r.$$

Here $N_{H/K}(\alpha_0) = \alpha_0^r$, because r is odd. But α_0^r is a primitive r^{a-1}st root, and all its rth roots are primitive r^ath roots. None belongs to $K = \mathbb{Q}(\zeta_{r^{a-1}})$, so the displayed equality is impossible.

This class exclusion also proves that L contains no rth root of α_0. In detail, if $\theta^r = \alpha_0$ with $\theta \in L$, then $\sigma \mapsto \sigma(\theta)/\theta$ is a character of G into μ_r. Regarding the embedding of G into a product μ_r^s as an inclusion of $\mathbb{F}_r$-vector spaces, extend that character to μ_r^s. A product of powers of the chosen radical generators then has the same character as θ. Their quotient lies in H, and taking rth powers expresses the class of α_0 as a product of the classes of the $w_{t,v,i}$, a contradiction. Consequently, for $H_1 = \mathbb{Q}(\zeta_{r^{a+1}})$, we have

$$L \cap H_1 = H.$$

Indeed H_1/H has prime degree r and is obtained by adjoining an rth root of α_0.

Let c denote complex conjugation on L, and let τ be the automorphism of H_1 with exponent $-1 + r^a \pmod{r^{a+1}}$. Both restrict to inversion on H. The restriction map from $\mathrm{Gal}(LH_1/\mathbb{Q})$ identifies it with the pairs of automorphisms of L and H_1 that agree on their intersection H. Thus there is an automorphism g restricting to (c, τ). By the qualitative Chebotarev theorem (Serre, 1981, Section 2.1, Theorem 1, p. 131), the unramified rational primes whose Frobenius is conjugate to g have natural density $|\mathrm{Cl}(g)|/[LH_1 : \mathbb{Q}] > 0$. Since $H_1/\mathbb{Q}$ is abelian, every such prime satisfies

$$p \equiv -1 + r^a \pmod{r^{a+1}}.$$

It follows that $\mathrm{ord}_r(p) = 2$, $v_r(p+1) = a$, and $v_r(p-1) = 0$. The restriction of Frobenius to L is a conjugate of c, of order 2. Hence every residue field of L above p is $\mathbb{F}_{p^2}$.

Write $t = m_t/n_t$ in lowest terms with $n_t > 0$. In addition to the ramified primes, omit the primes dividing the nonzero integer

$$r \prod_{t \in \mathcal{T}} n_t \prod_{\substack{t \in \mathcal{T} \\ v \in (\mathbb{Z}/R\mathbb{Z})^\times}} \left| N_{H/\mathbb{Q}}(m_t + n_t \alpha_0^v) \right|.$$

This is a finite exclusion and leaves the density positive. Every numerator and denominator of every $w_{t,v,i}$ is now a local unit: the exponents $v + iR/r$ are again units modulo R. The chosen rth roots of the $w_{t,v,i}$ are local units as well, so they reduce into the same $\mathbb{F}_{p^2}$ residue field and prove that every reduced ratio $w_{t,v,i}$ is an rth power there.

Reduction preserves the order R of α_0 because $p \nmid R$. For example, the discriminant of $X^R - 1$ is a unit at p, so its distinct integral roots remain distinct after reduction. The reductions of the powers α_0^v, with v prime to R, are therefore all the primitive Rth roots in $\mathbb{F}_{p^2}$. If α is the

reduction of α_0^v, the root ζ in the statement is the reduction of ζ_0^v. The ratio indexed by vi modulo r therefore tests its ith component; multiplication by v permutes those indices. The reduced ratios being rth powers proves the claimed character equality for every such root and every shift. The exclusions also prove nonvanishing.

Finally $X^r - \alpha^r$ has the distinct roots $\alpha\zeta^i$ in $\mathbb{F}_{p^2}$. Evaluation at these roots identifies B_α with a product of r copies of $\mathbb{F}_{p^2}$. The displayed character equality says that $(t+u)^{(p^2-1)/r}$ has equal coordinates, which is precisely the asserted scalar property. □

The statement concerns a fixed rational menu and the corresponding character-value tests. The common value may be a nontrivial rth root of unity. The earlier Proposition 30.8 concerns rth-power values of elements in $\mathbb{Q}(\zeta_r)$; the proposition above treats ratios at the higher level $\mathbb{Q}(\zeta_{r^a})$ while retaining the exact valuation a. Neither statement bounds the first prime exhibiting this behavior in terms of the menu, or rules out a menu chosen using p, a menu growing with $\log p$, or other additive operations.

Proposition 30.10 (A finite menu of fixed polynomials)**.** *Fix an odd prime r, an integer $a \geq 2$, and finitely many separable polynomials in $\mathbb{Q}[X]$, each of positive odd degree at most r. There is a positive-density set of primes p with $\operatorname{ord}_r(p) = 2$ and $v_r(p^2-1) = a$ such that every polynomial in the menu has a root modulo p. In particular, a fixed finite menu of degree-r polynomials cannot cover all primes in this exact-a branch by irreducible reductions.*

Proof. Let M be the compositum of $K = \mathbb{Q}(\zeta_r)$ and the splitting fields of the polynomials. The group $\mathrm{Gal}(M/K)$ embeds in a product of subgroups of S_r. The exponent of S_r is $\mathrm{lcm}(1,\dots,r)$, whose r-adic valuation is 1. Hence the r-primary part of the exponent of $\mathrm{Gal}(M/K)$ divides r. For $L = M\mathbb{Q}(\zeta_{r^a})$, the r-primary part of the group exponent divides r^{a-1}. The intersection and compatibility argument in the preceding proof therefore applies without requiring M/K to be abelian. It gives positive-density primes of the prescribed exact valuation whose Frobenius on M is a conjugate of complex conjugation.

Complex conjugation is an involution on the distinct roots of each polynomial. An involution on an odd finite set has a fixed point. After omitting the primes dividing denominators, leading coefficients, and discriminants, Frobenius fixed points are exactly the linear factors of the reduction. Thus every polynomial in the menu has a root. □

These statements concern fixed finite menus. They do not rule out an adaptive deterministic algorithm, a menu that grows with $\log p$, or an element whose coefficients depend on previously computed data.

Proposition 30.11 (Global r-units have small Teichmüller projection)**.** *Let r be an odd prime, $K = \mathbb{Q}(\zeta_r)$, and $u \in \mathcal{O}_K[1/r]^\times$. Let $p \neq r$ be prime, put*

$$s = \operatorname{ord}_r(p), \quad q = p^s, \quad R = r^a, \quad a = v_r(q-1),$$

$$w = (q-1)/R, \qquad D = \mathbb{F}_p[T]/\Phi_r.$$

Reduce u to a unit of D, and define

$$\mathcal{P}(u) = \prod_{b \in (\mathbb{Z}/r\mathbb{Z})^\times} \sigma_b(u^w)^{\omega_a(b)^{-1}}, \qquad \sigma_b(T) = T^b,$$

with inverse exponents taken modulo R. Then $\mathcal{P}(u)$ has order dividing r. It cannot be a maximal Teichmüller unit when $a \geq 2$. If in addition s is even and $a \geq 2$, then the residue of u is itself an rth power in every field component of D.

Proof. The unique prime of K above r is stable under complex conjugation. Since the divisor of u is supported there, all finite valuations of $u/\bar{u}$ vanish. Thus $u/\bar{u}$ is an algebraic unit. Every complex conjugate has absolute value 1, because K is a CM field and its embeddings commute with complex conjugation. An algebraic integer all of whose conjugates have absolute value 1 is a root of unity: the monic polynomials whose roots are the conjugates of its positive integral powers have bounded integer coefficients and bounded degree, so only finitely many such powers occur. Applied here, this gives

$$u/\bar{u} \in \mu(K) = \mu_{2r}.$$

For completeness, the last equality uses the familiar ramification and degree facts for prime-conductor cyclotomic fields. A root of unity of order divisible by an odd prime other than r, or by 4, would give a subfield ramified at a prime other than r. A root of order divisible by r^2 would give a subfield of degree at least $r(r-1) > [K : \mathbb{Q}]$. Neither is possible, and K already contains μ_{2r}.

The map $\mathcal{P}$ is a group homomorphism into the r-primary part of $D^\times$ and has Teichmüller equivariance. In particular, since $\omega_a(-1) = -1 \pmod{R}$,

$$\mathcal{P}(\bar{u}) = \mathcal{P}(u)^{-1}, \qquad \mathcal{P}(u)^2 = \mathcal{P}(u/\bar{u}).$$

The last term has order dividing r: the r-primary part of the reduction of a root in μ_{2r} has order dividing r. Squaring is an automorphism of any finite group of odd r-power order, so $\mathcal{P}(u)$ also has order dividing r.

For the final assertion, suppose s is even. The unique order-two element of $(\mathbb{Z}/r\mathbb{Z})^\times$ is -1, whence $p^{s/2} \equiv -1 \pmod{r}$. Consequently, in every component $\mathbb{F}_q$, complex conjugation acts as $p^{s/2}$th-power Frobenius. The equation $u/\bar{u} \in \mu_{2r}$ then implies that the r-primary order of $u^{p^{s/2}-1}$ is at most r. Since $r \nmid p^{s/2} - 1$, the same is true of u. The r-primary subgroup of $\mathbb{F}_q^\times$ is cyclic of order r^a. When $a \geq 2$, all of its elements of order at most r are rth powers, while the rth-power map is bijective on the prime-to-r part. Hence u is an rth power in that component. □

The proposition applies to the entire global unit group and to the units obtained by allowing denominators at r. For Gauss or Jacobi sums it applies only when the sum in question lies in K and has divisor supported above r; arbitrary such sums need not satisfy those hypotheses.

30.6 Jacobi sums and the semiprimitive case

Point counting can sometimes construct algebraic constants, but the splitting conditions in that use of point counting are consequential. For a fixed odd prime ℓ, Pila gives an unconditional algorithm polynomial in $\log p$ for the ℓth roots of unity in $\mathbb{F}_p$ when $p \equiv 1 \pmod{\ell}$ (Pila, 1990, Theorem D, author p. 2). His argument uses Jacobi sums to distinguish the prime ideals of $\mathbb{Q}(\zeta_\ell)$ above the completely split prime p (Pila, 1990, §5, especially Propositions 5.1–5.3, author pp. 14–18). The exponent and constants in his general point-counting theorem also depend on the defining equations and their degrees (Pila, 1990, Theorem A, author p. 1); the statement is not a uniform polynomial bound in a growing ℓ.

The following elementary Frobenius symmetry explains a further restriction on this particular Jacobi-sum strategy. It is the standard semiprimitive phenomenon; the same symmetry appears in Akiyama (1996, Lemma 1, author p. 3) for a quadratic second character. We give the two-character form needed here.

Proposition 30.12. *Let $q = p^{2j}$, where p is prime and $j \geq 1$, and let $m > 2$ divide $p^j + 1$. Let χ be a multiplicative character of exact order m on $\mathbb{F}_q^\times$. If $a, b, a+b$ are all nonzero modulo m, then*

$$J(\chi^a, \chi^b) := \sum_{x \in \mathbb{F}_q} \chi^a(x)\chi^b(1-x) \;=\; p^j,$$

where all characters are extended by zero at zero. Consequently these values generate no nontrivial number field, and their principal ideals do not distinguish primes above p.

Proof. Put $Q = p^j$. The substitution $x \mapsto x^Q$ permutes $\mathbb{F}_q$ and satisfies $(1-x)^Q = 1 - x^Q$. Therefore

$$J(\chi^a, \chi^b) = J(\chi^{Qa}, \chi^{Qb}) = J(\chi^{-a}, \chi^{-b}) = \overline{J(\chi^a, \chi^b)}.$$

The sum is real. To recall its absolute value, take any nontrivial additive character e of $\mathbb{F}_q$ and put $G(\alpha) = \sum_x \alpha(x)e(x)$. For a nontrivial multiplicative character α, substituting $x = uy$ and using additive and multiplicative orthogonality gives

$$|G(\alpha)|^2 = \sum_{u \neq 0} \alpha(u) \sum_{y \neq 0} e((u-1)y) = q.$$

If $\alpha, \beta, \alpha\beta$ are nontrivial, sorting the terms in $G(\alpha)G(\beta)$ according to $x + y$ gives $G(\alpha)G(\beta) = J(\alpha, \beta)G(\alpha\beta)$; the terms with $x + y = 0$ sum to zero. Thus $|J(\alpha, \beta)| = \sqrt{q}$. Apply this with $\alpha = \chi^a$, $\beta = \chi^b$. The sign can also be determined. Every power of χ is trivial on $\mathbb{F}_Q^\times$, since $m \mid Q + 1$. Use an additive character $e = \psi \circ \mathrm{Tr}_{q/Q}$ with ψ nontrivial, and take $v \neq 0$ on the trace-zero line. To compute $G(\alpha)$, sum over the $Q + 1$ one-dimensional $\mathbb{F}_Q$-spaces in $\mathbb{F}_q$. At their nonzero points the additive character sums to $Q - 1$ on the trace-zero line and to -1 on every other line. For $\alpha = \chi^a, \chi^b, \chi^{a+b}$, the sum of the multiplicative character over the quotient $\mathbb{F}_q^\times/\mathbb{F}_Q^\times$ is zero. Thus $G(\alpha) = Q\alpha(v)$. The Gauss-sum quotient above is exactly Q, also in characteristic two. Its principal ideal in $\mathbb{Z}[\zeta_m]$ is $(p)^j$. □

In the remaining Kummer case, suppose $s = \mathrm{ord}_r(p)$ is even and $a = v_r(p^s - 1)$. Then $r \nmid p^{s/2} - 1$, whereas $r \mid p^{s/2} + 1$; hence $r^a \mid p^{s/2} + 1$. Proposition 30.12 applies with $q = p^s$, $j = s/2$, and $m = r^a$. Thus extending the base field in Pila's argument does not preserve its Jacobi-sum separation mechanism in this case. This observation concerns those Jacobi sums; it does not exclude other curves, twists, or uses of point counting.

31 An auxiliary cyclotomic conductor for Jacobi sums

An enlargement of the cyclotomic conductor can remove the semiprimitive phenomenon and can make Jacobi ideals distinguish an entire r-power subfield. The statements below establish that algebraic distinction. They do not assert a uniform polynomial-time method for computing the required Jacobi generator.

Proposition 31.1. *Let p, r, d, a, q satisfy Lemma 6.1, let $2 \leq b \leq a$, and put*

$$M = r^b q, \qquad K_M = \mathbb{Q}(\zeta_M), \qquad G = (\mathbb{Z}/M\mathbb{Z})^\times, \qquad H = \langle p \rangle.$$

Let R_b be the unique subfield of $\mathbb{Q}(\zeta_{r^b})$ of degree r^{b-1} over $\mathbb{Q}$. Fix a prime $\mathfrak{P}$ of K_M above p, and let $\chi_\mathfrak{P}$ be the order-M power-residue character of its residue field. The Jacobi sum

$$J = -\sum_{z \in \mathcal{O}_{K_M}/\mathfrak{P}} \chi_\mathfrak{P}(z)\chi_\mathfrak{P}(1-z) \quad \textit{satisfies} \quad R_b \subseteq \mathbb{Q}(J).$$

Moreover the tuple of valuations of all Galois conjugates of J distinguishes any two primes above p whose restrictions to R_b are different. More precisely, $\mathbb{Q}(J)$ contains a cyclic CM field of degree $2^{v_2(r-1)}r^{b-1}$ containing R_b. The valuation tuple distinguishes any two primes with different restrictions to this larger field.

Proof. Write $r-1=2^u m$ with m odd. Choose a character τ modulo r of order 2^u; it is odd, and $\tau(p)$ has order 2^t. Choose a character η modulo q of 2-power order satisfying $\eta(p)=\tau(p)^{-1}$. This is possible because p is a quadratic nonresidue modulo q, so evaluation at p maps the 2-primary character group modulo q onto its full 2-power roots of unity. The character η has order 2^t, is nonprincipal, and is even: the 2-part of $\operatorname{ord}_q p$ is strictly larger than 2^t.

Let ν be a primitive character modulo r^b of order r^{b-1}, trivial on the Teichmüller subgroup. Because p has order d modulo r^b, we have $\nu(p)=1$. Thus

$$\psi=\tau\nu\eta$$

is an odd primitive character of conductor exactly M, and $\psi(p)=1$. The factors $\tau\eta$ and ν have coprime orders. In particular,

$$\ker\psi\subseteq\ker\nu, \tag{74}$$

where ν is inflated to G.

Use the convention

$$\theta_M(c)=\sum_{s\in G}\left\langle -\frac{cs}{M}\right\rangle\sigma_s^{-1},\qquad \alpha=2\theta_M(1)-\theta_M(2).$$

The composite-conductor Stickelberger formula gives

$$(J)=\mathfrak{P}^{\alpha}.$$

The identity holds for arbitrary residue degree; it is stated in Bernard and Kučera (2024, Proposition 5.1, author p. 17, proof p. 18), with precisely this convention for θ_M and for the sign of J. Here α has coefficients 0 or 1.

Evaluating at ψ, a substitution $v=cs \bmod M$ gives

$$\psi(\theta_M(c))=\psi(c)L(0,\overline{\psi})\quad((c,M)=1).$$

Indeed the constant term in $\langle -s/M\rangle=1-s/M$ sums to zero for a nonprincipal character. Consequently

$$\psi(\alpha)=(2-\psi(2))L(0,\overline{\psi})\neq 0.$$

The first factor is nonzero because $|\psi(2)|=1$; the second is nonzero for an odd primitive Dirichlet character, by its functional equation and nonvanishing at 1. Equivalently, the nonvanishing follows from the classical relative class-number product recorded in Bernard and Kučera (2024, equation (31), author p. 18).

The primes above p form the G-set G/H. Let $\overline{\alpha}\in\mathbb{Z}[G/H]$ be the valuation vector of (J). Since $\psi|_H=1$, the evaluation just computed descends to this permutation module and remains nonzero. If $g\in G$ stabilizes (J), then

$$\psi(\overline{\alpha})=\psi(g\overline{\alpha})=\psi(g)\psi(\overline{\alpha}),$$

so $g\in\ker\psi\subseteq\ker\nu$. Every automorphism fixing J fixes its ideal. Hence

$$\operatorname{Gal}(K_M/\mathbb{Q}(J))\subseteq\ker\nu.$$

In fact the preceding stabilizer calculation gives $K_M^{\ker\psi}\subseteq\mathbb{Q}(J)$. The character $\tau\eta$ has order 2^u by restriction to the r-part of the conductor, and ν has the coprime order r^{b-1}. Thus the former field is cyclic of degree $2^u r^{b-1}$. It is a CM field because ψ is odd, and p splits completely in it because $\psi(p)=1$. It contains $R_b=K_M^{\ker\nu}$. The valuation tuple consequently separates restrictions even to $K_M^{\ker\psi}$.

For the final assertion, if two primes have different restrictions to R_b, the group element taking one to the other has $\nu(g)\neq 1$. It cannot stabilize the entire valuation vector. Translating a coordinate where the valuations differ amounts to replacing J by a Galois conjugate. That conjugate distinguishes the given pair. □

An exact small example. For $p = 17$, $r = 3$, $q = 5$, and $M = 45$, the decomposition subgroup is $H = \{1, 17, 19, 8\}$. The kernel of reduction from conductor 45 to 15 is $\{1, 16, 31\}$. Three cosets in the orbit of H under $\{1, 16, 31\}$ are

$$\{1, 8, 17, 19\}, \quad \{2, 16, 34, 38\}, \quad \{4, 23, 31, 32\}.$$

Label primes by inverse cosets, consistently with the coefficient of σ_s^{-1} in α. That coefficient is 1 exactly when $s < M/2$, their valuation sums are $4, 2, 1$. Adjoining their negatives gives all six cosets and valuations $4, 2, 1, 0, 2, 3$, respectively. Only H has value 4, so the ideal stabilizer is contained in H. Conversely $\sigma_{17}J = J$ by the substitution $z \mapsto z^{17}$ in the finite sum. Thus $\mathbb{Q}(J) = K_{45}^H$, the sextic compositum of $\mathbb{Q}(\zeta_9 + \zeta_9^{-1})$ and $\mathbb{Q}(\sqrt{-15})$. Indeed both fields are fixed by 17, and their coprime degrees give six. This determines the field generated by J from its ideal valuations and Frobenius symmetry.

The conductor and computation requirements. Taking $b = 2$ gives $M \ll r^2 d^2 \log^2(2pd)$, polynomial in r and $\log p$. The proposition then distinguishes the first degree-r cyclotomic layer. It does not supply a maximal r-primary unit when $a > 2$. Taking $b = a$ distinguishes all the required r-power layers, but its numerical conductor contains r^a, which need not be polynomial in r and $\log p$. There is also a degree requirement intrinsic to this particular Jacobi generator:

$$[\mathbb{Q}(J) : \mathbb{Q}] \geq 2^{v_2(r-1)} r^{b-1}.$$

If J occurs as a Frobenius eigenvalue of an abelian variety of dimension g, then $2g \geq 2^{v_2(r-1)} r^{b-1}$.

Even at $b = 2$, the point-counting theorem of Pila (1990, Theorem A, author p. 1) has an exponent depending on the defining equations and dimension. It does not give a uniform polynomial bound when the auxiliary conductor q and the Fermat genus grow with $\log p$. Knowing the ideal exponent α, or its nonzero Fourier coefficient, also does not give a polynomial-time computation of a generator of $\mathfrak{P}^\alpha$. The positive distinction above therefore leaves a concrete computational requirement: construct a suitably separating Jacobi generator uniformly within the allowed input budget, and then give the required maximal-unit extraction. Neither requirement is supplied by the ideal calculation alone.

More explicitly, even when $a = b = 2$, the output needed for the finite-field problem is an explicitly represented element in the given target field whose r-primary part has order r^2. The proposition does not produce that element. A route through its Jacobi sum must still compute a representation of the sum, or a separating Frobenius polynomial with sufficiently explicit CM endomorphisms, and justify the reduction and factor-selection step in the presented finite field. The valuation assertion concerns the entire tuple of conjugates; it does not assert that the single valuation of one chosen sum distinguishes every prime.

Corollary 31.2. *For $r = 3$ and $b = 2$, the field cut out by ψ above is*

$$L = \mathbb{Q}(\zeta_9 + \zeta_9^{-1})\mathbb{Q}(\sqrt{-3q}).$$

It is cyclic of degree 6, p splits completely in it, and its absolute discriminant is $3^9 q^3$.

Proof. Indeed, the order-3 part is the real cubic field of conductor 9, and the order-2 part is the product of the quadratic characters of conductors 3 and q. Both characters take the value 1 at p. The real cubic field has discriminant 81, while $\mathbb{Q}(\sqrt{-3q})$ has discriminant $-3q$, since $q \equiv 1 \pmod 4$. The tensor product of integral bases gives an order in L, so its discriminant yields

$$|D_L| \leq 81^2 (3q)^3 = 3^{11} q^3.$$

The five nonprincipal characters have conductors $9, 9, 3q, 9q, 9q$. The conductor-discriminant identity therefore gives the exact value $3^9 q^3$. At 3, the wild cubic part cannot cancel the tame quadratic part; at q, the quadratic part is nontrivial. This proves the conductor list. □

The fixed-valuation construction of Theorem 1.3, with $r = 3$ and $b = 2$, realizes this dimension-three approach under A. The conductor calculation above alone supplies neither the CM model nor its finite-field descent, and the resulting order-nine unit is maximal only when $v_3(p+1) = 2$.

The distinction is material in the available CM factorization literature. Boyer (2014, Chapter III, printed p. 50/PDF p. 62) fixes the conductor, measures complexity only in $\log p$, and allows arbitrary finite preprocessing independent of p. His Proposition 2.6, printed p. 67/PDF p. 79, proves separation for the displayed CM curves with field $\mathbb{Q}(\zeta_l + \zeta_l^{-1}, i)$, for prime $l > 5$, $p \equiv 1 \pmod 4$, and $p \equiv -1 \pmod l$. The broader even-order claims are Conjecture 2.7 (printed p. 68/PDF p. 80) and Conjecture 2.11 (printed p. 72/PDF p. 84). Those results do not provide a uniform varying-q model for the sextic field just displayed.

31.1 Polarized ideals, local trace sums, and residue maps

The generator-recovery step for an explicit polarized ideal has an unconditional polynomial-time solution. More precisely, Lenstra and Silverberg (2019, Theorem 1.3, author manuscript p. 2) give a deterministic algorithm which, on input a CM-order A, an explicit fractional A-ideal I, and a totally positive $w \in A \otimes \mathbb{Q}$ satisfying $I\overline{I} = wA$, decides whether $I = \alpha A$ and $\alpha\overline{\alpha} = w$ for some α, and constructs such an α when one exists. Its bit cost is polynomial in the complete explicit input length, including the rank and the ideal basis. This theorem does not construct an ideal basis from a formal expression in unnamed primes above an input prime.

Tal (2026, Theorem 2.7 and §§2.6, 3.1–3.3, pp. 25–31) apply this theorem to polarized Jacobi ideals. The character and the compatible prime above p are inputs to that reconstruction; in the higher-index algorithm, a representation of the finite field is also supplied. The subsequent orientation step chooses among symbolic roots of unity in the cyclotomic number field, using proper-conductor data and integral coefficient conditions (Tal, 2026, Lemma 2.8, pp. 27–28; Appendix G, pp. 63–65). This follows the phase-selection framework of van Wamelen (2002, Theorem 9, pp. 9–10); his implementation first obtains the relevant prime from a factor modulo p (van Wamelen, 2002, §4, p. 4; §7, pp. 18–19). The ideal must therefore be made explicit before these generator algorithms apply. This is a restriction on this transfer of the algorithms, and does not affect their stated factorial or Jacobi-sum computations. In particular, Tal's deterministic preprocessing bound in the split case allows $p^{1/4+o(1)} + q^{O(1)}$ operations, and the higher-index bound has constants depending on the fixed index (Tal, 2026, Theorem 1.3, pp. 6–7; §1.3, pp. 12–14; §4.2, p. 33). These are not a uniform polynomial bound in $\log p$ and $\log q$.

The prime, factor, and root interfaces. For $m \geq 2$ with $p \nmid m$, put $f = \mathrm{ord}_m(p)$. The following three explicit data are interconvertible in deterministic time polynomial in m, $\log p$, and their representation lengths: an irreducible factor g of Φ_m modulo p; a basis for a prime $\mathfrak{P}$ of $\mathbb{Z}[\zeta_m]$ above p; and a presented field $\mathbb{F}_{p^f}$ with an element η of exact order m. Indeed, a factor gives $\mathfrak{P} = (p, g(\zeta_m))$ and the field $\mathbb{F}_p[X]/(g)$. Its distinguished element has order m, since the factors Φ_d for $d \mid m$ have disjoint roots modulo p, as follows from the separability of $X^m - 1$. A prime basis gives the residue field by integer and finite-field linear algebra. The image of ζ_m generates that field; its successive powers give its minimal polynomial. Conversely, the Frobenius orbit of η has length exactly f, so

$$\prod_{i=0}^{f-1} (X - \eta^{p^i})$$

is the required irreducible factor over $\mathbb{F}_p$. These conversions do not produce any of the three data from the integers p, m alone.

Proposition 31.3 (Local units of the trace-one sums)**.** *Let $p > 3$ be prime with $p \equiv -1 \pmod 3$, let $a = v_3(p+1)$, and suppose $q \geq 4$ divides $p+1$ with $(p+1)/q$ odd. Put $M = 2q$ and $k = \mathbb{F}_{p^2}$. Suppose a multiplicative character χ_M of exact order M and compatible reductions of its cyclotomic values into k are given. For each $m \mid M$, write*

$$\chi_m = \chi_M^{M/m}, \qquad L_m = \mathbb{Q}(\zeta_m), \qquad E_m = \sum_{\mathrm{Tr}_{k/\mathbb{F}_p}(x)=1} \chi_m(x).$$

The characters are extended by zero at zero, including the trivial character. Let $\mathfrak{P}_m$ be the prime specified by the given reduction ρ_m, and let $e_m = v_{\mathfrak{P}_m}(E_m)$, with $v_{\mathfrak{P}_m}(p) = 1$. Then every E_m is nonzero and

$$c_m = \rho_m(p^{-e_m} E_m) \in \mathbb{F}_p^*.$$

Here reduction of the displayed local unit is taken in the local ring at $\mathfrak{P}_m$.

Put $b = v_3(M)$. If $b < a$, every normalized local reduction of $\xi_m E_m$, for any root of unity $\xi_m \in L_m$, is a cube in k^*. Every product, quotient, or integer power formed from these nonzero reductions is again a cube. If $b = a$, the already supplied element $\eta_M = \rho_M(\zeta_M)$ gives an element of exact order 3^a by

$$\eta_M^{M/3^a}.$$

The case $b > a$ cannot occur. No restriction is imposed on the other prime factors of q.

Proof. First, $M \mid p^2 - 1$, so $p \nmid m$ for every $m \mid M$. Thus L_m is unramified at p. The trace-one fiber has p elements and contains no zero, giving $E_1 = p$.

Suppose $m > 1$. If $m \mid p+1$, the restriction of χ_m to $\mathbb{F}_p^*$ is trivial. Fix a nonzero z of trace zero. The sum of χ_m on the nonzero trace-zero elements is $(p-1)\chi_m(z)$. The sum on each trace-t fiber, for $t \in \mathbb{F}_p^*$, is E_m, by scaling the trace-one fiber by t. The sum of a nontrivial character over k is zero, and hence $E_m = -\chi_m(z) \neq 0$.

If $m \nmid p+1$, its divisibility into $2(p+1)$ shows that $m/\gcd(m, p+1) = 2$. Consequently $\chi_m|_{\mathbb{F}_p^*}$ is the quadratic character λ. Its sum on the trace-zero fiber is zero, and its sum on the trace-t fiber is $\lambda(t)E_m$ for $t \neq 0$. For a nontrivial additive character ψ of $\mathbb{F}_p$, this proves the Gauss-sum identity

$$G_k(\chi_m, \psi \circ \mathrm{Tr}) = E_m G_{\mathbb{F}_p}(\lambda, \psi).$$

For any nontrivial multiplicative character θ of a finite field of size Q, and any nontrivial additive character Ψ, substituting $x = uy$ in the squared absolute value of the Gauss sum gives

$$|G(\theta, \Psi)|^2 = \sum_{u \neq 0} \theta(u) \sum_{y \neq 0} \Psi((u-1)y) = Q.$$

The two Gauss sums therefore have absolute values p and $\sqrt{p}$, respectively, proving $|E_m|^2 = p$ and in particular $E_m \neq 0$. This also yields the algebraic identity $E_m \overline{E}_m = p$ in this case.

Let $\sigma_p(\zeta_m) = \zeta_m^p$. Because $\chi_m(x)^p = \chi_m(x^p)$ and $x \mapsto x^p$ permutes the trace-one fiber, $\sigma_p(E_m) = E_m$. On the residue field, $\rho_m(\sigma_p(y)) = \rho_m(y)^p$ for integral y. This identity shows that σ_p stabilizes $\mathfrak{P}_m$ and extends to its local ring. Since $p^{-e_m} E_m$ is a local unit fixed by σ_p, its residue satisfies $c_m^p = c_m$. It is nonzero, and therefore belongs to $\mathbb{F}_p^*$.

For completeness, a root of unity of order n in L_m has $n \mid \mathrm{lcm}(2, m)$. Indeed, with $l = \mathrm{lcm}(m, n)$, the two given roots generate $\mathbb{Q}(\zeta_l) = \mathbb{Q}(\zeta_m)$, so $\varphi(l) = \varphi(m)$. As $m \mid l$, the totient formula forces $l = m$, or $l = 2m$ with m odd. Thus every such root of unity has order prime to p, is a local unit, and its reduction has order dividing $\mathrm{lcm}(2, m)$.

Now $3 \nmid p - 1$ and $v_3(p^2 - 1) = a$. If $b < a$, then

$$\mathrm{lcm}(2, m) \mid (p^2 - 1)/3 \qquad (m \mid M).$$

In the cyclic group k^* every element of such an order is a cube. Every element of $\mathbb{F}_p^*$ is also a cube, since cubing is an automorphism of that subgroup. Hence the normalized reduction $\rho_m(\xi_m)c_m$ is a cube. Closure of the cube subgroup proves the assertion about products, quotients, and integer powers. If $b = a$, reduction preserves the exact order M of ζ_M, since $p \nmid M$ and the cyclotomic factors of $X^M - 1$ have disjoint roots modulo p. The order formula in a cyclic group then gives the asserted exact order of $\eta_M^{M/3^a}$. Finally $M \mid p^2 - 1$ implies $b \le a$. □

The preceding proposition concerns direct multiplicative use of the specified local units and their root-of-unity phases. It makes no assertion about arbitrary additions, global coefficient reconstruction, or the ability of global ideal valuations to distinguish residue maps. In fact the latter distinction becomes effective when the ideals themselves are available:

Proposition 31.4 (Explicit ideal profiles give a residue map)**.** *Let K have degree N, with explicitly given ring of integers O, and let p be unramified in K. Let $B \subseteq O$ be a unital order in a subfield of K, given by an integer basis and its embedding in O. Suppose every prime $\mathfrak{q}$ of B above p has residue field $B/\mathfrak{q} = \mathbb{F}_p$. Let $I_1, \dots, I_s$ be explicit integral O-ideals, given by integer bases and supported only above p, with a known bound*

$$0 \le v_{\mathfrak{P}}(I_j) \le h \quad (\mathfrak{P} \mid p,\ 1 \le j \le s).$$

Suppose their profiles distinguish different contractions to B:

$$\mathfrak{P} \cap B \ne \mathfrak{P}' \cap B \implies (v_{\mathfrak{P}}(I_j))_{j=1}^s \ne (v_{\mathfrak{P}'}(I_j))_{j=1}^s.$$

Then one can construct a unital $\mathbb{F}_p$-algebra homomorphism $B/pB \longrightarrow \mathbb{F}_p$ unconditionally and deterministically in time polynomial in $N, s, h, \log p$ and the total bit length of the input bases, multiplication data, and embedding. The exponent is absolute. No prime above p or finite-field factorization is an additional input. The order B need not be maximal, and p may divide $[O_F : B]$, where F is its fraction field and O_F is the ring of integers of F.

Proof. Put $A = O/pO$. Since p is unramified, A is a product of finite fields, although no such product decomposition is supplied or computed initially. For $1 \le j \le s$ and $1 \le t \le h$, compute the integral ideal

$$L_{j,t} = (p^{-(t-1)} I_j \cap O) + pO$$

by integer lattice intersection and sum, and let $J_{j,t}$ be its image in A. At a prime $\mathfrak{P} \mid p$, writing $v = v_{\mathfrak{P}}(I_j)$, its valuation is

$$v_{\mathfrak{P}}(L_{j,t}) = \min\{\max\{v - t + 1, 0\}, 1\}.$$

Thus the corresponding factor of $J_{j,t}$ is zero when $v \ge t$, and is the whole residue field when $v < t$.

The identity element $e_{j,t}$ of this direct-factor ideal is computable by linear algebra over $\mathbb{F}_p$. Represent an unknown $e \in J_{j,t}$ in a basis of that ideal, and impose $ev = v$ for every basis vector v of $J_{j,t}$. These are linear equations. In a product of fields they have the unique solution which is one on the supported factors and zero elsewhere. This also covers $J_{j,t} = 0$, whose identity in this sense is zero.

Starting with the single idempotent 1_A, refine a list of nonzero orthogonal idempotents by replacing each u with the nonzero members of $ue_{j,t}, u(1 - e_{j,t})$. Each refinement preserves their sum and their orthogonality. At most N nonzero members can occur, since their nonzero ideals in the N-dimensional algebra A form a direct sum. Once all pairs (j, t) have been processed, any two prime factors in the support of one member u have the same full valuation profile. The hypothesis implies that they contract to a single prime $\mathfrak{q}$ of B.

Choose any member u of this final list. At each field factor in its support the image of an element of B factors through $B/\mathfrak{q} = \mathbb{F}_p$. The embedding of that prime field into each factor is canonical. Consequently, for each $b \in B$ there is a unique $\lambda_b \in \mathbb{F}_p$ with

$$u\,\iota(b) = \lambda_b u \quad \text{in } A,$$

where $\iota : B \to A$ is reduction of the given embedding. To compute λ_b, choose one nonzero coordinate of u in the given $\mathbb{F}_p$-basis of A, and divide the corresponding coordinate of $u\iota(b)$ by it. Perform this operation for each vector of the given integer basis of B and extend $\mathbb{F}_p$-linearly modulo pB. The resulting map is additive, sends 1 to 1, and is multiplicative because $u^2 = u$:

$$u\iota(bb') = (u\iota(b))(u\iota(b')) = \lambda_b \lambda_{b'} u.$$

This constructs the claimed homomorphism. The map $B/pB \to A$ need not be injective. Its image is sufficient for the argument; in particular every nilpotent maps to zero in the reduced algebra A. Empty refinement lists when $s = 0$ or $h = 0$ cause no problem: the profile hypothesis then says that all contractions coincide.

Finally, all rational lattice denominators used above divide p^{h-1} when $h \geq 1$, with bit length $O(h \log p)$. Integer lattice intersections and sums, using Hermite normal forms and linear systems, have deterministic polynomial bit cost in their dimensions and their input bit lengths. Reduction modulo p, the sh linear systems for the idempotents, and the refinements of at most N members use a polynomial number of operations on $O(\log p)$-bit field elements. Input structure constants and bases control the remaining coefficient lengths. These bounds give an absolute polynomial in the parameters stated, without enumerating p elements or factoring a polynomial over $\mathbb{F}_p$. □

Thus explicit separating Jacobi ideals would already suffice for the residue-map step, even without recovering a global generator. The conclusion is conditional on their integer bases being available. A formal Stickelberger exponent vector, expressed relative to an unspecified prime above p, does not provide those bases. Neither proposition supplies that missing input for a growing conductor, and neither invokes a fixed-field residue-map theorem with a varying field.

The explicit ideal bases in the preceding residue-map interface remain substantive input. The following elementary observations describe some restrictions on constructing them by CM reflex norms and ideal arithmetic. The degree assertion has an explicit field-containment premise.

Proposition 31.5 (Restrictions on reflex fields and polarized ideal profiles)**.** *The following statements hold unconditionally.*

1. *Let r be a prime, $b \geq 1$, and let $F/\mathbb{Q}$ be cyclic of degree r^b. Suppose F is contained in the compositum of the Galois closures of number fields $E_1, \ldots, E_s$, where $n_i = [E_i : \mathbb{Q}]$ and $s \geq 1$. Then*

$$r^b \leq \max_i n_i.$$

In particular, if F is contained in the compositum of reflex fields of CM fields of dimensions g_i, then $r^b \leq \max_i 2g_i$. Here dimension means that the CM field has degree $2g_i$.

2. *Let K be a CM field with maximal order $\mathcal{O}$, and let p be unramified in K. Suppose every prime $\mathfrak{P}$ above p is fixed by complex conjugation. If a nonzero integral ideal I is supported above p and*

$$I\overline{I} = p^k\mathcal{O} \qquad (k \in \mathbb{Z}_{\geq 0}),$$

then k is even and $I = p^{k/2}\mathcal{O}$. More generally, if $I\overline{I} = w\mathcal{O}$ for a totally positive $w \in K^+$, then

$$2v_{\mathfrak{P}}(I) = v_{\mathfrak{P}}(w) \qquad (\mathfrak{P} \mid p).$$

Thus a family of such ideals can have separating valuation profiles only if the corresponding principal ideals of the polarizations already have separating profiles.

3. *Let r be an odd prime, $b \geq 1$, and let $K/\mathbb{Q}$ be an abelian CM extension with*

$$G = \mathrm{Gal}(K/\mathbb{Q}) = U \times V, \qquad U \simeq C_{r^b}, \qquad r \nmid |V|.$$

Put $F = K^V$ and let $U_1 \subset U$ be the subgroup of order r. For each of finitely many CM types Φ_j on K, write

$$H_j = \{h \in G : h\Phi_j = \Phi_j\}, \qquad E_j = K^{H_j}$$

for its stabilizer and reflex field. Suppose $F \not\subset E_j$ for every j. For any nonzero fractional ideal A_j of E_j, let $I_j = N_{\Phi_j}(A_j)$ be its CM reflex norm ideal in K. Then every I_j is fixed by U_1.

This remains true for any ideals obtained from the I_j and rational principal ideals by products, sums, intersections, fractional ideal colons, inverses, intersection with $\mathcal{O}_K$, multiplication or division by integer powers of p, and G-conjugation. It also remains true under the filters

$$I \longmapsto (p^{-t}I \cap \mathcal{O}_K) + p\mathcal{O}_K \qquad (t \in \mathbb{Z}_{\geq 0}).$$

If p is unramified in K and splits completely in F, no family of p-supported integral ideals obtained in this manner has valuation profiles separating all primes of K above different primes of F.

Still assuming that p splits completely in F, if instead a reflex field E_j contains F and a prime ideal of E_j above p is supplied explicitly, its contraction to $\mathcal{O}_F$ already specifies a residue map $\mathcal{O}_F/p\mathcal{O}_F \to \mathbb{F}_p$.

Proof. For the first assertion, let L_i be the Galois closure of E_i and $L = L_1 \cdots L_s$. The action of $\mathrm{Gal}(L_i/\mathbb{Q})$ on the n_i embeddings of E_i is faithful: an automorphism fixing all their images fixes the field they generate, namely L_i. Thus

$$\mathrm{Gal}(L/\mathbb{Q}) \hookrightarrow \prod_i \mathrm{Gal}(L_i/\mathbb{Q}) \hookrightarrow \prod_i S_{n_i}.$$

The exponent of S_n is $\mathrm{lcm}(1, \ldots, n)$, since permutations decompose into disjoint cycles and every cycle length at most n occurs. Its r-part is consequently $r^{\lfloor \log_r n \rfloor}$. A subgroup of the displayed product has exponent dividing the exponent of that product, whose r-part is $r^{\lfloor \log_r(\max_i n_i) \rfloor}$. Restriction to the assumed subfield F is a surjection onto $\mathrm{Gal}(F/\mathbb{Q}) \simeq C_{r^b}$. The exponent of a quotient divides the exponent of its source. It follows that $b \leq \lfloor \log_r(\max_i n_i) \rfloor$, as required. The reflex field of a CM field is a subfield of its Galois closure: its defining type stabilizer is a subgroup of that Galois group (Milne, 2020, Proposition 1.16 and Example 1.19, pp. 13–14). This proves the stated reflex-field consequence without assuming that the degree of a reflex field is at most the degree of its CM field.

For the second assertion, write $I = \prod_{\mathfrak{P}|p} \mathfrak{P}^{e_{\mathfrak{P}}}$. Complex conjugation fixes each factor, and p is unramified, so comparison of valuations in $I\bar{I} = p^k\mathcal{O}$ gives $2e_{\mathfrak{P}} = k$ for every $\mathfrak{P} \mid p$. Hence k is even and all the exponents are $k/2$. The same comparison with $w\mathcal{O}$ proves the more general formula. Equalities of valuation vectors for a family of ideals are therefore exactly the equalities of the corresponding vectors for their polarizations, with all entries multiplied by two. For example, in $K = \mathbb{Q}(\zeta_{r^a})$ and $p \equiv -1 \pmod{r^a}$, the Frobenius at p is complex conjugation; its decomposition group fixes each prime above p. The scalar-polarization conclusion consequently applies at this direct cyclotomic conductor.

For the third assertion, the description $E_j = K^{H_j}$ follows from the same reflex-field definition. The hypothesis $F \not\subset E_j$ is equivalent to $H_j \not\subset V$. A subgroup of the finite abelian group $U \times V$

is the product of its primary parts. Since $r \nmid |V|$, the failure of $H_j \subset V$ means that H_j has a nontrivial r-primary part in the cyclic group U. Every nontrivial subgroup of U contains its unique subgroup U_1 of order r. Thus

$$U_1 \subset H_j \qquad \text{for every } j.$$

Put $A'_j = A_j \mathcal{O}_K$. In the ideal convention for the CM reflex norm, its defining identity in this abelian case is

$$I_j^{|H_j|} = \prod_{\phi \in \Phi_j} \phi^{-1}(A'_j) \quad \text{in the group of fractional ideals of } K;$$

see Milne (2020, Corollary 1.27, p. 17), with its field k equal to K. Every $h \in H_j$ fixes A'_j. Since G is abelian, it commutes with each ϕ and therefore fixes every factor on the right. Hence $(hI_j)^{|H_j|} = I_j^{|H_j|}$. Unique prime-ideal factorization makes the fractional ideal group torsion-free, so $hI_j = I_j$. In particular U_1 fixes I_j.

The listed ideal operations preserve this invariance directly. For instance, $h(I : J) = (hI : hJ)$ and $h(I \cap J) = hI \cap hJ$; the other formulas follow from the same definitions. Rational principal ideals and $\mathcal{O}_K$ are fixed. Because G is abelian, conjugating a U_1-fixed ideal by any element of G again gives a U_1-fixed ideal. These observations also prove invariance of the displayed filters.

Now let $\mathfrak{P}$ be any prime above p, and choose a nonidentity $u \in U_1$. Its contraction $\mathfrak{q} = \mathfrak{P} \cap \mathcal{O}_F$ satisfies $u\mathfrak{q} \neq \mathfrak{q}$: the restriction of u to F is nonidentity, while a completely split prime in the Galois extension $F/\mathbb{Q}$ has trivial decomposition group. Thus $\mathfrak{P}$ and $u\mathfrak{P}$ have different contractions to F. Nevertheless every U_1-fixed ideal I has

$$v_{u\mathfrak{P}}(I) = v_{\mathfrak{P}}(u^{-1}I) = v_{\mathfrak{P}}(I).$$

No tuple of these profiles separates that pair.

Finally, if $F \subset E_j$ and an explicitly given prime $\mathfrak{Q} \subset \mathcal{O}_{E_j}$ above p is supplied, then $\mathfrak{q} = \mathfrak{Q} \cap \mathcal{O}_F$ is a prime of F above p. Complete splitting gives $\mathcal{O}_F/\mathfrak{q} \simeq \mathbb{F}_p$, canonically as an $\mathbb{F}_p$-algebra. Composing reduction with this quotient gives the asserted map. With explicit integral bases and the field embedding, contraction is an integer-lattice intersection; the values of the map on a basis are then obtained by linear algebra modulo p. These operations have polynomial bit cost in the supplied basis data, field degrees, and $\log p$. □

Remark 31.6. A direct Frobenius ideal also detects residue degree one without selecting among equal-degree components. Let p be unramified in a number field K, with integral basis $b_1, \dots, b_N$, and put

$$J = (p, b_1^p - b_1, \dots, b_N^p - b_N) \subset \mathcal{O}_K.$$

At a prime $\mathfrak{P} \mid p$, all the differences vanish exactly when its residue field is $\mathbb{F}_p$: the basis images span the residue field over $\mathbb{F}_p$, and the fixed elements of its p-power Frobenius are precisely $\mathbb{F}_p$. If some difference is nonzero, it generates the unit ideal in that residue field. Thus, when $K/\mathbb{Q}$ is Galois, all residue degrees agree, and $J = p\mathcal{O}_K$ if that degree is one, while $J = \mathcal{O}_K$ otherwise. This proves that J cannot separate the primes in this case. Its basis is inexpensive to compute: reduce binary powers modulo p in the integral multiplication table, lift their coefficient vectors, and form their lattice with $p\mathcal{O}_K$. The cost is polynomial in the table's bit length and $\log p$; no large integral power is needed.

The proposition concerns the stated field containments, polarizations, reflex norms, and ideal operations. It gives no obstruction to arbitrary new elements formed by addition, to relative towers, or to general correspondences. In particular it does not replace the missing construction of explicit separating ideal bases from the order and p alone. Lenstra–Silverberg's generator algorithm still takes the ideal itself as input (Lenstra and Silverberg, 2019, Theorem 1.3, author manuscript p. 2).

32 The size of the wild field and the limits of its units

The next bounds concern the characteristic-zero representation and cyclotomic units used above, rather than arbitrary finite-field algorithms.

Proposition 32.1. *Let R_b be as above. If R_b is contained in the compositum of finitely many number fields, each of degree at most D, then $D \geq r^{b-1}$. If all the fields are abelian CM fields of degree at most $2h$, then $h \geq r^{b-1}$.*

Proof. The Galois group of the normal closure of a degree-d number field embeds in the symmetric group S_d. The r-part of its exponent is therefore at most $r^{\lfloor \log_r D \rfloor}$: the exponent of S_D is $\operatorname{lcm}(1, \ldots, D)$. The Galois group of the compositum of the normal closures embeds in their direct product and has the same upper bound on the r-part of its exponent. Since $R_b/\mathbb{Q}$ is cyclic of degree r^{b-1}, containment would give a quotient of that order and exponent. Hence $r^{b-1} \leq r^{\lfloor \log_r D \rfloor} \leq D$.

In the abelian CM case the fields are already Galois. The r-part of the exponent of each Galois group divides its degree. A quotient of exponent r^{b-1} of a subgroup of their product consequently requires some factor's degree to be divisible by r^{b-1}. That degree is even, and thus at least $2r^{b-1}$, proving the second assertion. □

This proposition restricts the characteristic-zero containment mechanism used in the construction. In particular, arbitrarily many bounded-dimensional abelian CM fields cannot replace its growing wild field while retaining that mechanism. It does not rule out finite-field computations using such varieties that avoid characteristic-zero field containment.

The limit of this representation. This is a polynomial-time statement on each fixed stratum, not a uniform polynomial-time bound for variable b. To see the representation growth on prime inputs, choose a prime

$$p \equiv 2 \cdot 3^b - 1 \pmod{3^{b+1}}.$$

It has $v_3(p+1) = b$. Under A, the least-prime-in-a-progression bound gives $p \ll (\varphi(3^{b+1}) \log(3^{b+1}))^2 = 3^{O(b)}$ (Lamzouri et al., 2015, Corollary 1.2). Also $p \geq 2 \cdot 3^b - 1$, so $\log p = \Theta(b)$. The dimension $g = 3^{b-1}$, and especially the explicit 8^g or 16^g coordinates, grow faster than every polynomial in $\log p$. This is a limitation of this written CM representation, not a lower bound for other algorithms.

Proposition 32.2. *Let $2 \leq b \leq a = v_3(p+1)$, with $p > 3$, and set $K_b = \mathbb{Q}(\zeta_{3^b})$. The reduction of every $v \in \mathcal{O}_{K_b}[1/3]^\times$ at a prime over p has three-primary part of order at most 3^b.*

Proof. The unique prime above three is fixed by conjugation, so $v/\overline{v}$ is a global unit. All its complex absolute values equal one; Kronecker's theorem makes it a root of unity. The roots of unity of K_b have order dividing $2 \cdot 3^b$. Since $p \equiv -1 \pmod{3^b}$, conjugation on the residue field is p-Frobenius. On its three-primary multiplicative subgroup Frobenius is inversion. The three-primary part of $v/\overline{v}$ is therefore the square of that of v; squaring preserves its order, which is at most 3^b. □

The proposition restricts the whole cyclotomic three-unit group. It does not restrict arbitrary sums, which may acquire other prime divisors, or varying Frobenius integers supported at p. Also, adjoining a finite extension cannot make the order-nine seed a noncube when $a \geq 3$: $v_3(p^{2m} - 1) = a + v_3(m)$, and norms cannot increase its order. The uniform lifting problem remains the selection of roots in the successive promised split trace cubics.

33 A positive density of primes with only two-power residue degrees for fixed CM torsion data

Proposition 33.1. *Fix distinct primes $q_i \equiv 1 \pmod 4$, $q_i \neq 3$, and put $R = \mathbb{Q}(\zeta_9 + \zeta_9^{-1})$ and $L_i = R(\sqrt{-3q_i})$. For each i, fix algebraic coefficients for a CM theta model, finitely many torsion levels, and finitely many algebraic endomorphism data. There is an effectively computable constant a_0, depending on this finite data, such that, for every $a \geq a_0$, a positive-density set of rational primes p satisfies*

$$v_3(p+1) = a, \qquad (p/q_i) = -1 \quad \textit{for every } i,$$

and all residue fields generated by the fixed coefficients and torsion coordinates have degree a power of two over $\mathbb{F}_p$. In particular, their finite tensor products, subalgebras and finite etale factor quotients have no field component of degree divisible by three.

Proof. Let

$$F_0 = R(\sqrt{-3}, \sqrt{q_1}, \ldots, \sqrt{q_m}).$$

The independent quadratic square classes define an involution c_0 fixing R and negating every displayed square root. It fixes every L_i. Choose a finite Galois extension $E/\mathbb{Q}$ containing F_0 and all fixed data fields. In the preimage of $\langle c_0 \rangle$ in $\mathrm{Gal}(E/\mathbb{Q})$, a Sylow two-subgroup maps onto $\langle c_0 \rangle$. Hence it contains an element s restricting to c_0; its order is a power of two.

The intersection of E with the three-power cyclotomic tower is a fixed finite field contained in $\mathbb{Q}(\zeta_{3^c})$. The integer c can be bounded effectively from $[E : \mathbb{Q}]$. This intersection contains $\mathbb{Q}(\zeta_9)$, since F_0 contains R and $\sqrt{-3}$. The restriction of s to the intersection is complex conjugation: its order is a power of two, and it is nontrivial on ζ_3, while the three-power cyclotomic Galois groups have a unique involution.

For $a \geq c$, the residue class

$$\tau_a = -1 + 3^a \pmod{3^{a+1}}$$

restricts to complex conjugation on that intersection. Thus (s, τ_a) defines an element of $\mathrm{Gal}(E\mathbb{Q}(\zeta_{3^{a+1}})/\mathbb{Q})$. The qualitative Chebotarev theorem supplies a positive-density set of primes with this conjugacy class (Serre, 1981, Section 2.1, Theorem 1, p. 131). The cyclotomic restriction gives $v_3(p+1) = a$. The action on $\sqrt{q_i}$ gives $(q_i/p) = -1$, and quadratic reciprocity gives $(p/q_i) = -1$. The action on each L_i is trivial, so p splits in each CM field.

After omitting finitely many ramified primes and denominators, the residue degree in E is the order of s, a power of two. Residue degrees in subfields divide that number. Composita of finite fields of two-power degrees again have two-power degree, and the same holds for subfields. This proves the assertion for the stated finite-algebra operations. □

This proposition concerns fixed coordinate-algebra witnesses. It does not restrict an auxiliary-prime menu or torsion level growing with p, or adaptive calculations with the varying Frobenius integers. It is not an impossibility theorem for finite-field construction.

34 Further information from a fixed CM Frobenius

We now test whether one fixed CM Frobenius element, its local torsion, or its canonical isogenies provide the missing cubic factor.

Lemma 34.1. *In Proposition 12.1, the original p^e-Frobenius F already belongs to the embedded CM field L. Its degree over $\mathbb{Q}$ is six, and e is even.*

Proof. The support argument there proves $F^6 = \pi \in L$ and $\mathbb{Q}(\pi) = L$. Since F commutes with its power π, it centralizes L; the rank-one L-Tate representation has centralizer L. Thus $F \in L$, and its primitive support gives $\mathbb{Q}(F) = L$. The prime three is totally ramified in L, with residue field $\mathbb{F}_3$. Reducing $F\overline{F} = p^e$ at that prime gives $1 = (-1)^e$, so e is even. □

In particular, the tempting normalized sixth-power radical is already trivial in characteristic zero:

$$\frac{\pi}{p^{3e}} = \left(\frac{F^2}{p^e}\right)^3.$$

The same support argument applies to every positive power of F. Every geometric endomorphism is defined over a finite extension of the residue field and therefore commutes with some F^m. Since $\mathbb{Q}(F^m) = L$, the centralizer argument gives $\mathrm{End}^0_{\overline{\mathbb{F}}_p}(A) = L$. Thus this ordinary reduction does not acquire additional noncommuting endomorphisms after extending its constant field. There is a broader local restriction on radicals of elements of a field splitting at the target prime.

Lemma 34.2. *Let a number field L split completely at a prime $p \equiv -1 \pmod{3^a}$. If $1 \le k \le a$, then adjoining 3^k-th roots of finitely many elements of $L^\times$, together with their conjugates and μ_{3^k}, produces only residue degrees dividing two at primes over p.*

Proof. Each completion of L at p is $\mathbb{Q}_p$. Write $\alpha = p^v u$, where $u \in \mathbb{Z}_p^\times$. The map $x \mapsto x^{3^k}$ is an automorphism on $\mathbb{F}_p^\times$, and Hensel's lemma lifts its unique inverse to $\mathbb{Z}_p^\times$. Thus every local radical lies in

$$\mathbb{Q}_p(p^{1/3^k}, \mu_{3^k}).$$

The cyclotomic part is unramified quadratic, because $\mathrm{ord}_{3^k}(p) = 2$; over it the displayed radical of p is totally ramified. The containing field has residue degree two, proving the claim for all subfields and finite composita in question. □

This excludes a residue-degree witness from these radicals alone. It does not exclude new combinations with ζ_3 before radical extraction, or arbitrary arithmetic searches inside the quadratic residue field.

34.1 Exact orbit counts on a full local CM module

For $L = \mathbb{Q}(\theta, \delta)$, with $\theta^3 - 3\theta + 1 = 0$ and $\delta^2 = -3q$, put $\lambda = \delta/(\theta - 2)$. Direct elimination gives

$$\lambda^6 + 15q\lambda^4 + 18q^2\lambda^2 + 3q^3 = 0.$$

This polynomial is Eisenstein at three. Thus $O = \mathcal{O}_L \otimes \mathbb{Z}_3 = \mathbb{Z}_3[\lambda]$, $v_\lambda(3) = 6$, $\overline{\lambda} = -\lambda$, and $O/3 = \mathbb{F}_3[\lambda]/(\lambda^6)$. The next proposition assumes a Tate lattice free of rank one over this full local order; the original index-three product order does not automatically give this hypothesis.

Proposition 34.3. *Suppose Frobenius acts on $O/(\lambda^h)$ by multiplication by ϵu, where $\epsilon \in \{1, -1\}$ and $u \equiv 1 \pmod{\lambda}$. Put $t_j = v_\lambda(u^{3^j} - 1)$, allowing infinity. For $j \ge 1$, the orbit count at length 3^j when $\epsilon = 1$, or length $2 \cdot 3^j$ when $\epsilon = -1$, is respectively*

$$\frac{3^{\min(t_j,h)} - 3^{\min(t_{j-1},h)}}{3^j}, \qquad \frac{3^{\min(t_j,h)} - 3^{\min(t_{j-1},h)}}{2 \cdot 3^j}.$$

A nonzero such count is prime to three only if $t_0 = 1$ and $j = 1$. Conversely that case has a count prime to three whenever $h > 1$.

Proof. The kernel of multiplication by $u^{3^j}-1$ has $3^{\min(t_j,h)}$ elements. Subtract successive fixed sets and divide by the orbit length. For $\epsilon=-1$, every nonzero orbit has even length, and multiplication by $u^{2\cdot 3^j}-1$ has the same kernel as multiplication by $u^{3^j}-1$. When the numerator is nonzero, the count has three-adic valuation $t_{j-1}-j$. The binomial expansion and $v_\lambda(3)=6$ give

$$t_j \geq \min(3t_{j-1}, t_{j-1}+6) \geq t_{j-1}+2.$$

Consequently $t_{j-1}-j \geq t_0+j-2$, proving the necessity and the asserted converse. □

For $h=6$, depth one gives counts $N_3=8, N_9=78$ when $\epsilon=1$, or $N_6=4, N_{18}=39$ when $\epsilon=-1$. The norm relation modulo three permits depths $1,3,5$, as well as scalar reduction: its first nonzero term must have odd degree in λ. Depth three instead gives $N_3=234$ or $N_6=117$, both divisible by three. Nonsemisimplicity alone therefore does not give the required factor count.

34.2 A stronger quotient test with explicit CM action

Proposition 34.4. *Let a presented finite field k_0, a fixed cyclic CM torsion module $O/(\lambda^h)$, and its integral CM action be explicitly given. Put $U_h=(O/(\lambda^h))^\times$, and let $u_F\in U_h$ be its Frobenius unit. An invariant algebra of degree three, obtained from the primitive torsion point algebra using an index-three subgroup of U_h, is a field for at least one such subgroup if and only if $u_F\notin U_h^3$. For fixed h, all such invariant algebras are computable by linear algebra from the given endomorphism action.*

Proof. Primitive torsion points form a torsor under the finite abelian group U_h. For an index-three subgroup H, the quotient is a three-element torsor for U_h/H, and its Frobenius is translation by u_FH. Its coordinate algebra is a field exactly when that translation is nontrivial. The intersection of all index-three subgroups is U_h^3. The invariant coordinate algebra is the simultaneous kernel of $g-1$ for the given permutations $g\in H$. No torsion point or residue-field component is selected. □

For example $1+\lambda^5\in O/3$ is a norm-one noncube, whereas its full point-algebra counts above are divisible by three. Thus the quotient test is stronger than the full point count. The endomorphism units used here act on the finite torsion scheme; they need not be automorphisms of the abelian variety.

Over a coefficient product with t components of degree e, identical successful quotient outcomes give t factors of absolute degree $3e$. The existing factor-count transfer applies when $3\nmid t$. When $3\mid t$ this count gives no procedure for selecting a component. No uniform bound is currently proved for finding an auxiliary q satisfying this test and the component-count condition. Consequently this is a sufficient conditional construction, not a smaller established uniform gate.

34.3 The canonical CM three-isogeny preserves the local order

Under the full local-order hypothesis above, the rank-one O-module has a unique O-stable subgroup of order three. Quotienting by it replaces the local Tate lattice by its λ^{-1}-multiple. This has the same multiplier order and the same Frobenius multiplier. Consequently the canonical CM ideal isogeny preserves the depth and cube-class tests above; it does not provide a descent to a smaller local order. Moreover the unique prime $\mathfrak{p}_3$ of L over three satisfies $\mathfrak{p}_3^6=(3)$. Repeating this ideal isogeny six times returns to a homothetic global lattice, hence to the same unpolarized CM isomorphism class. This observation concerns the O-stable ideal-isogeny chain. It does not exclude choosing a different, non-O-stable subgroup by a separate coordinate algorithm.

35 The original CM lattice removes the depth-one case

The integral lattice of the original period model imposes a stronger restriction than its rational endomorphism field alone.

Proposition 35.1. *Consider the specific sextic CM model with period lattice*

$$\mathcal{O} = \mathbb{Z}[\theta, (1+\delta)/2], \qquad \theta^3 - 3\theta + 1 = 0, \quad \delta^2 = -3q,$$

and a good ordinary reduction at the target prime p. Let $F \in L = \mathbb{Q}(\theta, \delta)$ be its p^e-Frobenius as in Lemma 34.1. With $\lambda = \delta/(\theta - 2)$ and $O = \mathcal{O}_L \otimes \mathbb{Z}_3$, one has

$$\mathcal{O} \otimes \mathbb{Z}_3 = \mathbb{Z}_3 + \lambda^2 O, \qquad F \in (\mathcal{O} \otimes \mathbb{Z}_3)^\times.$$

In particular, for the sign $\epsilon \in \{1, -1\}$ with $F \equiv \epsilon \pmod{\lambda}$, $v_\lambda(\epsilon F - 1) \geq 2$. After an isogeny saturates the local Tate lattice to a rank-one O-lattice, no full torsion-point algebra at any fixed three-power level has a nonzero count prime to three at a factor degree divisible by three, measured over its field of definition. Modulo three one has exactly

$$F = \epsilon(1 + c_3\lambda^3 + c_5\lambda^5), \qquad c_3, c_5 \in \mathbb{F}_3.$$

The invariant-unit quotient test on the saturated three-torsion module succeeds if and only if $c_5 \neq 0$.

Proof. The full local ring is $O = \mathbb{Z}_3[\lambda]$, with $v_\lambda(3) = 6$. The elements $\theta - 2$ and δ have valuations two and three. The product order is therefore contained in $\mathbb{Z}_3 + \lambda^2 O$. The former has index three in the maximal order, by the discriminant calculation for the period model. The latter also has index three, since reduction modulo λ^2 leaves precisely one additional $\mathbb{F}_3$-coordinate. The inclusion is consequently an equality.

The argument uses the integral Tate lattice. The characteristic-zero uniformization with lattice $\mathcal{O}$ gives $T_3(A) = \mathcal{O} \otimes \mathbb{Z}_3$, compatibly with the CM action. To compare with the good reduction, pass to a strictly henselian valuation ring. For every n, the kernel of multiplication by 3^n on the abelian scheme is finite etale: multiplication is finite flat and its differential is invertible because $p \neq 3$. A finite etale scheme over a strictly henselian local ring is constant, so its geometric generic and special fibres are identified compatibly with all CM endomorphisms. Taking inverse limits preserves the specified $\mathcal{O}$-lattice.

Since $F \in L$ and acts integrally on this lattice, it belongs to its multiplier ring $(\mathcal{O}_3 : \mathcal{O}_3) = \mathcal{O}_3$. It is a unit because its norm relation is $F\overline{F} = p^e$ and $p \neq 3$. Its constant residue is ϵ, and a scalar in $\mathbb{Z}_3$ differs from that sign by a multiple of three, of valuation at least six. Thus its normalized depth is at least two.

An isogeny over the residue field does not change the element F of the rational CM algebra. If its field of definition is enlarged, Frobenius becomes a positive power of F, whose normalized depth is still at least two. On a saturated local lattice the orbit-count formula in Proposition 34.3 therefore excludes a count prime to three at every three-divisible orbit length. This assertion is about that full-order module; the original module $\mathcal{O}_3/3\mathcal{O}_3$ is a different integral module and is not identified with $O/3O$.

For the final assertion write $\epsilon F = 1 + c_2\lambda^2 + c_3\lambda^3 + c_4\lambda^4 + c_5\lambda^5$ in $O/3O = \mathbb{F}_3[\lambda]/(\lambda^6)$. The relation $F\overline{F} = 1$ there uses the already proved evenness of e. As $\overline{\lambda} = -\lambda$, its coefficients of λ^2 and then λ^4 give $c_2 = c_4 = 0$. Finally, the cubes in the unit group of this ring are precisely

$$\{\epsilon'(1 + c\lambda^3) : \epsilon' \in \{1, -1\},\ c \in \mathbb{F}_3\},$$

by the characteristic-three binomial identity. The unit is therefore a noncube exactly when $c_5 \neq 0$, and Proposition 34.4 applies. □

This identifies a concrete remaining condition on the varying auxiliary prime: its original CM Frobenius must have nonzero λ^5-coefficient, and the coefficient algebra must also satisfy the required component-count condition. Neither condition has been forced uniformly by rational Dirichlet GRH.

36 A class-number test for varying the auxiliary quadratic field

The quadratic class number gives a computable auxiliary-prime test, but the integral order contributes a separate factor of three.

Proposition 36.1. *Let $q \equiv 1 \pmod 4$ be prime, $q \neq 3$, and put $E = \mathbb{Q}(\sqrt{-3q})$, $R = \mathbb{Q}(\zeta_9+\zeta_9^{-1})$, and $L = ER$. Then*

$$3 \nmid h(L) \quad \Longleftrightarrow \quad 3 \nmid h(E).$$

For the product order $\mathcal{O} = \mathbb{Z}[\theta, (1+\sqrt{-3q})/2]$ used in the period construction, one nevertheless has

$$|\operatorname{Pic}(\mathcal{O})| = 3h(L).$$

Thus making the maximal-order class number prime to three does not remove the three-part introduced by the product order.

Proof. The extension L/E is cyclic cubic, totally ramified at the unique prime of E over three, and unramified at every other finite prime. Both fields are totally imaginary. The units of E are $\{1,-1\}$, and both are norms from L, since the extension degree is odd. The ambiguous class-number formula therefore gives

$$|\operatorname{Cl}(L)^{\operatorname{Gal}(L/E)}| = h(E).$$

For an exact primary proof and formula, see (Yu, 2016, Theorem 1.1, p. 2). A nonzero finite three-primary group acted on by a group of order three has a nonzero fixed subgroup: count its orbits, whose nontrivial sizes are three. Applying this to the three-primary part of $\operatorname{Cl}(L)$ proves the first equivalence.

The product order has index three and local conductor $\mathfrak{p}_3^2$, with no conductor away from three. The standard exact sequence for extension of invertible ideals from an order to its normalization is

$$\mathcal{O}_L^\times \longrightarrow (\mathcal{O}_L/\mathfrak{p}_3^2)^\times/(\mathcal{O}/\mathfrak{p}_3^2)^\times \longrightarrow \operatorname{Pic}(\mathcal{O}) \longrightarrow \operatorname{Cl}(L) \longrightarrow 1.$$

The last map is surjective because every ideal class has an ideal representative coprime to the conductor. An invertible ideal whose extension is principal is obtained by gluing the trivial modules off the conductor and on its local support; their gluing ratio is a unit of the normalization modulo the conductor. Changing either trivialization divides by a residue unit of the order or by a global unit of the normalization. This proves exactness in the displayed form. The middle residue-unit quotient has order three, since the local rings are $\mathbb{F}_3[\lambda]/(\lambda^2)$ and $\mathbb{F}_3$, respectively. Its global-unit image is trivial. Indeed L has only the roots of unity $\{1,-1\}$. For every unit u, Kronecker's theorem gives $u/\overline{u} \in \{1,-1\}$. Reduction at the unique prime above three gives $u/\overline{u} = 1$ in $\mathbb{F}_3$, excluding the negative sign. Thus every unit lies in the real subfield R, whose full integer ring is $\mathbb{Z}[\theta] \subset \mathcal{O}$. The exact sequence now gives the asserted factor three. □

The quadratic class number $h(-3q)$ is computable in time polynomial in q, for example by enumerating reduced positive binary quadratic forms. This supplies a concrete test during a small-q search. No uniform bound under rational Dirichlet GRH has been proved here for an auxiliary q passing this class-number test together with the Frobenius λ^5-coefficient and component-count tests. In particular, an ideal-class cardinality has not been identified with a chosen theta coefficient-field degree without the additional reflex-norm and marking calculation.

37 What sparse cyclotomic compression leaves to select

The conductor r^{a+1} has a short binary description even when its cyclotomic field has large degree. The following calculation isolates the coefficient selection left by using this sparse presentation directly in characteristic p. It is a statement about this proposed compression method, rather than a lower bound for finite-field construction.

Proposition 37.1. *Let $p \neq r$ be primes with r odd, and put $d = \mathrm{ord}_r(p)$, $a = v_r(p^d - 1)$, and $g = r^{a-1}$. For $j = a, a+1$, write*

$$C_j = \mathbb{F}_p[X_j]/\Phi_{r^j}(X_j),$$

and let the tame group $H \cong (\mathbb{Z}/r\mathbb{Z})^\times$ act by $X_j \mapsto X_j^{\omega_j(c)}$. Set $A = C_a^H$ and $B = C_{a+1}^H$. Then

$$A \cong \mathbb{F}_p^g, \qquad B \cong (\mathbb{F}_{p^r})^g.$$

Under the embedding $X_a \mapsto X_{a+1}^r$, the algebra A is exactly the subalgebra of B fixed by the absolute Frobenius. Consequently, for $b \in B$, the polynomial

$$\prod_{i=0}^{r-1} (T - b^{p^i})$$

has coefficients in A, which need not be scalar in $\mathbb{F}_p$.

Proof. The order of p modulo r^a is d, whereas its order modulo r^{a+1} is dr. The first assertion follows from the valuation definition and the second from lifting the exponent, since r is odd and $r \nmid d$. The geometric points of C_j form a regular set for $(\mathbb{Z}/r^j\mathbb{Z})^\times$. Taking H-invariants passes to its orbits, a regular set for a cyclic wild group of order r^{j-1}. On that quotient, the Frobenius is trivial for $j = a$ and has exact order r for $j = a+1$. The finite-field product descriptions follow. The image of A in B is Frobenius-fixed and has dimension g, which is the dimension of the fixed algebra just computed; therefore the two algebras coincide. Frobenius cyclically permutes the factors of the displayed polynomial, proving the coefficient assertion. □

A consistent evaluation of the full coefficient algebra A is equivalent to a normalized maximal Teichmüller unit. For a particular b, whether its coefficients suffice depends on the subalgebra they generate. For instance $b = 1$ gives only scalars. The correspondence can itself be implemented with short arithmetic circuits; it is not an artifact of a large dense power basis.

Proposition 37.2. *Use A, C_a, g as above and put $D = \mathbb{F}_p[Z]/\Phi_r(Z)$. There is a natural bijection between unital $\mathbb{F}_p$-algebra homomorphisms $\rho : A \to \mathbb{F}_p$ and elements $U \in D$ satisfying*

$$U^g = Z, \qquad U^{r^a} = 1, \qquad \sigma_c(U) = U^{\omega_a(c)} \quad (c \bmod r).$$

From such a U, evaluation of an invariant arithmetic circuit in C_a under its corresponding ρ costs polynomial time in the circuit size, r, and $\log p$. Conversely, a consistent evaluation oracle for one such ρ, on $r-1$ explicitly constructible invariant circuits of size polynomial in $r, \log p$, gives U within the same polynomial bound apart from oracle time.

Proof. There is an isomorphism

$$A \otimes_{\mathbb{F}_p} D \longrightarrow C_a, \qquad Z \longmapsto X_a^g.$$

Over an algebraic closure, each tame orbit of primitive r^a-th roots maps bijectively to the primitive r-th roots under $x \mapsto x^g$. This proves the isomorphism of finite étale algebras, without selecting any component.

An evaluation ρ consequently sends X_a to an element $U \in D$ with all the stated identities. Conversely, such a U defines the equivariant homomorphism $C_a \to D$, $X_a \mapsto U$. Its restriction to A lands in $D^H = \mathbb{F}_p$, giving ρ. The two constructions are inverse.

For the computational assertion, use the basis $1, Z, \dots, Z^{r-2}$ of D, and let $b_0(Z), \dots, b_{r-2}(Z)$ be its dual basis for the nondegenerate trace pairing of the étale algebra D. This basis is computed by linear algebra of dimension $r-1$. Under the displayed isomorphism, write

$$X_a = \sum_{j=0}^{r-2} A_j Z^j, \qquad A_j = \sum_{c \in H} \sigma_c(X_a b_j(X_a^g)) \in A.$$

These are the usual dual-trace coefficient formulas. The powers have binary exponents of length $O(a \log r) = O(r \log p)$; the tame lifts are computed by Hensel lifting. Thus the $r-1$ displayed invariant circuits have polynomial size. Their consistent evaluations give $U = \sum_j \rho(A_j) Z^j$. In the reverse direction, substitute the known U into an invariant circuit. Its value in D is scalar by equivariance, so its scalar coefficient gives the value of ρ. The asserted bounds follow. □

The word “consistent” is essential: independently choosing a root for each coordinate polynomial need not preserve products or define a homomorphism. The proposition identifies the remaining operation in this specific sparse-algebra proposal with the maximal-unit input. It provides no new algorithm for that operation.

For example, when $r = 3$ and $p \equiv -1 \pmod{3^a}$ with exact valuation a, put $x = X_{a+1}$ and $b = x + x^{-1}$. Then $b \in B$, and its Frobenius polynomial is

$$T^3 - 3T - (x^3 + x^{-3}).$$

Its apparently single scalar parameter is an element of $A \cong \mathbb{F}_p^{3^{a-1}}$. In this cubic example a single consistent value y of $x^3 + x^{-3}$ suffices to recover the normalized unit. Indeed, any specialization W of X_a at this value satisfies $W + W^{-1} = y$ and has order 3^a. The two choices are W and W^{-1}, so their gth powers are the two primitive cube roots. Define $S_0 = 0$, $S_1 = 1$ and $S_{n+1} = yS_n - S_{n-1}$. Then $W^g = S_g W - S_{g-1}$, with $S_g = (W^g - W^{-g})/(W - W^{-1}) \neq 0$. In the formal algebra D the choice with gth power Z is consequently

$$U = (Z + S_{g-1})/S_g.$$

Binary powering of the two-by-two recurrence matrix uses $O(\log(g+1))$ field operations. The remaining operation is to find one such y; representing the coefficient by a short circuit does not select its value.

38 Cyclic certificates for descent

An auxiliary cyclotomic algebra supplies roots of unity, but a root constructed in that algebra must still be descended to the desired base. We now give two sufficient conditions for this descent. The first uses the cyclic period of a root and a residue map from a small real cyclotomic subfield. The second tests whether a cyclic automorphism lifts with unchanged order to a Kummer algebra. A successful test gives the required prime-degree field without selecting a component of the auxiliary algebra. The conditions are computable, but the construction of the auxiliary root does not guarantee that they hold.

Formal cyclotomic rings and the extension of the tame action to Kummer algebras occur in Lenstra (1991, Proposition (4.3), printed p. 337; Theorems (5.1)–(5.2), pp. 338–340). Kummer elements, cyclotomic algebra maps, and the alternative between a proper factor and a structured algebra also occur in Ivanyos et al. (2012, Lemmas 6.1–6.2, author pp. 30–31; Theorem 6.3, author p. 32), in their setting of completely split algebras over a prime field with smooth $p-1$. We make the descent conditions, their costs, and their limitations explicit below.

Arithmetic conventions. An element has *strict order* h if its image has order exactly h in every field component. All finite algebras have explicit bases and multiplication tables over their indicated fields. A split algebra of rank n over F means an algebra isomorphic to F^n; the component maps are not part of the input. Write $\mu_h(J) = \{x \in J^\times : x^h = 1\}$. Norms of finite free algebras are determinant norms.

The following elementary operations will be used throughout. For a nonzero nonunit b in a finite étale algebra J, multiplication by b gives complementary ideals

$$J = bJ \oplus \operatorname{ann}(b).$$

Compute its image and kernel, and decompose 1 in these two subspaces. The resulting identities are complementary idempotents, with explicit projection and inclusion maps. Treat both factors and patch their answers by addition. For a vector of coefficients, intersect their zero supports before deciding that the vector is nonzero on every component. Polynomial Euclid and module elimination split on nonunit pivots and then use unit pivots. A splitting tree beginning with dimension M has at most M nonzero leaves and $M-1$ proper splits. Restarting a polynomially bounded calculation after such a split adds polynomial overhead.

Fixed algebras are kernels of explicit linear maps. Constant-rank module quotients have bases on the base idempotent pieces; compute their determinant norms there and patch over the base. These procedures do not require a factorization or a monogenic presentation of a general product algebra. Nor do they average by a group order in the coefficient field. Thus characteristic dividing a cyclic or tame group order causes no omitted exception.

Proposition 38.1 (Descent from a short cyclic period)**.** *Fix an odd prime r and an integer $t \geq 0$. Let $p \neq r$ be prime, put $d = \operatorname{ord}_r(p)$, and let $\ell \neq r$ be an odd prime dividing $\Phi_d(p)$. For a supplied nonconstant monic divisor $g \mid \Phi_r$, put*

$$\mathcal{A} = \mathbb{F}_p[Z]/g(Z), \qquad B = \mathcal{A}[T]/\Phi_\ell(T), \qquad n = \ell - 1.$$

Choose a primitive residue g_ℓ modulo ℓ, and let σ fix $\mathcal{A}$ and send T to T^{g_ℓ}. Suppose $z \in \mathcal{A}^\times$ and $y \in B^\times$ are supplied with $y^r = z$. Let n' be the least positive integer for which $\sigma^{n'}(y) = y$.

If $v_r(n') \leq t$, then under A a simultaneous root $x \in \mathcal{A}$, $x^r = z$, can be constructed deterministically in bit time

$$(\ell + \dim_{\mathbb{F}_p} \mathcal{A} + \log p)^{O_{r,t}(1)}.$$

If $r \nmid n'$, the construction is unconditional. The period and the displayed condition are computable within these bounds. No factorization of g or Φ_ℓ, primitive ℓ-th root in $\mathcal{A}$, or component field is supplied. No bound on the period produced by an earlier root construction is asserted.

Proof. Every field component of $\mathcal{A}$ has cardinality $q = p^d$. Since $\ell \mid p^d - 1$, one has $p \neq \ell$, and B is split of rank n over each component of $\mathcal{A}$. The σ-action on its n geometric points is regular cyclic. Trial division of the numerical integer n and testing primitive residues cost polynomial time in ℓ. The least period n' divides n. Starting with n, remove a prime divisor s from the candidate period whenever $\sigma^{n'/s}(y) = y$; the resulting integer is the least period. All equality tests take place in the whole algebra.

Set

$$C = B^{\langle \sigma^{n'} \rangle}.$$

It is finite free of rank n' over $\mathcal{A}$ and contains y. An explicit basis is particularly simple: the elements $T, T^2, \ldots, T^{\ell-1}$ form an $\mathcal{A}$-basis of B, permuted regularly by σ, and sums over the orbits of $\sigma^{n'}$ form a basis of C. This argument remains valid when the characteristic divides the subgroup order.

If $r \nmid n'$, let $b = N_{C/\mathcal{A}}(y)$. Then $b^r = z^{n'}$. Choose integers u, v with $un' + vr = 1$; the unit $b^u z^v$ is the desired root.

Suppose $j = v_r(n')$ is positive and at most t, and put $m = r^j$. Let K_j be the unique degree-m subfield of $\mathbb{Q}(\zeta_\ell)$. It is cyclic and totally real: its defining subgroup has odd index, so it contains complex conjugation. If $f = \mathrm{ord}_\ell(p)$, then $f \mid d < r$. The Frobenius of p in the quotient of order m has order dividing both f and m, and is therefore trivial. Thus p splits completely in K_j.

For completeness, the integral data needed here are explicit. The elements ζ_ℓ^a, $1 \le a \le \ell - 1$, form an integral basis of $\mathbb{Q}(\zeta_\ell)$. The defining subgroup permutes them, so its orbit sums form an integral basis of $\mathcal{O}_{K_j}$. Counting residue pairs computes the multiplication table, with coefficients of absolute value at most ℓ. Its trace determinant gives a discriminant bounded by $\ell^{O_{r,t}(1)}$. This is the period construction used in Lemma 16.3. Proposition 16.1, under A, therefore constructs a residue map

$$\lambda : \mathcal{O}_{K_j} \longrightarrow \mathbb{F}_p$$

in time $(\ell + \log p)^{O_{r,t}(1)}$. Here $m \le r^t$ is fixed; no degree-uniform bound is inferred.

The period subalgebra

$$C_0 = B^{\langle \sigma^m \rangle}$$

is the scalar extension to $\mathcal{A}$ of $\mathcal{O}_{K_j}/p\mathcal{O}_{K_j}$. This follows directly from the same orbit-sum bases, without division by a subgroup order. Since $m \mid n'$, one has $C_0 \subset C$. Impose the values given by λ on all period basis elements:

$$C_\lambda = C \otimes_{C_0,\lambda} \mathcal{A}.$$

The ideal generated by the period-value differences and this quotient are computed by linear algebra. Over each component of $\mathcal{A}$, the quotient selects one fiber of a regular cyclic quotient. It is consequently finite free and split of rank $h = n'/m$ over $\mathcal{A}$, with $r \nmid h$. Constant-rank elimination uses the support refinements described above; no nonzero component of C_λ is discarded.

The image of y still has r-th power z everywhere. Its full norm $b = N_{C_\lambda/\mathcal{A}}(y)$ satisfies $b^r = z^h$. For $uh + vr = 1$, the element $b^u z^v$ is a simultaneous root in $\mathcal{A}$. All algebra dimensions are at most $(\ell - 1)\dim_{\mathbb{F}_p} \mathcal{A}$; the remaining matrix, integer, and powering operations have polynomial cost in these dimensions and $\log p$. This proves the asserted bound. □

Proposition 38.2 (A normalized cyclic lift)**.** *Let $F = \mathbb{F}_q$ have a supplied power-basis presentation. Let $r \mid q - 1$ be an odd prime, put $a = v_r(q - 1)$, and supply $\zeta \in F$ of order r. Let B be an explicitly presented split finite étale F-algebra of rank n, with a supplied F-automorphism σ acting as one n-cycle on its geometric points. Suppose*

$$u^{r^{a-1}} = \zeta \quad \text{for a supplied } u \in B.$$

There is an unconditional deterministic algorithm, polynomial in n, r and the bit length of the presentation of F, that constructs $c \in \mu_{r^a}(B)$ with $c^r = \sigma(u)/u$. For

$$E = B[X]/(X^r - u), \qquad \rho|_B = \sigma, \qquad \rho(X) = cX,$$

the map ρ is an automorphism and

$$w = \prod_{i=0}^{n-1} \sigma^i(c) \in \mu_r(F), \qquad \rho^n(X) = wX.$$

If $w = 1$, then ρ has exact order n, and E^ρ is a field of degree r over F. Its multiplication table and a monic irreducible degree-r polynomial over F are computable in the same bound. If $q = p^d$ with $\gcd(d, r) = 1$, an irreducible degree-r polynomial over $\mathbb{F}_p$ is also computable.

If $w \ne 1$, replacing c by cv, $v \in \mu_r(B)$, makes the lift have order n exactly when

$$N_{B/F}(v) = w^{-1}.$$

Finding such a v is polynomially equivalent to finding an idempotent $e \in B$ for which $r \nmid \dim_F(eB)$. When $r \nmid n$, a constant $v \in \mu_r(F)$ suffices.

Proof. The power equation implies that u is a unit of strict order r^a. Put $s = \sigma(u)/u$. Then $s^{r^{a-1}} = 1$; on each component of B, the element u^r generates the full subgroup of this order. Its discrete logarithm may differ between components.

A simultaneous digit construction. For $a = 1$, one has $s = 1$ and takes $c = 1$. Otherwise set $s_0 = s$. At stage $i = 0, \dots, a-2$, the remaining logarithm is divisible by r^i, so

$$t_i = s_i^{r^{a-2-i}} \in \mu_r(B).$$

For $0 \le j < r$, form

$$e_{i,j} = \frac{1}{r} \sum_{k=0}^{r-1} (t_i \zeta^{-j})^k, \qquad b_i = \sum_{j=0}^{r-1} e_{i,j} u^{j r^i}, \qquad s_{i+1} = s_i b_i^{-r}.$$

Since the characteristic is different from r, these expressions are defined. The elements $e_{i,j}$ are exactly the zero-support idempotents for $t_i - \zeta^j$: they are orthogonal and sum to 1. Thus b_i removes the i-th base-r digit of the componentwise logarithm of s to the base u^r. At termination $s_{a-1} = 1$, and

$$c = \prod_{i=0}^{a-2} b_i$$

satisfies $c^r = s$. Each b_i, and hence c, has order dividing r^a.

The construction uses $O(ar^2)$ algebra operations for the displayed idempotent sums, together with polynomially many repeated squarings whose exponents have $O(a \log r)$ bits. It does not scan r^a possibilities. It is functorial under an automorphism fixing the inputs u, s, ζ, because the support idempotents are uniquely specified. In particular, if $\sigma^m(u) = u$, then $\sigma^m(c) = c$. This specific digit choice will be used below. Although s is killed by r^{a-1}, c can have order r^a; for $a = 2$, a ratio equal to ζ can give $c = u$.

The lift and the fixed field. The equation $(cX)^r = \sigma(u)$ proves that ρ respects the defining relation; applying σ^{-1} and the corresponding inverse coefficient gives its inverse. Its n-th power is as displayed, with $\sigma(w) = w$. Regularity gives $B^\sigma = F$, and telescoping gives

$$w^r = \prod_{i=0}^{n-1} \frac{\sigma^{i+1}(u)}{\sigma^i(u)} = 1.$$

Thus $w \in \mu_r(F)$, and its exponent relative to ζ is found by at most r comparisons.

Write $q - 1 = r^a V$, where $r \nmid V$. On a component of B, a root $\alpha^r = u$ satisfies

$$\alpha^q = \alpha u^{(q-1)/r} = \zeta^V \alpha.$$

The q-power Frobenius cycles all r roots. Consequently $X^r - u$ is irreducible on every component, and E is a product of n copies of $\mathbb{F}_{q^r}$.

If $w = 1$, the order of ρ is exactly n, because its restriction to B already has that order. Every geometric ρ-orbit has length n, so E^ρ has F-rank r. Let τ be q-power Frobenius. If a nontrivial τ^j, $0 < j < r$, preserved a ρ-orbit, then for some geometric point P one would have $\tau^j P = \rho^k P$. Projection to B forces $k \equiv 0 \pmod{n}$, contradicting the free order-r action of τ above a base point. Thus Frobenius is an r-cycle on the r points of E^ρ, which is therefore the field $\mathbb{F}_{q^r}$.

Compute the fixed algebra as $\ker(\rho - 1)$, and obtain its multiplication table by solving for products of basis vectors. Any element outside its scalar F-line generates it, since r is prime. The characteristic polynomial of multiplication by that element is monic irreducible of degree r. For $q = p^d$, this field is $\mathbb{F}_{p^{dr}}$; the kernel of $x \mapsto x^{p^r} - x$ is its degree-r subfield over $\mathbb{F}_p$, from which the prime-field polynomial is obtained in the same way.

Normalization and its exact input. Replacing c by cv changes w to $wN_{B/F}(v)$. Write $w = \zeta^j \neq 1$. If e is an idempotent of rank k with $r \nmid k$, choose h with $hk \equiv -j \pmod r$. Then

$$v = 1 - e + \zeta^h e, \qquad N_{B/F}(v) = \zeta^{hk} = w^{-1}.$$

Conversely, the spectral idempotents e_h of any normalizing $v \in \mu_r(B)$ are computable by the same finite sums. Writing $k_h = \dim_F(e_h B)$, one has

$$N_{B/F}(v) = \zeta^{\sum_h hk_h} \neq 1.$$

At least one k_h is nonzero modulo r. All these ranks are matrix ranks, so both reductions are polynomial. For $r \nmid n$, use a constant $v = \zeta^h$ with $hn \equiv -j \pmod r$. For $r \mid n$, constants have norm 1; the equivalence supplies no idempotent by itself.

Finally, $a \log r \leq \log q$, and the largest F-dimension is nr. All matrix operations, support refinements, digit steps, and powers therefore satisfy the asserted absolute polynomial bound. □

Proposition 38.3 (A cyclic lift over the full cyclotomic base)**.** *Let $p \neq r$ be primes, with r odd, and put*

$$d = \operatorname{ord}_r(p), \quad q = p^d, \quad a = v_r(q-1), \quad D = \mathbb{F}_p[Z]/\Phi_r(Z).$$

Let $\ell \neq r$ be an odd prime dividing $\Phi_d(p)$, and set

$$B = D[T]/\Phi_\ell(T), \qquad n = \ell - 1, \qquad \sigma(T) = T^{g_\ell}, \quad \sigma|_D = 1,$$

where g_ℓ is a primitive residue modulo ℓ. Suppose $u_0 \in B$ is supplied with $u_0^{r^{a-1}} = Z$.

There is an unconditional deterministic algorithm of bit cost $(r + \ell + \log p)^{O(1)}$, with an absolute exponent, that tame-projects u_0 to u with the same bottom power, constructs the simultaneous digit root $c_0^r = \sigma(u)/u$, and tame-projects c_0 at level a to c. On $E = B[X]/(X^r - u)$, the maps

$$\rho|_B = \sigma, \qquad \rho(X) = cX,$$
$$\Gamma_b|_B = \gamma_b, \quad \Gamma_b(X) = X^{\omega_{a+1}(b)} \quad (b \in \mathbb{F}_r^\times),$$

where $\gamma_b(Z) = Z^b$ and $\gamma_b(T) = T$, commute and obey their indicated group laws. Their obstruction has the form

$$\prod_{i=0}^{n-1} \sigma^i(c) = Z^j \quad (j \in \mathbb{F}_r),$$

with j computable.

If $j = 0$, their simultaneous fixed algebra is a field of degree r over $\mathbb{F}_p$. The algorithm returns a monic irreducible degree-r polynomial and $U \in D$ satisfying

$$U^{r^{a-1}} = Z, \qquad \gamma_b(U) = U^{\omega_a(b)}, \qquad \operatorname{ord}(U) = r^a \text{ in every component.}$$

No field component of D or B is selected. If m is the least σ-period of the tame-projected u, then

$$r \mid n/m \quad \Longrightarrow \quad j = 0.$$

The sufficient condition is equivalently $v_r(m) < v_r(n)$; it does not require either valuation to be bounded.

Starting with only p, r, ℓ, the required u_0 can be constructed with the separately charged cofactor

$$C_\ell = \frac{\Phi_d(p)}{2^{v_2(\Phi_d(p))}\, r^a\, \ell^{v_\ell(\Phi_d(p))}}$$

in time $(r + \ell + C_\ell + \log p)^{O(1)}$, with an absolute exponent. The resulting test still need not give $j = 0$. Additional factors may be omitted only under the supplied-root or coprime-rank hypotheses, with the numerical charges, of Theorem 5.1.

Proof. All powers and actions below are computed in the displayed product algebras; their field components are used to justify the identities.

Tame normalization at the full level. Every component of D has size q, and B is split of rank n over each such component. Let $\Delta = \mathbb{F}_r^\times$, acting by the maps γ_b. For an element killed by r^a, define

$$\mathcal{P}(x) = \left(\prod_{b \in \Delta} \gamma_b(x)^{\omega_a(b)^{-1}} \right)^e, \qquad e(r-1) \equiv 1 \pmod{r^a}.$$

Exponents are read modulo r^a. The integer inverse e is not division by $r - 1$ in the coefficient field. Reindexing proves

$$\gamma_b(\mathcal{P}(x)) = \mathcal{P}(x)^{\omega_a(b)}.$$

The map $\mathcal{P}$ is a homomorphism and is the identity on elements already satisfying this character relation. The lifts $\omega_k(b)$ are computable by Hensel lifting, or by $b^{r^{k-1}} \bmod r^k$.

Set $u = \mathcal{P}(u_0)$. The r^{a-1}-st power of the inner product is Z^{r-1}, so $u^{r^{a-1}} = Z$, and u remains strict maximal. Because σ commutes with Δ, $s = \sigma(u)/u$ already has the tame character. The digit construction in Proposition 38.2, applied to the whole product algebra with $\zeta = Z$, gives $c_0^r = s$ and $c_0 \in \mu_{r^a}(B)$. Set $c = \mathcal{P}(c_0)$. Then

$$c^r = \mathcal{P}(s) = s, \qquad \gamma_b(c) = c^{\omega_a(b)}.$$

The projector is at level a, not $a - 1$, since c_0 can have order r^a.

The actions and the global obstruction. The element X in E has strict order r^{a+1}. The formula for Γ_b respects $X^r = u$, since $\omega_{a+1}(b)$ reduces to $\omega_a(b)$ modulo r^a. Multiplicativity of these lifts modulo r^{a+1} proves $\Gamma_b \Gamma_{b'} = \Gamma_{bb'}$. This is the tame extension of the action in Lenstra (1991, Proposition (4.3), printed p. 337). The maps commute with ρ: on B this is immediate, and on X it is the identity

$$\gamma_b(c) = c^{\omega_{a+1}(b)},$$

which follows from $c^{r^a} = 1$ and the displayed level-a character relation.

The cyclic product $w = \prod_i \sigma^i(c)$ is in D, satisfies $w^r = 1$, and obeys $\gamma_b(w) = w^b$. Evaluate at a geometric point $Z = \zeta$, and write $w(\zeta) = \zeta^j$. The character relation then gives $w(\zeta^b) = (\zeta^b)^j$ at every other point. Thus $w = Z^j$ for one global $j \in \mathbb{F}_r$; testing the r possibilities computes it.

Why the simultaneous fixed algebra is a field. Suppose $j = 0$. Then ρ has exact order n. The commuting group $C_n \times \Delta$ acts freely on the geometric points of E: if $\rho^k \Gamma_b$ fixes a point, its Z-coordinate first gives $b = 1$, and its T-coordinate then gives $k \equiv 0 \pmod{n}$. The algebra

$$H = E^{\langle \rho, \Gamma_b : b \in \Delta \rangle}$$

therefore has dimension r over $\mathbb{F}_p$. Write $q - 1 = r^a V$, $r \nmid V$. Frobenius of degree d fixes B and sends

$$X \longmapsto Z^V X.$$

No nontrivial power of this order-r action can preserve one $C_n \times \Delta$-orbit: projection to Z and T forces the compensating group element to be the identity, as above. It acts as an r-cycle on the r points of H. Prime-field Frobenius commutes with this cycle and has it as its d-th power; hence it too is an r-cycle. Consequently $H \cong \mathbb{F}_{p^r}$. Intersecting the kernels of the displayed actions computes H, its multiplication table, and an irreducible degree-r polynomial from any basis vector outside the scalar line.

A simultaneous Hilbert–90 construction of U. The additional output can be obtained directly in B. Solve the D-linear equation

$$L = \{h \in B : \sigma(h) = c^{-1}h\}.$$

Over one field component of D, the n split coordinates of B give a cyclic recurrence. It is consistent exactly because $N(c) = 1$; its solution space is one-dimensional, and a nonzero solution is nonzero on all n coordinates. Thus L is a free rank-one D-module, whose generators are units of B.

Compute an $\mathbb{F}_p$-basis of L by linear algebra. For each vector, its zero support over D is the common zero support of its coefficients in the displayed D-basis of B. Starting with the zero vector, add a basis vector only on the as-yet-unfilled D-support where that vector is nonzero. These additions use orthogonal idempotents, so there is no cancellation on a filled support. The basis spans every fiber of L, and after at most $\dim_{\mathbb{F}_p} D$ support refinements the resulting h is nonzero on every D-component. The cyclic recurrence then makes it a unit on every B-component. This constructs a generator without selecting a component as an output.

The recurrence and the root equation give $\sigma(h^r u) = c^{-r} h^r \sigma(u) = h^r u$, so

$$U_0 = h^r u \in B^\sigma = D.$$

On every field component, multiplying by the r-th power h^r preserves the nontrivial r-power class of u. Hence the r-primary part of U_0 has exact order r^a. Choose $e_1 V \equiv 1 \pmod{r^a}$, and put $U_1 = U_0^{Ve_1}$. This kills the prime-to-r part and preserves the r-primary part. The bottom power $U_1^{r^{a-1}}$ is strict primitive of order r, possibly Z^k with different nonzero k on different components of D. Its spectral idempotents identify these $r-1$ possibilities. On the k-support raise U_1 to the inverse of k modulo r^a, and patch to obtain W with $W^{r^{a-1}} = Z$. Finally $U = \mathcal{P}(W)$ has this bottom power, strict order r^a, and the required tame character.

The period certificate. Let $m \mid n$ be the least period of the tame-projected u. The maps σ^m fix u, s, Z, so the specified functorial digit construction fixes c_0. The tame maps commute with σ; therefore $\sigma^m(c) = c$ as well. Put $W_m = \prod_{i=0}^{m-1} \sigma^i(c)$. Cyclically shifting the product shows $W_m \in D$, and

$$W_m^r = \frac{\sigma^m(u)}{u} = 1, \qquad w = W_m^{n/m}.$$

Thus $r \mid n/m$ implies $w = 1$. An arbitrarily chosen root of s would not justify the period assertion; the digit construction is part of the algorithm. Factoring the numerical n and testing divisors computes m.

Absolute cost from the supplied unit. The prime-field dimensions of D, B, E are

$$r - 1, \qquad (r-1)n, \qquad r(r-1)n.$$

Moreover

$$a \log r \le d \log p \le (r-1) \log p.$$

The digit construction, both tame projections, all action matrices, invariant kernels, the rank-one module, and characteristic polynomials therefore have absolute polynomial bit cost in $r, \ell, \log p$. Exponents at levels a and $a+1$ have polynomial bit length. Support refinements have polynomial overhead, and no step enumerates a numerical primary power.

The separately charged construction of u_0. It remains to justify using Theorem 5.1 over the larger base B, rather than silently supplying u_0. Its root-extraction proof applies to any explicitly presented finite étale $\mathbb{F}_p$-algebra J of dimension M whose field components all have size $q = p^d$, provided a strict primitive r-th root and the stated selected prime-order roots are supplied. More explicitly, put $N = \Phi_d(p)$, and let $\mathcal{S}, \mathcal{T}$ be disjoint sets of odd primes other than r dividing N. For each $\ell' \in \mathcal{S}$, supply an element of strict order ℓ' in J; for $\ell' \in \mathcal{T}$, require $r \nmid \ell' - 1$, with no corresponding root supplied. The residual cofactor is

$$C = \frac{N}{2^{v_2(N)} r^a \prod_{\ell' \in \mathcal{S} \cup \mathcal{T}} \ell'^{v_{\ell'}(N)}}.$$

No factorization of C is required. Here are the operations and their costs.

In a field component, the proof uses only its cardinality, the strict primitive root, and the promise $z^{(q-1)/r} = 1$. For the binomial algebra it uses the cyclotomic projection with exponent

$$E_d = \prod_{\lambda \mid d,\ \lambda \text{ prime}} (1 - p^{d/\lambda}).$$

There are $h = 2^{\omega(d)} \le d$ Frobenius terms ($h = 1$ for $d = 1$). The projected values lie in $\Phi_d(p)$-torsion, and the rational-function count uses fewer than $rhC \le r^2 C$ prime-field shifts, where C is the residual numerical cofactor. This count and its degree estimate do not depend on M.

Frobenius, powering, inversions, coefficient comparisons, and polynomial Euclid are operations in the explicitly based algebra. The support refinements at the beginning of this section handle all nonunit pivots and component-dependent equalities. If a proper monic factor of $X^r - z$ has degree j, $0 < j < r$, its signed constant coefficient b satisfies $b^r = z^j$; for $uj + vr = 1$, the root is $b^u z^v$. These componentwise identities do not require a special coordinate presentation of J.

A missing selected prime ℓ' with $r \nmid \ell' - 1$ adjoins one algebra of rank $\ell' - 1$. The root is assembled on all its components before taking the full norm back to J. The largest simultaneous binomial and witness dimension is at most $Mr(\ell' - 1)$. No rational-shift search is nested inside this auxiliary-prime adjunction, and no dimension is a product over the selected primes. Each encountered split partitions a bounded positive dimension, so retaining both factors and restarting the bounded arithmetic adds only polynomial overhead.

When $p < rhC$, the proof instead uses absolute Frobenius. For an étale algebra J', compute $\ker(x \mapsto x^p - x) \cong \mathbb{F}_p^k$, where k is its number of field factors. Comparing a basis of this fixed algebra with all p prime-field scalars separates those factors. The cost is polynomial in $p, \dim J', \log p$. After applying this to the base and its split binomial algebra, the root in each degree-d quotient is pulled back by linear algebra and patched across the base. The enumeration is over $p < r^2 C$ scalars, not over q or over all elements of J.

Consequently the larger-base root interface has cost

$$\left(M + r + C + \sum_{\ell' \in \mathcal{S} \cup \mathcal{T}} \ell' + \log p \right)^{O(1)}$$

with an absolute exponent, and exactly the selected-root and coprime-rank hypotheses of the cited theorem. On $J = B$, the formal element T is a strict primitive ℓ-th root. It supplies that

selected root, allowing the full ℓ-primary factor to be omitted from the construction's cofactor. With no other odd prime selected, the residual is precisely C_ℓ.

Starting with $W_1 = Z$, extract roots $W_{i+1}^r = W_i$ for $1 \le i < a$. An element of order r^i, $i < a$, is an r-th power in every component of size q, and every extracted root has order r^{i+1}. Thus $u_0 = W_a$ has the required bottom power and strict order. There are at most $a - 1 \le \log_2 q$ calls. All valuations and exponents use polynomially many bits, and the numerical $\Phi_d(p)$ has $O(r \log p)$ bits. This proves the separately stated cofactor cost.

After this construction the algorithm computes j. Neither the cofactor estimate nor the existence of u_0 forces $j = 0$; the sufficient condition remains a test on the constructed algebraic data. □

Proposition 38.4 (A primitive cyclic obstruction)**.** *In Proposition 38.2, suppose $r \mid n$ and $w = \zeta^j$ with $j \not\equiv 0 \pmod r$. Then ρ has exact order nr, acts regularly on the geometric points of E, and*

$$E^{\langle \rho^r \rangle} = B^{\langle \sigma^r \rangle}.$$

This is a split rank-r algebra over F. If $q - 1 = r^a V$ and $jk \equiv V \pmod r$, then $\mathrm{Frob}_q = \rho^{nk}$ on E.

Given an F-algebra map $\lambda : B^{\langle \sigma^r \rangle} \to F$, form the corresponding quotients B_λ, E_λ. The induced maps σ^r, ρ^r have exact orders $n/r, n$, respectively, and their residual lift obstruction is still w. These quotients and actions are computable by linear algebra once λ is supplied.

For an odd prime $\ell \mid q - 1$, take $B = F[T]/\Phi_\ell(T)$ and $\sigma(T) = T^{g_\ell}$, with g_ℓ primitive modulo ℓ. Successive compatible selections of s such quotients select a residue map on the canonical degree-r^s subfield of $\mathbb{Q}(\zeta_\ell)$, with values in F. In the setting of Propositions 38.1 and 38.3, these subfields split completely at p, so the number-field residue maps may be taken with values in $\mathbb{F}_p$. This describes the canonical cyclic-subgroup recursion; it is not an impossibility statement for other descent methods.

Proof. Restriction to B makes the order of ρ a multiple of n, and ρ^n has order r. Since $\rho^{nr} = 1$, its order is exactly nr. Its orbit on a geometric point traverses all n base points; after these n steps, ρ^n cycles the r points in the fiber. The action on all nr points is therefore regular.

The subgroup generated by ρ^r has order n, so its fixed algebra has rank r. The included algebra $B^{\langle \sigma^r \rangle}$ also has rank r and is fixed by ρ^r, proving equality. It is split because it is a subalgebra of the split algebra B. There is also a direct Frobenius verification. Both Frob_q and ρ^{nk} fix B and multiply X by ζ^V. Since $r \mid n$, this Frobenius belongs to $\langle \rho^r \rangle$, as the split fixed algebra requires.

The geometric fiber selected by λ in E is one orbit of ρ^r, of length n. The base fiber has n/r points, regularly cycled by σ^r. On this quotient,

$$(\rho^r)^{n/r}(X) = \rho^n(X) = wX.$$

Thus the primitive obstruction is unchanged. If r still divides the residual base rank n/r, the same situation recurs. When that rank is prime to r, a constant correction is available as in Proposition 38.2. No procedure for finding λ is supplied by this identity.

For the cyclotomic algebra, the fixed algebras

$$B^{\langle \sigma^{r^s} \rangle} \qquad (r^s \mid \ell - 1)$$

are the scalar extensions of the reductions of the degree-r^s subfields, by their orbit-sum integral bases. Choosing successive compatible fibers is exactly choosing one point of the largest such fixed algebra. In general the residue map has values in F. In the specialized setting, $\mathrm{ord}_\ell(p) \mid d < r$, so its image in every r-power cyclic quotient is trivial; the subfields split at p and their residue values lie in $\mathbb{F}_p$.

Over the full base D, the first fixed algebra is likewise split of rank r over D. On it, prime-field Frobenius equals γ_p, because the T-action of p has order prime to r and is trivial in the cyclic quotient of order r. Its tame fixed algebra is therefore split over $\mathbb{F}_p$. Tame descent does not turn this canonical first quotient into the required field.

Repeatedly applying Proposition 16.1 to these total degrees $r, r^2, \dots$ gives no fixed exponent when their depth grows. A theorem controlled by relative degree would be an additional ingredient. The assertion here concerns only this specific invariant-subgroup recursion. □

Proposition 38.5 (Balanced cyclic data)**.** *Let $F = \mathbb{F}_q$ contain ζ of odd prime order r, and let $z \in F^\times$ be an r-th power. Put*

$$R = F[Y]/(Y^r - z), \quad \tau(Y) = \zeta Y, \quad B = R^\kappa, \quad n = r\kappa, \quad y = (Y, \dots, Y),$$

where $\kappa \geq 1$, and define

$$\sigma(b_0, \dots, b_{\kappa-1}) = (b_1, \dots, b_{\kappa-1}, \tau(b_0)).$$

These presentations are constructed without finding a root of z in F. The action of σ is regular cyclic of order n; the element y has exact period n and characteristic polynomial $(X^r - z)^\kappa$. The algebra generated by the cyclic coboundaries $\sigma^j(y)/y$ is exactly $F^\kappa \subset B$.

For the grading

$$B = \bigoplus_{w \bmod r} F^\kappa y^w,$$

multiplication, powers, inversion of homogeneous units, cyclic translates, addition and comparison within one degree, and support refinements from those comparisons produce only idempotents in F^κ. Their ranks in B are multiples of r. On every surviving union of blocks, y retains a characteristic polynomial of the form $(X^r - z)^{\kappa'}$. This assertion concerns only the stated operations.

If $n = r^b h$, $r \nmid h$, then for $0 \leq j < b$,

$$B^{\langle \sigma^{r^j} \rangle} \subset F^\kappa, \qquad \mathrm{Tr}_{B/B^{\langle \sigma^{r^j} \rangle}}(y^w) = 0 \quad (1 \leq w < r),$$

and

$$N_{B/B^{\langle \sigma^{r^j} \rangle}}(y) = z^{\,n/r^{j+1}}.$$

If $a = v_r(q-1) \geq 2$, z has exact order r^{a-1}, and $\zeta = z^{r^{a-2}}$, the specified digit lift for $u = y$ has primitive obstruction $w = \zeta$.

There is also an existence example in the actual algebra of Proposition 38.3. Whenever $a \geq 2$ and $r \mid \ell - 1$, put $f = \mathrm{ord}_\ell(p)$. A tame-equivariant maximal unit u can have exact period $(\ell - 1)/f$ and make the digit-and-tame construction return $w = Z^f$, which is primitive. This does not assert that the construction from the original inputs returns that particular unit.

Proof. Write $E_0 = F^\kappa$ for the blockwise scalar subalgebra of B.

The constructible carry algebra. The algebra R is split and étale, but its polynomial presentation uses only the supplied z. The restriction of σ to F^κ is the κ-cycle on blocks, and σ^κ is τ on each block. Hence σ has order $r\kappa$ and acts regularly on the geometric points. For a shift $k = q_0\kappa + s$, $0 \leq s < \kappa$, the components of $\sigma^k(y)$ have multipliers ζ^{q_0} and, when $s > 0$, ζ^{q_0+1}. Equality with y requires $s = 0$ and $r \mid q_0$. Thus its exact period is n.

The first coboundary is

$$\sigma(y)/y = (1, \dots, 1, \zeta).$$

Subtracting 1 and dividing by $\zeta - 1$ gives the last block idempotent. Its cyclic translates give all block idempotents. Every cyclic coboundary is blockwise scalar, proving the assertion about

their generated algebra. The characteristic polynomial of multiplication by Y on R is $X^r - z$, so the characteristic polynomial on the product has the displayed balanced form.

The grading is direct because $1, Y, \dots, Y^{r-1}$ is an F-basis of R. The listed arithmetic operations preserve homogeneity, and σ preserves each degree. A same-degree difference has the form cy^w with $c \in F^\kappa$. Since y is a unit, its zero support is a union of whole R-blocks. Thus every listed support refinement has rank divisible by r, and the balanced polynomial persists on each union of blocks. For example, the norm of a homogeneous unit on κ' blocks is

$$\left(\prod_i c_i\right)^r z^{w\kappa'}.$$

This identity does not exclude an independent scalar root search.

For $0 \le j < b$, the divisibility $r^j \mid \kappa$ implies that $\langle \sigma^{r^j} \rangle$ contains σ^κ, whose fixed algebra is F^κ. The trace of y^w, $1 \le w < r$, over this order-r subgroup vanishes because $\sum_i \zeta^{iw} = 0$. Its norm of y is z, since $\zeta^{r(r-1)/2} = 1$ for odd r. Transitivity of trace and norm gives the asserted formulas.

Under the additional maximal-order hypotheses, $y^{r^{a-1}} = \zeta$. The digit construction takes

$$c = (1, \dots, 1, Y^{r^{a-2}}).$$

Its r-th power is the displayed coboundary, and

$$\prod_{i=0}^{n-1} \sigma^i(c) = N_{R/F}(Y)^{r^{a-2}} = z^{r^{a-2}} = \zeta.$$

Thus the obstruction is primitive. Every μ_r-valued element of F^κ has norm 1 from B, because each block has rank r; such elements do not normalize this lift.

An inhomogeneous polynomial test can escape the stated class. If it yields a factor of $Y^r - z$ of degree j with $0 < j < r$, the signed constant coefficient has r-th power z^j, and Bezout exponents already give a scalar root. No bound or impossibility for finding such a test follows from the grading argument.

An example inside the full cyclotomic algebra. Assume now the hypotheses of Proposition 38.3, with $a \ge 2$ and $r \mid n = \ell - 1$. Put $f = \mathrm{ord}_\ell(p)$, $M = n/f$, and $\kappa_0 = M/r$. Since $f \mid d < r$, one has $r \nmid f$ and κ_0 is an integer. Choose, only to describe geometric values, a primitive r-th root ζ, an element $\xi \in \mathbb{F}_q$ of order r^a with $\xi^{r^{a-1}} = \zeta$, and a primitive ℓ-th root $\eta \in \mathbb{F}_q$.

For s modulo n, define

$$k(s) = \left\lfloor \frac{s \bmod M}{\kappa_0} \right\rfloor \in \{0, \dots, r-1\},$$

and assign, at the geometric point $Z = \zeta^b,\ T = \eta^{g_\ell^s}$,

$$u(Z, T) = \xi^{\omega_a(b)} \zeta^{bk(s)}.$$

These values define an element of the original $\mathbb{F}_p$-algebra B. Indeed, if $p = g_\ell^h \bmod \ell$, then $\gcd(n, h) = M$, so multiplying the T-exponent by p shifts s by a multiple of M and preserves $k(s)$. Also $p^d \equiv 1 \pmod{r^a}$ with $d \mid r - 1$, whence

$$p \equiv \omega_a(p \bmod r) \pmod{r^a}.$$

Thus the values at the Frobenius-transformed geometric points are the p-th powers of the displayed values. Frobenius-compatible functions on the geometric points of a finite étale algebra are exactly its elements. The same formula directly gives $\gamma_b(u) = u^{\omega_a(b)}$ and $u^{r^{a-1}} = Z$.

The block word $k(s)$ has exact period M: a smaller shift either moves a boundary inside a constant block or rotates the r distinct block values nontrivially. Hence u has this exact period, and $v_r(M) = v_r(n)$. The ratio $\sigma(u)/u$ is Z at the r block boundaries in each M-period and 1 elsewhere. The specified digit construction therefore takes $c_0 = u^{r^{a-2}}$ on those boundaries and 1 elsewhere. The boundary support is tame-invariant, so this c_0 already has the tame character and its projector does not change it.

At a fixed $Z = \zeta^b$, the starting values at the r boundaries are $\xi^{\omega_a(b)}\zeta^{bj}$, $0 \le j < r$. Their product, raised to r^{a-2}, is

$$\left(\xi^{r\omega_a(b)}\zeta^{br(r-1)/2}\right)^{r^{a-2}} = \xi^{r^{a-1}\omega_a(b)} = \zeta^b.$$

This includes $a = 2$, and uses oddness of r. Thus the product of c_0 over one period is Z. There are f periods in the full cycle, giving

$$w = Z^f \neq 1.$$

This is an existence example in the actual cyclotomic algebra. The choices of ξ, η describe that example; they are not additional data claimed to be computable from the original inputs within the desired bound.

Small explicit instances. For the carry algebra take $F = \mathbb{F}_{19}$, $r = 3$, $a = 2$, $\zeta = z = 7$, and $\kappa = 1$. The roots of $Y^3 - 7$ are $4, 9, 6$, whose product is $216 \equiv 7 \pmod{19}$. Thus $c = Y$ has order 9, the obstruction is $w = 7$, the lifted automorphism has order 9, and its cube fixes the split rank-three algebra R. This also shows concretely why projecting c at level $a - 1$ would be invalid.

The full cyclotomic example can already retain two powers of the target prime in its period: take

$$p = 683, \qquad r = 3, \qquad \ell = 19.$$

Here $d = 2$, $a = v_3(683^2 - 1) = 2$, $f = \mathrm{ord}_{19}(683) = 2$, and $n = 18$. The displayed construction has period $M = 9$ and obstruction $w = Z^2$. These instances illustrate the algebraic assertions; they impose no behavior on a different choice of the simultaneous maximal unit. □

Proposition 38.6 (A quotient by the small Frobenius orbit)**.** *Let $p \neq r$ be primes with r odd, $d = \mathrm{ord}_r(p)$, and let $\ell \neq r$ be an odd prime dividing $\Phi_d(p)$. For a nonconstant monic $g \mid \Phi_r$, put*

$$\mathcal{A} = \mathbb{F}_p[Z]/g, \qquad B = \mathcal{A}[T]/\Phi_\ell(T), \qquad f = \mathrm{ord}_\ell(p).$$

Let ψ fix $\mathcal{A}$ and send T to T^p. Then B^ψ has $\mathcal{A}$-rank $(\ell - 1)/f$. Any supplied simultaneous root $y \in B^\times$, $y^r = z \in \mathcal{A}^\times$, descends unconditionally to a simultaneous root in B^ψ, in absolute polynomial time in $r, \ell, \dim_{\mathbb{F}_p} \mathcal{A}, \log p$. However,

$$v_r((\ell - 1)/f) = v_r(\ell - 1).$$

The prime-field fixed algebra

$$D_0 = \left(\mathbb{F}_p[T]/\Phi_\ell(T)\right)^{\mathrm{Frob}_p} \cong \mathbb{F}_p^{(\ell-1)/f}$$

is computable by linear algebra. Selecting a degree-f field component of $\mathbb{F}_p[T]/\Phi_\ell(T)$ is polynomially equivalent to selecting an $\mathbb{F}_p$-algebra map $D_0 \to \mathbb{F}_p$. The quotient preserves the r-primary depth. When several components occur, choosing one still requires a point of the fixed algebra.

Proof. Since $\ell \mid p^d - 1$, the integer f divides $d < r$. The automorphism ψ has order f, and its orbits on the $\ell - 1$ primitive cyclotomic points all have length f. Its fixed algebra therefore has the asserted rank. For the supplied y, set

$$b = \prod_{i=0}^{f-1} \psi^i(y).$$

Then $b \in B^\psi$ and $b^r = z^f$. Choose integers u, v with $uf + vr = 1$. The element $b^u z^v$ is the required simultaneous root. All products and powers have polynomial cost in the displayed dimensions; $f < r$, and no field component is used as an input. Since $r \nmid f$, the valuation identity follows.

The prime-field cyclotomic algebra is a product of $(\ell - 1)/f$ copies of $\mathbb{F}_{p^f}$. Its absolute Frobenius fixed algebra is the product of their prime fields, computed as the kernel of $x \mapsto x^p - x$. A field-component map restricts to an $\mathbb{F}_p$-valued point of D_0. Conversely, a supplied point of D_0 identifies its primitive idempotent by the kernel and annihilator construction. That idempotent cuts out the corresponding degree-f field in the original cyclotomic algebra, with an explicit basis and multiplication table. Both reductions are linear algebra. Thus removing the small Frobenius orbit leaves precisely the stated point-selection problem. □

These propositions give additional computable sufficient conditions for descent. The cost of constructing the auxiliary unit and the condition that its obstruction vanish are separate requirements. Proposition 38.5 shows that a regular cyclic action, maximal order, and tame equivariance do not force that vanishing. The numerical input classes established earlier are therefore unchanged by these conditional criteria.

39 Questions remaining

The remaining problem is a bound uniform in the prime degree and the valuation for a prescribed hard input. The separate polynomial-time output algorithms of Tang, accepting OpenAI's 2026 prime-field preprint and CT with Tang's revision of Lemma 8.1, give the output under those named preprint inputs. Theorem 1.3 and Corollary 1.4 prove the construction using A alone for every fixed odd prime degree and fixed valuation. The unsettled uniform case is

$$r \nmid p - 1, \qquad a = v_r(p^{\mathrm{ord}_r(p)} - 1) \geq 2,$$

with r and a allowed to grow and a requested bound polynomial in $r, \log p$ with an absolute exponent. Theorem 1.1 also handles arbitrary valuation whenever, at a polynomially bounded cutoff L, its residual cofactor C_L is polynomially bounded in $r, \log p$. This removes the full primary parts of every prime at most L that is not one modulo r, as well as those of two and r. It includes Theorem 5.3 and the classical smooth cubic family of Corollary 5.5. Under A, Corollary 5.6 also removes every prime in any fixed finite set in the cubic case. More generally, Theorem 1.2 removes every prime $\ell \leq L$ with $v_r(\ell - 1) \leq t$, for each fixed r, t, leaving $C_{r,t,L}(p)$. Its exponent depends on r, t, although the valuation a and the removed primary exponents remain unbounded. The uniform problem remains where the applicable residual cofactor is large at every allowed cutoff; no density or infinitude assertion about these input classes is needed for the algorithms. The CM dimension r^{a-1} and its explicit theta representation prevent the displayed fixed-parameter proof from giving that bound. Proposition 32.1 restricts one attempt to replace its characteristic-zero field by a compositum of bounded-degree fields; it imposes no lower bound on other finite-field algorithms.

The classical short-auxiliary-prime estimate remains a separate open derivation even in the solved output strata. The distinction concerns the ingredients of the requested derivation. We make no assertion of logical independence from A. Accepting CT, the 2026 preprint

theorem with Tang's revised Lemma 8.1, gives a separate conditional route to the output and auxiliary-prime conclusions.

The analytic question. The classical sufficient RH hypothesis is the single induced Hecke function in (22), equivalently the pure-field zeta function under A. A weaker sufficient input is the one-sided window of Proposition 18.6: a fixed gap below twice the principal term for the smoothed pure-field prime sum at a scale polynomial in r and $\log p$. The residual-zero identity in Section 24.2 makes strict positivity equivalent to a good prime in its support. It therefore describes precisely what must be proved, but does not establish that positivity. The full tensor cone permits equality at the wrong endpoint, and arbitrary rational twists preserve a conductor contribution of size at least p^{r-1} when $\mathrm{ord}_r(p) > 1$. The quoted twist-density bound does not certify a fixed zero-free width at a polynomial-logarithmic twist scale. A successful analytic argument from A must add information beyond these displayed constraints. The zero-free strip for Hecke L-functions of OpenAI (2026b, Theorem 1.2, p. 3) is information of another kind, and OpenAI (2026a, Proposition 11.3, p. 46) derives the auxiliary prime from it.

The algebraic question. Theorem 1.7 identifies an exact alternative input: a maximal r-primary unit with the stated Teichmüller equivariance in the explicit cyclotomic algebra. A field of degree divisible by r already provides it, but this reverse reduction does not construct that field from its order. An explicit factor-count witness with component count prime to r also suffices. Complete cyclotomic products, their symmetric compressions, fixed candidate menus, and the semiprimitive Jacobi sums considered here do not produce such a witness uniformly. These are limitations of the specified constructions, with their hypotheses; none is a general impossibility theorem.

The cyclic certificates in Section 38 make two further sufficient conditions effective. At each fixed prime r, a root whose cyclic period has bounded r-adic valuation descends through a fixed-degree real residue map; a maximal unit whose normalized cyclic lift has unchanged order gives the field directly. A strict drop in the r-adic period is one sufficient condition for the latter test. The construction of the auxiliary root does not guarantee these conditions. Correcting a nontrivial lift obstruction requires a prime-to-r-rank idempotent, and the canonical successive quotient selections have growing degree. Thus these certificates do not remove the remaining uniform input condition.

Proposition 31.4 gives another explicit interface: separating ideal lattices determine a residue map by valuation-support idempotents. Polarized generator recovery applies once such ideals are given, but does not construct their lattices from a formal prime-factor exponent vector. The local trace-one residues of Proposition 31.3 do not by themselves supply the missing maximal primary unit through the specified multiplicative operations.

Elliptic and class-field inputs. The elliptic route still needs an effectively constructed curve whose Frobenius discriminant has r-adic valuation one. In the cubic prime-field Tate family, its success is exactly the noncubic class of $k - 3\zeta_3$. Given an ordinary prime-field parent, the restricted nonparent selector is polynomial-time equivalent to general promised-cube extraction. The two-trial construction over $\mathbb{F}_{p^4}$ supplies an ordinary curve with rational three-torsion, but changes that parent domain and does not force the necessary noncube. The CM route needs a polynomially bounded order with prime class nonzero modulo r; for $r = 3$ its stated budget is equivalently an inert complex cubic field within the same discriminant bound. Neither quadratic splitting nor existence of the relevant isogeny class supplies these effective choices.

What the average theorems settle. For fixed r, the short auxiliary search succeeds outside $O_{r,C_r}(\sqrt{P}\log\log P)$ primes in a dyadic interval. Theorem 1.12 gives a separate bound uniform in the degree for every selected menu of the stated size, and an almost-all result inside the

hard-input subclass; it still leaves prescribed exceptional primes unresolved. The simultaneous $12r$ table gives both algebraic algorithmic interfaces for almost all characteristics in the displayed degree range, and the cubic Tate theorem gives a direct output after $O(\log p)$ tests outside the displayed exceptional set. The supersingular point formulation is equivalent to finding a maximal three-primary unit; the two unconditional ordinary quartic-field trials do not by themselves supply a noncube. These results retain exceptional input primes. The character-product lower comparisons show that exceptional primes with long unsuccessful searches do occur; their lower scales still allow a uniform bound with a larger logarithmic power. The fixed-parameter CM construction removes exceptional inputs from every fixed prime-degree and valuation stratum. A bound uniform in both parameters is known in outcome: OpenAI (2026a, Theorem 1.1, p. 2), Shoup's reduction and Theorem 1.7 give the maximal equivariant unit in time polynomial in r and $\log p$, from a zero-free strip for Hecke L-functions. Such a bound from A alone, or a construction of that unit with no analytic input, remains the central question for these methods.

Method and records. The mathematics and text were produced by GPT-6 Astra in ChatGPT Pro mode. Separate statement-only re-proofs used the same model, also in ChatGPT Pro mode. Claude Opus 5.5 at effort high made targeted proof checks of the fixed-degree average theorem and its character lemmas, of the logarithmic cubic search and its quadratic-value and independence lemmas, and of the root-extraction lemma behind the cofactor theorem, with the recorded finite computations. GPT-6.2 Astra at reasoning effort ultra (OpenAI Codex) edited the paper, checked the two field-construction interfaces against Tang's source, re-derived the cyclic comparison and norm steps, and reran the fourteen finite checks. The program titles in the text are listed in the separate computational supplement, *Finite examples for constructing prime-degree fields*. The examples test exact finite calculations; the CM examples exercise the extraction from Frobenius data. The full theta-recognition and point-counting algorithm is proved at fixed parameters. These checks are not formal certificates of the whole paper. No human has checked this paper. The original portions of this text are dedicated to the public domain under CC0 1.0. The exposition of Tang's reductions in Section 21 adapts Qiyue Tang's *Deterministic Polynomial-Time Factorization over Finite Fields*, version 2, with notation and explanations adapted here. The underlying source is licensed under CC BY 4.0; its original attribution and licence remain applicable.

Checks recorded for version 2. Nine reviews of the 161-page edition, by GPT-6.2 Astra at reasoning effort ultra (OpenAI Codex), covered the cofactor theorems, the analytic sections, the CM construction, the cyclic splitting (independently, from its statement), the elliptic and late sections, and two whole-text readings; their corrections are in this text. Claude Opus 5.5 at effort high, in intakes of 23 and 24 November 2026, made targeted proof checks of Theorems 5.1 and 1.1, of Lemma 16.3, of the proof of Theorem 1.2 given Proposition 16.1, and of the conditional step, the CM type and the separating masks of that proposition; it also checked the auxiliary field in the proof of Theorem 1.3, the five editorial corrections of the 191-page edition (one, a wording change, by reading only), the genus-two point counts, the sextic Weil-number test of the CM review, and Remark 16.4. These are targeted checks, not a complete verification.

References

Adleman, L. M., & Lenstra, H. W., Jr. (1986). Finding irreducible polynomials over finite fields. In *Proceedings of the eighteenth annual ACM symposium on theory of computing* (pp. 350–355). ACM. https://doi.org/10.1145/12130.12166

Akiyama, S. (1996). On the pure Jacobi sums. *Acta Arithmetica, 75*(2), 97–104. Checked

author manuscript, Lemma 1, p. 3, https://www.math.tsukuba.ac.jp/~akiyama/papers/Jacobi_priv.pdf.

Alpöge, L. (2026, November 20). *Announcement concerning the Riemann hypothesis* [Author's four-page note]. https://alpo.ge/remark.pdf

Arene, C., Kohel, D., & Ritzenthaler, C. (2012). Complete addition laws on abelian varieties. *LMS Journal of Computation and Mathematics, 15*, 308–316. https://doi.org/10.1112/S1461157012001027

Arthur, J., & Clozel, L. (1989). *Simple algebras, base change, and the advanced theory of the trace formula* (Annals of Mathematics Studies, Vol. 120). Princeton University Press. https://www.claymath.org/library/cw/arthur/pdf/30.pdf.

Bach, E., & Sorenson, J. (1996). Explicit bounds for primes in residue classes. *Mathematics of Computation, 65*(216), 1717–1735. https://doi.org/10.1090/S0025-5718-96-00763-6

Bach, E., von zur Gathen, J., & Lenstra, H. W., Jr. (2001). Factoring polynomials over special finite fields. *Finite Fields and Their Applications, 7*(1), 5–28. https://doi.org/10.1006/ffta.2000.0306.

Balakrishnan, J. S., Ionica, S., Lauter, K., & Vincent, C. (2016). Constructing genus-3 hyperelliptic Jacobians with CM. *LMS Journal of Computation and Mathematics, 19*(Suppl. A), 283–300. https://doi.org/10.1112/S1461157016000322. Author version: https://arxiv.org/pdf/1603.03832v2.

Bartoli, D., Marino, G., Neri, A., & Vicino, L. (2024). Exceptional scattered sequences. *Algebraic Combinatorics, 7*(5), 1405–1431. https://doi.org/10.5802/alco.377.

Beckwith, O., Raum, M., & Richter, O. K. (2024). Imaginary quadratic fields with ℓ-torsion-free class groups and specified split primes. *International Mathematics Research Notices, 2024*(16), 11582–11596. https://doi.org/10.1093/imrn/rnae130

Belabas, K., Diaz y Diaz, F., & Friedman, E. (2008). Small generators of the ideal class group. *Mathematics of Computation, 77*(262), 1185–1197. https://doi.org/10.1090/S0025-5718-07-02003-0. Author version consulted: https://www.math.u-bordeaux.fr/~kbelabas/research/OnBach.pdf.

Bernard, O., & Kučera, R. (2024). A short basis of the Stickelberger ideal of a cyclotomic field. *Mathematics of Computation, 93*(346), 887–909. https://doi.org/10.1090/mcom/3863. Author version consulted: https://arxiv.org/pdf/2109.13329v1, Proposition 5.1, pp. 17–18.

Bhargava, M., Taniguchi, T., & Thorne, F. (2024). Improved error estimates for the Davenport–Heilbronn theorems. *Mathematische Annalen, 389*(4), 3471–3512. https://doi.org/10.1007/s00208-023-02684-w. Checked author manuscript, https://www.math.kobe-u.ac.jp/HOME/tani/dh-improved-public.pdf, Theorem 1.3 and Table 1, pp. 2–3; Theorem 1.4 and Corollary 1.5, p. 4.

Bhargava, M., & Varma, I. (2016). The mean number of 3-torsion elements in the class groups and ideal groups of quadratic orders. *Proceedings of the London Mathematical Society, 112*(2), 235–266. https://doi.org/10.1112/plms/pdv062. Checked author-hosted journal offprint, https://www.math.toronto.edu/~ila/BV.pdf, §5.1, Lemmas 33–34 and Proposition 35, pp. 256–258.

Bhargava, V., Ivanyos, G., Mittal, R., & Saxena, N. (2017). Irreducibility and deterministic r-th root finding over finite fields. In *Proceedings of the 2017 ACM International Symposium on Symbolic and Algebraic Computation*. https://doi.org/10.1145/3087604.3087620 Author manuscript: https://www.cse.iitk.ac.in/users/nitin/papers/irred-and-rth-root.pdf.

Biasse, J.-F., & Erukulangara, M. R. (2023). A proof of the conjectured run time of the Hafner–McCurley class group algorithm. *Advances in Mathematics of Communications, 17*(6), 1388–1405. https://doi.org/10.3934/amc.2021055. Author version consulted: https://par.nsf.gov/servlets/purl/10324557.

Booher, J., Howe, E. W., Sutherland, A. V., & Voloch, J. F. (2026). *Doubly isogenous curves of genus two with a rational action of D_6* [Preprint, arXiv:2402.08853v3]. https://arxiv.org/pdf/2402.08853v3. Conjecture 10.1, p. 39; Theorem 10.4, p. 41.

Boyer, I. (2014). *Variétés abéliennes et jacobiennes de courbes hyperelliptiques, en particulier à multiplication réelle ou complexe* [Doctoral dissertation, Université Paris Diderot–Paris VII]. https://www.normalesup.org/~iboyer/files/TheseIBoyer.pdf.

Bröker, R., Lauter, K., & Sutherland, A. V. (2012). Modular polynomials via isogeny volcanoes. *Mathematics of Computation, 81*(278), 1201–1231. https://doi.org/10.1090/S0025-5718-2011-02508-1. Checked author manuscript: arXiv:1001.0402v2, §3, equations (3) and (5), pp. 4–5; §4.1, pp. 6–7.

Buchmann, J. A., & Lenstra, H. W., Jr. (1994). Approximating rings of integers in number fields. *Journal de Théorie des Nombres de Bordeaux, 6*(2), 221–260. https://doi.org/10.5802/jtnb.113. https://www.numdam.org/item/JTNB_1994__6_2_221_0.pdf.

Buchmann, J., & Shoup, V. (1996). Constructing nonresidues in finite fields and the extended Riemann hypothesis. *Mathematics of Computation, 65*(215), 1311–1326. https://doi.org/10.1090/S0025-5718-96-00751-X. Checked author manuscript, https://www.shoup.net/papers/nonresidues.pdf, pp. 2–5 and §4, Fact 6, p. 9.

Cardona, G., & Quer, J. (2002). *Curves of genus 2 with group of automorphisms isomorphic to D_8 or D_{12}* [Preprint, arXiv:math/0203165v1]. https://arxiv.org/pdf/math/0203165. Section 5, Proposition 5.1, pp. 19–20 of the consulted preprint.

Cilleruelo, J. (2011). The least common multiple of a quadratic sequence. *Compositio Mathematica, 147*(4), 1129–1150. https://doi.org/10.1112/S0010437X10005191. Lemmas 1, 3, and 4 and §2.4, pp. 1131–1134.

Cohen, H., & Pazuki, F. (2009). Elementary 3-descent with a 3-isogeny. *Acta Arithmetica, 140*(4), 369–404. https://doi.org/10.4064/aa140-4-6.

Colmez, P. (1993). Périodes des variétés abéliennes à multiplication complexe. *Annals of Mathematics, 138*(3), 625–683. https://doi.org/10.2307/2946559. Author-hosted scan: https://webusers.imj-prg.fr/~pierre.colmez/artin.pdf.

Cornelissen, G., Hokken, D., & Ringeling, B. (2026). *The asymptotic Mahler measure of Gaussian periods* (arXiv:2507.09303v3), §8.3, Lemma 8.3.1. https://arxiv.org/html/2507.09303v3.

Couveignes, J.-M., & Lercier, R. (2013). Fast construction of irreducible polynomials over finite fields. *Israel Journal of Mathematics, 194*, 77–105. https://doi.org/10.1007/s11856-012-0070-8 Author manuscript, September 30, 2009: https://www.math.u-bordeaux.fr/~jcouveig/publi/couveignes-lercier.pdf.

Cremona, J. E., & Sutherland, A. V. (2024). Computing the endomorphism ring of an elliptic curve over a number field. In *LuCaNT: LMFDB, computation, and number theory* (*Contemporary Mathematics, 796*, pp. 75–102). American Mathematical Society. https://doi.org/10.1090/conm/796/15998. Checked author manuscript dated 30 June 2023, §2.3, pp. 4–5; §2.4, Proposition 4, p. 5.

Cureton, D., & Thorner, J. (2026). *A quasi-Riemann hypothesis for Dedekind zeta functions* (Version 2). Hexagon, 2610.00200v2. https://hexagonmath.org/2610.00200v2. Theorem 1.1, p. 3; Lemma 8.1, pp. 21–22.

Duke, W. (2006). Modular functions and the uniform distribution of CM points. *Mathematische Annalen, 334*(2), 241–252. https://doi.org/10.1007/s00208-005-0706-7. The j-series and discriminant -15 singular modulus are displayed on p. 1 of https://math.ucla.edu/~wdduke/preprints/modud.pdf.

Enge, A. (2009). The complexity of class polynomial computation via floating point approximations. *Mathematics of Computation, 78*(266), 1089–1107. Checked author manuscript arXiv:cs/0601104v2, 25 July 2008, §1, Theorems 1.1–1.2 and heuristic Corollary 1.3, pp. 2–3. https://arxiv.org/abs/cs/0601104.

Evans, R., Lemmermeyer, F., Sun, Z.-H., & Van Veen, M. (2025). Ring class fields and a result of Hasse. *Journal of Number Theory, 266*, 33–61. https://doi.org/10.1016/j.jnt.2024.07.001. Checked April 2024 author manuscript, https://mathweb.ucsd.edu/~revans/hasse8.pdf, Theorems 1.1–1.4, pp. 1–2, and Theorem 1.6, p. 2.

Evans, R., & Van Veen, M. (2026). Cubic residuacity of real quadratic integers. *Journal of Number Theory, 289*, 35–45. https://doi.org/10.1016/j.jnt.2026.05.005. Available online; assigned to the December 2026 issue. Checked November 2025 author manuscript, https://mathweb.ucsd.edu/~revans/cubicresidue4.pdf, Theorem 1.1, p. 2, which explicitly assumes $p \equiv 1 \pmod 3$.

Evdokimov, S. A. (1989). Factoring a solvable polynomial over a finite field and Generalized Riemann Hypothesis [Russian]. *Zapiski Nauchnykh Seminarov LOMI, 176*, 104–117. https://www.mathnet.ru/eng/znsl4535. English translation (1992), *Journal of Soviet Mathematics, 59*(3), 842–849. https://doi.org/10.1007/BF01104107. Hypothesis, original p. 105; Theorem 1.1, pp. 107–108.

Friesen, C., Spearman, B. K., & Williams, K. S. (1986). Another proof of Eisenstein's law of cubic reciprocity and its supplement. *Rocky Mountain Journal of Mathematics, 16*(2), 395–402. https://people.math.carleton.ca/~williams/papers/pdf/141.pdf. Equation (1.1), p. 395; §2, pp. 396–397.

Griffon, R., & Pazuki, F. (2022). Isogenies of elliptic curves over function fields. *International Mathematics Research Notices, 2022*(19), 14697–14740. https://doi.org/10.1093/imrn/rnab033. Author version 2, https://math.richardgriffon.me/articles/GriffonPazuki_Isogenies_V2.pdf, Remark 5.5, p. 13, for the monic modular-polynomial relation and equal support of poles; journal pagination was not used for that locator.

Guo, Z. (2017). *$\mathcal{P}$-schemes and deterministic polynomial factoring over finite fields* [Doctoral dissertation, California Institute of Technology]. Version arXiv:1706.10028v2, §1.1, Hypothesis $(*)$, p. 3. https://arxiv.org/abs/1706.10028. Appendix B locators refer to the Caltech dissertation, pp. 240–243, https://thesis.caltech.edu/10241/1/Guo_Zeyu_2017.pdf.

Hasanalizade, E., Shen, Q., & Wong, P.-J. (2022). Counting zeros of Dedekind zeta functions. *Mathematics of Computation, 91*(333), 277–293. https://doi.org/10.1090/mcom/3665. Author version consulted: arXiv:2102.04663v3, https://arxiv.org/pdf/2102.04663v3.

Hashimoto, Y., & Aoki, M. (2024). Normal integral bases and Gaussian periods in the simplest cubic fields. *Annales mathématiques du Québec, 48*, 157–173. https://doi.org/10.1007/s40316-022-00204-x

Hashimoto, Y., & Nuida, K. (2026). *Ordinary 3-isogeny graphs and improvement of supersingularity testing for twisted Hessian curves over prime fields* [Preprint]. arXiv:2609.11037v1. https://doi.org/10.48550/arXiv.2609.11037.

Hooley, C. (1967). On Artin's conjecture. *Journal für die reine und angewandte Mathematik, 225*, 209–220. Published translation consulted: Hooley, C. (1968). O gipoteze Artina [On Artin's conjecture] (A. I. Vinogradov, Trans.). *Matematika, 12*(1), 37–47. https://www.mathnet.ru/eng/mat455.

Humphries, P., & Brumley, F. (2019). Standard zero-free regions for Rankin–Selberg L-functions via sieve theory. *Mathematische Zeitschrift, 292*, 1105–1122. https://doi.org/10.1007/s00209-018-2136-8. Publisher online-first PDF (Appendix A locators): https://discovery.ucl.ac.uk/id/eprint/10060648/1/Humphries_Brumley2018_Article_StandardZero-freeRegionsForRan.pdf.

Ivanyos, G., Karpinski, M., Rónyai, L., & Saxena, N. (2012). Trading GRH for algebra: Algorithms for factoring polynomials and related structures. *Mathematics of Computation, 81*(277), 493–531. https://doi.org/10.1090/S0025-5718-2011-02505-6. Author pagination consulted: https://www.cse.iitk.ac.in/users/nitin/papers/nogrh.pdf.

Jao, D., Miller, S. D., & Venkatesan, R. (2009). Expander graphs based on GRH with an application to elliptic curve cryptography. *Journal of Number Theory, 129*(6), 1491–1504. https://doi.org/10.1016/j.jnt.2008.11.006. Author version consulted: https://arxiv.org/pdf/0811.0647v2.

Kannan, R., Lenstra, A. K., & Lovász, L. (1988). Polynomial factorization and nonrandomness of bits of algebraic and some transcendental numbers. *Mathematics of Computation, 50*(181), 235–250. https://www.jstor.org/stable/2007927. Original article in the consulted course-reader PDF, pp. 7–22: https://www.math.cmu.edu/users/af1p/Teaching/AdditiveCombinatorics/LLLL.pdf.

Kazhdan, D. A., & Patterson, S. J. (1984). Metaplectic forms. *Publications Mathématiques de l'IHÉS, 59*, 35–142. https://doi.org/10.1007/BF02698770. Theorem I.6.4(c), p. 104; Theorem II.1.4, pp. 114–117; Theorem II.2.1 and the preceding paragraph, p. 118.

Koiran, P. (1996a). Hilbert's Nullstellensatz is in the polynomial hierarchy. *Journal of Complexity, 12*(4), 273–286. https://doi.org/10.1006/jcom.1996.0019.

Koiran, P. (1996b). *Hilbert's Nullstellensatz is in the polynomial hierarchy* (DIMACS Technical Report 96-27, extended version). https://perso.ens-lyon.fr/pascal.koiran/Publis/DIMACS.96-27.ps.

Koiran, P. (1997, August 8). *Erratum to Hilbert's Nullstellensatz is in the polynomial hierarchy.* https://perso.ens-lyon.fr/pascal.koiran/Publis/hn.erratum.ps.

Kolountzakis, M. N., & Révész, S. Gy. (2006). Turán's extremal problem for positive definite functions on groups. *Journal of the London Mathematical Society, 74*(2), 475–496. https://doi.org/10.1112/S0024610706023234. Author version: https://eigen-space.org/mk/ps/turan-groups.pdf.

Kriebel, A. (2026). *Output-sensitive sparse factorization over finite fields from dense factorization* [Preprint]. https://doi.org/10.5281/zenodo.23205034. Author version dated October 6, 2026; Theorem 2, pp. 5–6, Section 4 and Corollary 3, p. 7.

Lamzouri, Y., Li, X., & Soundararajan, K. (2015). Conditional bounds for the least quadratic non-residue and related problems. *Mathematics of Computation, 84*(295), 2391–2412. https://doi.org/10.1090/S0025-5718-2015-02925-1. Corrected author version consulted: arXiv:1309.3595v3, December 2016, https://arxiv.org/pdf/1309.3595v3.

Lemke Oliver, R. J., & Thorner, J. (2019). Effective log-free zero density estimates for automorphic L-functions and the Sato–Tate conjecture. *International Mathematics Research Notices, 2019*(22), 6988–7036. https://doi.org/10.1093/imrn/rnx309. Author version consulted: 19 August 2017, https://lemkeoliver.github.io/papers/15-AutoLFZDE.pdf.

Lenstra, A. K. (1983). Factoring polynomials over algebraic number fields. In J. A. van Hulzen (Ed.), *Computer algebra: EUROCAL '83* (Lecture Notes in Computer Science, Vol. 162, pp. 245–254). Springer. https://ir.cwi.nl/pub/2541/2541D.pdf.

Lenstra, H. W., Jr. (1991). Finding isomorphisms between finite fields. *Mathematics of Computation, 56*(193), 329–347. https://pub.math.leidenuniv.nl/~lenstrahw/PUBLICATIONS/1991b/art.pdf

Lenstra, H. W., Jr., & Silverberg, A. (2019). Testing isomorphism of lattices over CM-orders. *SIAM Journal on Computing, 48*(4), 1300–1334. doi:10.1137/17M115390X. Author manuscript: https://www.math.uci.edu/~asilverb/bibliography/CMorders.pdf.

Miller, G. L. (1976). Riemann's hypothesis and tests for primality. *Journal of Computer and System Sciences, 13*(3), 300–317. https://www.cs.cmu.edu/~glmiller/Publications/Papers/Mi76.pdf.

Milne, J. S. (1992). The points on a Shimura variety modulo a prime of good reduction. In R. P. Langlands & D. Ramakrishnan (Eds.), *The zeta function of Picard modular surfaces* (pp. 151–253). Les Publications CRM. Author manuscript: https://jmilne.org/math/articles/1992aP.pdf.

Milne, J. S. (1999). Lefschetz motives and the Tate conjecture. *Compositio Mathematica*, 117, 45–76. https://doi.org/10.1023/A:1000776613765.

Milne, J. S. (2020, July 14). *Complex multiplication* (Version 0.10) [Author's notes]. https://www.jmilne.org/math/CourseNotes/CM.pdf.

Montgomery, H. L., & Vaughan, R. C. (2007). *Multiplicative number theory I: Classical theory.* Cambridge University Press. https://doi.org/10.1017/CBO9780511618314.

Mumford, D. (1966). On the equations defining abelian varieties. I. *Inventiones Mathematicae, 1*, 287–354. https://doi.org/10.1007/BF01389737

Mumford, D. (1967). On the equations defining abelian varieties. II. *Inventiones Mathematicae, 3*, 75–135. https://doi.org/10.1007/BF01389741

Mumford, D. (1970). *Abelian varieties.* Oxford University Press. Chapter II, Section 6, Application 1, pp. 60–61.

OpenAI. (2026a, October 4). *Deterministic polynomial factorization over prime fields* [Preprint]. https://github.com/openai/math/tree/main/preprints/Deterministic-Polynomial-Factorization-over-Prime-Fields-October-4-2026.

Theorem 1.1, p. 2; Theorem 1.2, p. 3; Proposition 10.1, pp. 41–44; Theorem 11.1, p. 44; Lemma 11.2, p. 45; Proposition 11.3, p. 46.

OpenAI. (2026b, October 4). *Primitive roots for every admissible integer base* [Preprint]. https://raw.githubusercontent.com/openai/math/main/preprints/Primitive-roots-for-every-admissible-integer-base-October-4-2026/primitive-roots-all-integer-bases.pdf.

PARI Group. (2026). *polmodular.c source coverage snapshot: PARI/GP v2.19.0, development 31092-e6893b0017* [Numerical coefficient data]. https://pari.math.u-bordeaux.fr/lcov-report/basemath/polmodular.c.gcov.html. Snapshot dated 21 November 2026; Φ_3 coefficient list before phi3_ZV, source lines 2081–2093.

Pazuki, F. (2012). Theta height and Faltings height. *Bulletin de la Société Mathématique de France, 140*(1), 19–49. https://doi.org/10.24033/bsmf.2623. https://www.numdam.org/item/10.24033/bsmf.2623.pdf.

Pila, J. (1990). Frobenius maps of abelian varieties and finding roots of unity in finite fields. *Mathematics of Computation, 55*(192), 745–763. https://doi.org/10.1090/S0025-5718-1990-1035941-X. Checked author manuscript, https://people.maths.ox.ac.uk/pila/Frobenius.pdf, Theorems A and D, pp. 1–2; §5, pp. 14–18.

Poonen, B. (2019). Using zeta functions to factor polynomials over finite fields. In *Arithmetic geometry: computation and applications* (*Contemporary Mathematics, 722*, pp. 141–147). American Mathematical Society. https://doi.org/10.1090/conm/722/14528. Author version: https://math.mit.edu/~poonen/papers/fq-factoring-ams1.pdf, Hypothesis Z and Theorem 5.1, pp. 4–5.

Rubinstein, M., & Sarnak, P. (1994). Chebyshev's bias. *Experimental Mathematics, 3*(3), 173–197. https://www.math.uwaterloo.ca/~mrubinst/publications/Chebyshev.pdf.

Schoof, R. (1985). Elliptic curves over finite fields and the computation of square roots mod p. *Mathematics of Computation, 44*(170), 483–494. https://doi.org/10.1090/S0025-5718-1985-0777280-6.

Schoof, R. (2008). Four primality testing algorithms. In J. P. Buhler & P. Stevenhagen (Eds.), *Algorithmic number theory: Lattices, number fields, curves and cryptography* (MSRI Publications, Vol. 44, pp. 101–126). Cambridge University Press. https://pub.math.leidenuniv.nl/~stevenhagenp/ANTproc/05rene.pdf.

Serre, J.-P. (1981). Quelques applications du théorème de densité de Chebotarev. *Publications Mathématiques de l'IHÉS, 54*, 123–201. https://doi.org/10.1007/BF02698692.

Serre, J.-P., & Tate, J. (1968). Good reduction of abelian varieties. *Annals of Mathematics, 88*(3), 492–517. https://doi.org/10.2307/1970722.

Rónyai, L. (1992). Galois groups and factoring polynomials over finite fields. *SIAM Journal on Discrete Mathematics, 5*, 345–365. https://doi.org/10.1137/0405026.

Shoup, V. (1990). New algorithms for finding irreducible polynomials over finite fields. *Mathematics of Computation, 54*, 435–447. https://www.shoup.net/papers/detirred.pdf.

Shoup, V. (1992). Searching for primitive roots in finite fields. *Mathematics of Computation, 58*(197), 369–380. https://shoup.net/papers/primroots.pdf.

Silverman, J. H. (2009). *The arithmetic of elliptic curves* (2nd ed.). Graduate Texts in Mathematics, 106. Springer. https://doi.org/10.1007/978-0-387-09494-6.

Silvester, A. K., Spearman, B. K., & Williams, K. S. (2006). Cyclic cubic fields of given conductor and given index. *Canadian Mathematical Bulletin, 49*(3), 472–480. https://people.math.carleton.ca/~williams/papers/pdf/292.pdf. The cyclic-cubic conductor and discriminant facts are on p. 472.

Solovay, R., & Strassen, V. (1977). A fast Monte-Carlo test for primality. *SIAM Journal on Computing, 6*(1), 84–85. https://doi.org/10.1137/0206006. Correction (1978), *7*(1), 118, https://doi.org/10.1137/0207009.

Spaenlehauer, P.-J. (2026). *Deterministic root-finding over finite fields with elliptic curves* (arXiv:2610.32508v1). https://arxiv.org/abs/2610.32508v1. Theorem 1.3, p. 2; Conjecture 1.5 and Problem 1.7, p. 3; Proposition 6.4, p. 35 (printed pagination).

Sutherland, A. V. (2011). Computing Hilbert class polynomials with the Chinese remainder theorem. *Mathematics of Computation, 80*(273), 501–538. https://doi.org/10.1090/S0025-5718-2010-02373-7. Corrected author version consulted: arXiv:0903.2785v4, November 2013, https://arxiv.org/pdf/0903.2785v4.

Sutherland, A. V. (2012). Accelerating the CM method. *LMS Journal of Computation and Mathematics, 15*, 172–204. https://doi.org/10.1112/S1461157012001015. Checked author manuscript arXiv:1009.1082v4, 3 September 2012, §2.5, pp. 5–6; §2.4, pp. 4–5 distinguishes its Dedekind-GRH assumption for the CRT algorithm.

Sutherland, A. V. (2013). Isogeny volcanoes. In E. W. Howe & K. S. Kedlaya (Eds.), *Algorithmic number theory: Tenth international symposium* (The Open Book Series, Vol. 1, pp. 507–530). Mathematical Sciences Publishers. https://doi.org/10.2140/obs.2013.1.507. https://msp.org/obs/2013/1-1/obs-v1-n1-p25-s.pdf.

Sze, T.-W. (2011). *On taking r-th roots without r-th nonresidues over finite fields and its applications* (arXiv:1105.5852v1). https://arxiv.org/abs/1105.5852v1. The arXiv header is dated 30 May 2011; the retrieved version's internal title page says "Preliminary version, November 21, 2018."

Tal, Y. (2026). *Computing modular factorials below the square-root barrier*. Electronic Colloquium on Computational Complexity, Report TR26-211. https://eccc.weizmann.ac.il/report/2026/211/. September 2026 preprint.

Tang, Q. (2026a). *Deterministic polynomial-time factorization over finite fields* (Version 2). Hexagon, 2611.00135v2. https://hexagonmath.org/2611.00135v2. Theorem 1.1 and Remarks 1.2–1.3, p. 2; Lemma 2.1, pp. 2–4; Proposition 3.1, pp. 4–5; Proposition 4.1 and the final proof, pp. 5–7.

Tang, Q. (2026b). *Cureton–Thorner v2 review* [Review repository and separate revised manuscript]. https://github.com/tqyqiqi1509/cureton-thorner-v2-review. Replacement proof of the residual norm estimate in Lemma 8.1.

Thorner, J., & Zaman, A. (2019). A unified and improved Chebotarev density theorem. *Algebra & Number Theory, 13*(5), 1039–1068. https://doi.org/10.2140/ant.2019.13.1039. Publisher version: https://msp.org/ant/2019/13-5/ant-v13-n5-p02-p.pdf.

van Wamelen, P. (2002). Jacobi sums over finite fields. *Acta Arithmetica, 102*(1), 1–20. doi:10.4064/aa102-1-1.

Vélu, J. (1971). Isogénies entre courbes elliptiques. *Comptes rendus de l'Académie des Sciences de Paris, Série A, 273*, A238–A241. English translation by Alexandru Ghitza, *Isogenies between*

elliptic curves, https://aghitza.github.io/publications/translation-velu.pdf, pp. 3–4, equations (8)–(11) and the base-field remark.

von Känel, R., & Kret, A. (2023). *Integral points on coarse Hilbert moduli schemes* [Preprint, arXiv:2307.06944v1]. https://arxiv.org/pdf/2307.06944.

Waterhouse, W. C. (1969). Abelian varieties over finite fields. *Annales scientifiques de l'École Normale Supérieure, 2*(4), 521–560. https://doi.org/10.24033/asens.1183

Wei, X. (2023). Smallness of Faltings heights of CM abelian varieties. *Journal of Number Theory, 242*, 73–101. https://doi.org/10.1016/j.jnt.2022.07.007. Author version consulted: arXiv:2111.00617v1 (2021), https://arxiv.org/pdf/2111.00617v1.

Yu, C.-F. (2016). *A remark on Chevalley's ambiguous class number formulas* (arXiv:1412.1458v2). https://arxiv.org/abs/1412.1458v2. The stable version is dated 18 May 2016; its internal manuscript also displays 18 October 2021.

Zaman, A. (2018). The least unramified prime which does not split completely. *Forum Mathematicum, 30*(3), 651–661. https://doi.org/10.1515/forum-2017-0081. Author version consulted: arXiv:1704.03451v1, https://arxiv.org/pdf/1704.03451v1.