ECC RESEARCH PROGRAM · P4.3

The exTNFS Complexity Is Still Heuristic: A Complete Smoothness-Dependency Audit, a Boundary Relation-Supply Theorem, and the Degree-Uniformity Barrier

What is and is not proved about the extended tower number field sieve in medium characteristic — a negative report with retained partial theorems

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

with GPT 5.6 Sol · Claude Fable 5

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

ABSTRACT

Problem P4.3 asked for an unconditional $L_{p^n}(1/3, c)$ complexity theorem for the extended tower number field sieve (exTNFS) of Kim–Barbulescu in medium characteristic with composite extension degree — possibly for a modified algorithm, possibly with a worse constant. **This target was not achieved, and the task record closes it as failed.** This paper reports, with the record's own epistemic tags, exactly what stands between the community and such a theorem. We decompose the exTNFS pipeline into ten numbered analytic inputs S-01–S-10, each stated as a precise uniform density claim about the actual iterated-resultant norm forms; two are theorems (rigorous smooth-candidate factorization, and the Canfield–Erdős–Pomerance random-integer benchmark), and eight are open — joined by two adjacent non-smoothness gaps, polynomial-selection irreducibility and relation-matrix rank. We prove that the minimal missing relation statement is a **joint** two-sided smoothness lower bound (RC) that marginal densities do not imply; that the minimal missing rank statement is a hyperplane-escape bound (R2) that smoothness counts do not imply; and a **degree-uniformity barrier**: keeping relation norms on the optimized $L_Q(2/3)$ scale in the strict medium-characteristic interior forces the tower degree $\eta \to \infty$, outside the quantifiers of every fixed-form smooth-value theorem in the audited literature. On the boundary $\ell_p = 2/3$ we retain a genuine partial positive: an unconditional, explicitly randomized $L_Q(1/3, O(1))$ **relation-supply** theorem for the restricted family $p \equiv 3 \bmod 4$, $\eta = 2$, $R = \mathbb{Z}[i]$, with an explicit nonoptimized constant — which provably does not extend to the interior by the same proof, and which is **not** a DLP theorem: target splitting, special-q descent, and rank remain open, and we prove structural barriers (a subspace-counting bound and a resultant-bidegree identity) explaining why the boundary construction does not linearize either remaining stage. An exact toy experiment ($p = 5$, $\eta = 2$, $\kappa = 3$, all 5,856 primitive coefficient vectors factored completely) falsifies the literal random-integer model at finite scale: joint smoothness exceeds the matched dyadic baseline by factors 8.50–3.05. The honest bottom line: today's fully rigorous alternative remains Bender–Pomerance at $L_Q\left(1/2, \sqrt{2}\right)$-type costs, and every path to $1/3$ runs through the specific open statements catalogued here.

Keywords. discrete logarithm · exTNFS · number field sieve · smooth numbers · rigorous analysis · medium characteristic · L-notation

1 Introduction

The extended tower number field sieve computes discrete logarithms in $\mathbb{F}_{p^n}^*$ for composite $n = \eta\kappa$ by working in a tower $R = \mathbb{Z}[t]/(h)$ of degree η over which two polynomials f, g share an irreducible factor of degree κ modulo p; its announced complexities, down to $L_Q\big(1/3, (48/9)^{1/3}\big)$ and $L_Q\big(1/3, (32/9)^{1/3}\big)$ for the best constructions, made every medium-characteristic pairing field re-evaluated after 2016 [1]. Those complexities are theorems **under classical NFS heuristics**: each structured norm value is modelled as a uniformly random integer of the same size, two correlated norms are modelled as independent, and the descent tree is modelled by Dickman probabilities. Problem P4.3 asked whether this can be made unconditional — for exTNFS itself, or for an identified modification, possibly with a worse constant.

Outcome, stated plainly. PROVED The formal objective **failed**: no unconditional $L_Q(1/3, c)$ complexity theorem was proved for exTNFS or any complete modified DLP algorithm, and the research record closes the task as abandoned at the operator's direction. What the record retains — and what this paper reports — is (i) a complete ten-input dependency audit with every smoothness assertion stated precisely against the actual tower norm forms (§3); (ii) proofs that the minimal missing statements are a joint smoothness bound (RC), an adaptive descent bound (SQ), and a rank-escape bound (R2), none implied by the others (§3–4); (iii) a degree-uniformity barrier showing fixed-degree smooth-form theorems cannot reach the optimized interior parameter range (§5); (iv) an unconditional boundary relation-supply theorem at $\ell_p = 2/3$, $\eta = 2$ with an explicit constant (§6); (v) proved structural barriers for target splitting and special-q descent (§7); and (vi) an exact toy falsification of the literal random-integer model (§8).

We use the record's epistemic tags throughout: PROVED, CITED, CONJECTURE for precisely stated open inputs, CONDITIONAL: GRH where relevant, and EMPIRICAL: scope for finite computations. A negative report has one job — to make the obstruction exact — and we have tried to write every open statement so that a future proof (or refutation) of it is well-posed.

2 The algorithm being audited

2.1 Notation

Write

$$L_Q(\alpha, c) = \exp\big((c + o(1))(\log Q)^{\alpha}(\log\log Q)^{1-\alpha}\big), \quad Q = p^n, \quad n = \eta\kappa, \quad \gcd(\eta, \kappa) = 1, \tag{1}$$

and parameterize medium characteristic by $p = L_Q\big(\ell_p, c_p\big)$ with $1/3 < \ell_p < 2/3$. exTNFS chooses $h \in \mathbb{Z}[t]$ of degree η irreducible modulo p, works in $R = \mathbb{Z}[t]/(h)$, and chooses $f, g \in R[x]$ whose reductions modulo p share an irreducible degree-κ factor $k(x)$ over $R/pR \cong \mathbb{F}_{p^\eta}$, so both tower number fields map onto $\mathbb{F}_Q$ [1].

A relation candidate is $r = a(\iota) - b(\iota)x$ with $a(t) = \sum_{i<\eta} a_i t^i$, $b(t) = \sum_{i<\eta} b_i t^i$, $|a_i|, |b_i| \le A$, and its two integer norms are the iterated resultant forms in the 2η variables $\boldsymbol{z} = \big(a_0, \ldots, a_{\eta-1}, b_0, \ldots, b_{\eta-1}\big)$:

$$F_f(z) = |\mathrm{Res}_t(\mathrm{Res}_x(a(t) - b(t)x, f(x)), h(t))|, \quad F_g(z) = |\mathrm{Res}_t(\mathrm{Res}_x(a(t) - b(t)x, g(x)), h(t))|. \tag{2}$$

A row is kept when both $F_f(z)$ and $F_g(z)$ are B-smooth with $B = L_Q(1/3, \beta)$; accepted rows are factored over prime-ideal factor bases, extended by Schirokauer coordinates, and the sparse system is solved modulo the large prime $\ell \mid Q - 1$ [1], [2].

2.2 The heuristic cost calculation

CITED With optimized selection the two norms obey $F_f \le L_Q\big(2/3, \gamma_f + o(1)\big)$, $F_g \le L_Q\big(2/3, \gamma_g + o(1)\big)$, and Kim–Barbulescu write relation collection plus linear algebra as

$$L_Q\left(1/3,\, \beta + \frac{\gamma_f + \gamma_g}{3\beta} + o(1)\right) + L_Q(1/3,\, 2\beta + o(1)), \tag{3}$$

modelling each norm as an arbitrary integer of its size and multiplying the two probabilities [1].

Lemma 1 (balancing). PROVED Equating the two displayed exponents gives $\beta^2 = (\gamma_f + \gamma_g)/3$ and total constant $2\sqrt{(\gamma_f + \gamma_g)/3}$.

Every appearance of "the probability of an arbitrary integer" above is an unproved transfer; making that exact is the audit of the next section.

3 The ten smoothness inputs

Let $\mathcal{S}_Q(A)$ be the nonzero coefficient vectors with both norms nonzero and $P^+(m)$ the largest prime factor of $|m|$. We state the inputs S-01–S-10 of the record's dependency audit; the full formal versions are in the research file `SMOOTHNESS_ASSUMPTIONS.md`.

S-01/S-02 (marginal relation densities). CONJECTURE Uniformly along every admissible parameter sequence actually output by the selection algorithm,

$$\frac{\#\{z \in \mathcal{S}_Q(A) : P^+(F_s(z)) \le B\}}{\#\mathcal{S}_Q(A)} = L_Q\left(1/3, -\frac{\gamma_s}{3\beta} + o(1)\right), \quad s \in \{f, g\}. \tag{4}$$

S-03 (joint density — the minimal relation statement). CONJECTURE The single statement sufficient for the upper bound is the uniform **lower** bound

$$\frac{\#\{z \in \mathcal{S}_Q(A) : P^+(F_f(z)F_g(z)) \le B\}}{\#\mathcal{S}_Q(A)} \ge L_Q\left(1/3, -\frac{\gamma_f + \gamma_g}{3\beta} + o(1)\right), \tag{5}$$

paired with the deterministic count $\#\mathcal{S}_Q(A) = L_Q\big(1/3, \beta + (\gamma_f + \gamma_g)/(3\beta) + o(1)\big)$. We call this (RC).

Proposition 2 (marginals do not imply (RC)). PROVED S-01 and S-02 do not imply S-03: both events are functions of the **same** coefficient vector and may be correlated at the full L-exponent scale. Multiplying marginal densities is an additional assumption, not algebra. (§8 exhibits positive dependence exactly at toy scale.)

S-04/S-05 (initial target splitting). CONJECTURE For a target s and randomizing exponents $e \le E$, the JLSV2 lift norm $M_e \le X_0$ (or the Waterloo numerator/denominator pair $U_e V_e$, $|U_e V_e| \le Q^{1+o(1)}$) must contain squarefree Y_0-smooth values at random-integer frequency:

$$\#\{e \le E : P^+(M_e) \le Y_0,\, \mu^2(M_e) = 1\} \ge E \cdot \frac{\Psi(X_0, Y_0)}{X_0} \cdot L_Q(1/3, o(1)), \tag{6}$$

uniformly over all allowed targets — a smooth-value claim about subgroup-indexed structured lifts, not about random integers.

S-06 (squarefreeness). CONJECTURE The squarefree co-condition needs a positive uniform proportion **after conditioning on smoothness**; a smooth-values theorem alone does not discharge it.

S-07 (one special-q step). CONJECTURE For a prime ideal $\mathfrak{q}$ of norm ν with enumeration set $\mathcal{C}_{\mathfrak{q}}(D)$ (bounded combinations of an LLL basis of its lattice) and child bound $y = \nu^c$, with cofactor sizes $u_f = \log X_{f,\mathfrak{q}}/\log y$, $u_g = \log X_{g,\mathfrak{q}}/\log y$:

$$\frac{\#\left\{z \in \mathcal{C}_{\mathfrak{q}}(D) : P^+\left(G_{f,\mathfrak{q}}(z)\, F_g(z)\right) \le y\right\}}{\#\mathcal{C}_{\mathfrak{q}}(D)} \ge \rho_D(u_f)\, \rho_D(u_g)\, L_Q(1/3, o(1)), \tag{7}$$

where $G_{f,\mathfrak{q}} = F_f/\nu$ and ρ_D is Dickman's function. We call this (SQ).

S-08 (adaptive uniformity). CONJECTURE (SQ) must hold — or fail with summable probability — simultaneously over every $\mathfrak{q}$, every LLL basis, and every child selected by the random descent history, with the accumulated trial count inside the $L_Q(1/3, c + o(1))$ budget. A pointwise statement for one fixed lattice is insufficient.

S-09 (smoothness testing — known). CITED Hyperelliptic smoothness testing factors a y-smooth $m \le x$ in expected $y^{o(1)}$ time when $y = (\log x)^{\omega(1)}$ [3], [4]. PROVED The exTNFS regime satisfies the hypothesis ($x \le L_Q(2/3, O(1))$, $y = B = L_Q(1/3, \beta)$), so **testing cost is not the obstruction** — supply is.

S-10 (the benchmark — known). CITED Canfield–Erdős–Pomerance give, uniformly in their range, the density $\Psi\big(L_Q(b,d), L_Q(a,c)\big)/L_Q(b,d) = L_Q(b-a, -d(b-a)/c + o(1))$ [5]. PROVED This counts **integers**; invoking it after calling a norm "random" is precisely the unproved transfer isolated in S-01–S-08.

3.1 Status against the strongest known partial results

Table 1: The audit table. CONDITIONAL: GRH Buchmann–Hollinger's smooth-ideal lower bound [6] does not change the picture: even under GRH it counts ideals, not short principal generators satisfying a second-side condition or a special-q constraint.

Input	Closest unconditional result	What is missing
S-01–S-02	fixed binary forms: positive-proportion friable values [7]; fixed quadratic form and fixed-field smooth ideals [8]	parameter-dependent 2η-variable iterated-resultant forms; the far sparser $L(1/3)$ range; uniformity as fields, degrees, discriminants vary
S-03	products of finitely many **fixed** forms in dense ranges [8]	joint L-scale lower bound for two **varying** tower forms
S-04–S-06	rigorous index calculus randomizes integers/polynomials with countable distributions [9], [10]	control of subgroup-indexed tower lift norms, jointly smooth and squarefree, uniformly over targets
S-07–S-08	Dickman-type smooth-ideal counts in one **fixed** field [8]	short principal generators in varying special-q lattices, paired second side, adaptive uniformity
S-09	hyperelliptic smoothness test [3]	nothing — this input is closed
S-10	CEP density [5]	nothing — but it is only the benchmark

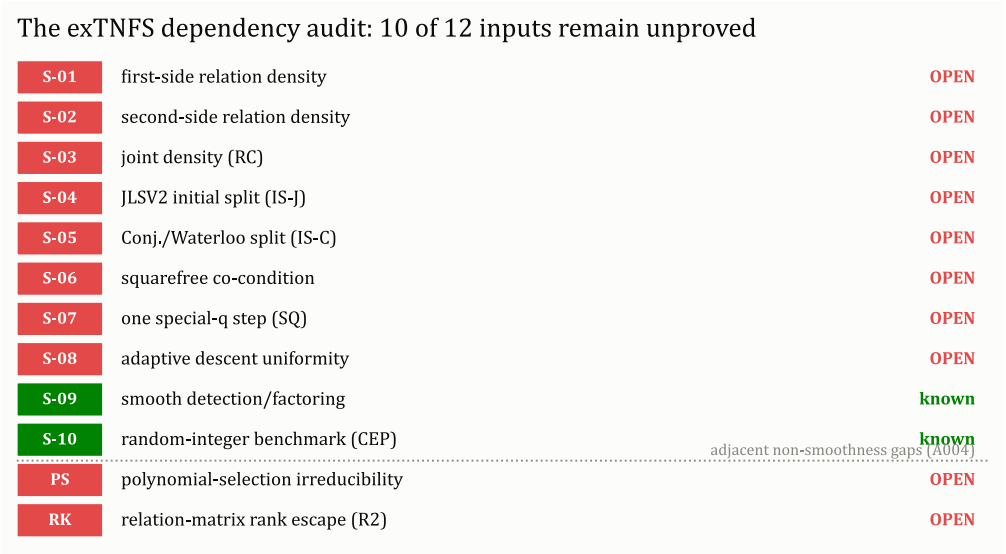

Figure 1: The dependency audit at a glance: of the ten smoothness inputs plus the two adjacent gaps of §4, exactly two are theorems — the cost of factoring accepted candidates (S-09) and the random-integer benchmark (S-10). Every supply statement is open.

4 The two adjacent non-smoothness gaps

Even a full proof of (RC), (IS), and (SQ) would not prove the DLP complexity. Two further inputs are unproved and are **not** smoothness statements.

4.1 Polynomial-selection irreducibility

CITED The Conjugation selection expects suitable irreducible polynomials after a small number of trials; the source labels this step heuristic [1]. An unconditional proof must construct the sequence uniformly or account for the search rigorously.

4.2 Relation-matrix rank: the minimal escape condition

PROVED Counting accepted relations is not enough: all $L_Q(1/3,\beta)$ rows could lie in a proper subspace of the admissible row space H_Q (the annihilator of the true virtual-logarithm functionals [2]). The record's attempt A004 isolates the **minimal** sufficient rank input.

Theorem 3 (rank escape suffices and preserves the constant). PROVED Let μ_Q be the distribution of the valuation/Schirokauer row R of an accepted relation, and define

$$\delta_Q = \inf_{\varphi \in H_Q^* \setminus \{0\}} \Pr_{R\sim\mu_Q}[\varphi(R) \neq 0]. \tag{8}$$

If $\delta_Q \geq L_Q(1/3, -o(1))$, then adaptively retaining independent rows spans H_Q after at most $\dim(H_Q)/\delta_Q$ accepted rows in expectation; since $\dim H_Q = L_Q(1/3,\beta+o(1))$, the leading $L(1/3)$ constant is unchanged.

Proof. If the current span $W \subsetneq H_Q$ is proper, pick $\varphi \neq 0$ vanishing on W; the event $\varphi(R) \neq 0$ forces $R \notin W$, so the expected number of accepted rows per dimension increase is at most $1/\delta_Q$; sum over at most $\dim H_Q$ increases. □

PROVED For **uniform** rows in an r-dimensional space, $r+s$ rows fail to span with probability below $\ell^{-s}/(\ell-1)$ — but exTNFS rows are sparse, share (a,b) across their two sides, and are conditioned on simultaneous smoothness, so the uniform benchmark is exactly that: a benchmark. The usable analytic form of the escape condition is a character-cancellation bound over the **smooth-conditioned** candidate set — i.e., it is logically downstream of the open (RC), and no audited fixed-form theorem supplies the required equidistribution. The record labels this condition (R2); it is open. The deterministic prerequisites (no empty factor-base column, connected support hypergraph, symmetries quotiented into H_Q, duplicate rows discarded) are catalogued as proved checklist items.

5 The degree-uniformity barrier

Could one instead **prove** smoothness for a restricted tower — say quadratic $\eta = 2$, where the outer norm is a binary quadratic form covered by Barbulescu–Lachand's fixed-form theorem [8]? The record's attempt A003 answers with exact identities and a scaling obstruction.

Proposition 4 (exact quadratic norm identities). PROVED For $h(T) = T^2 + h_1 T + h_0$ with root ι and $\Delta_h = h_1^2 - 4h_0$: $N(x+y\iota) = x^2 - h_1xy + h_0y^2$; for $z_i = x_i + y_i\iota$ the pullback $Q(U,V) = N(Uz_0 + Vz_1) = AU^2 + BUV + CV^2$ has

$$B^2 - 4AC = \Delta_h(x_0y_1 - x_1y_0)^2, \tag{9}$$

and the inhomogeneous slice $N(z_0 + Uz_1 + Vz_2)$ has homogeneous part of discriminant $\Delta_h(x_1y_2 - x_2y_1)^2$, constant term $N(z_0)$, and linear coefficients $\mathrm{Tr}(z_0\overline{z_1})$, $\mathrm{Tr}(z_0\overline{z_2})$.

PROVED Auditing Theorem 4.2 of [8] against these slices: the one-parameter affine slice is a single row of the averaged form (no lower bound follows); the two-parameter affine slice is a lattice

coset, outside the theorem; only the homogeneous slice matches, and then only when the square lattice-index factor in the discriminant disappears — after which the removed fixed divisor must itself be B-smooth. A genuine algebraic opening survives:

Proposition 5 (kernel randomization is surjective). PROVED Randomizing the selected polynomial inside the kernel of reduction, $f_{U,V} = f_0 + pU + KV$ with K a lift of the common factor k, shifts the relative resultant by the ideal $\left(pb^{\kappa}, K^{[\kappa]}(a,b)\right)$, which equals the whole ring R whenever $d = \kappa$, K is monic, and $(a,b) = R$ with (a,b) nonzero modulo p. The natural coefficient randomization is thus **algebraically** surjective onto the quadratic tower order for ideal-coprime candidates.

But surjectivity is not a distribution theorem, and the decisive obstruction is quantitative:

Theorem 6 (degree-uniformity barrier). PROVED Let $p = L_Q\left(\ell_p, c_p\right)$ with $1/3 < \ell_p < 2/3$ and let the relation search have $L_Q(1/3, \tau)$ candidates. The iterated resultant is homogeneous of degree $\eta d \geq n$ in the coefficients of (a,b), so the standard box-wide norm bound carries an A^n term with

$$\log(A^n) = \frac{\tau + o(1)}{2c_p\eta}\,(\log Q)^{4/3-\ell_p}(\log\log Q)^{\ell_p-1/3}. \tag{10}$$

For **fixed** η this is an $L_Q\left(4/3-\ell_p, O(1)\right)$ scale — strictly above $L_Q(2/3)$ throughout the interior — and keeping the norm on the $L_Q(2/3)$ scale forces

$$\eta = \Theta\left((\log Q/\log\log Q)^{2/3-\ell_p}\right) \to \infty, \quad \kappa = \Theta\left((\log Q/\log\log Q)^{1/3}\right). \tag{11}$$

Hence every fixed-degree, fixed-form smooth-value theorem misses the optimized interior regime: the theorem's quantifiers and the algorithm's parameter range are incompatible.

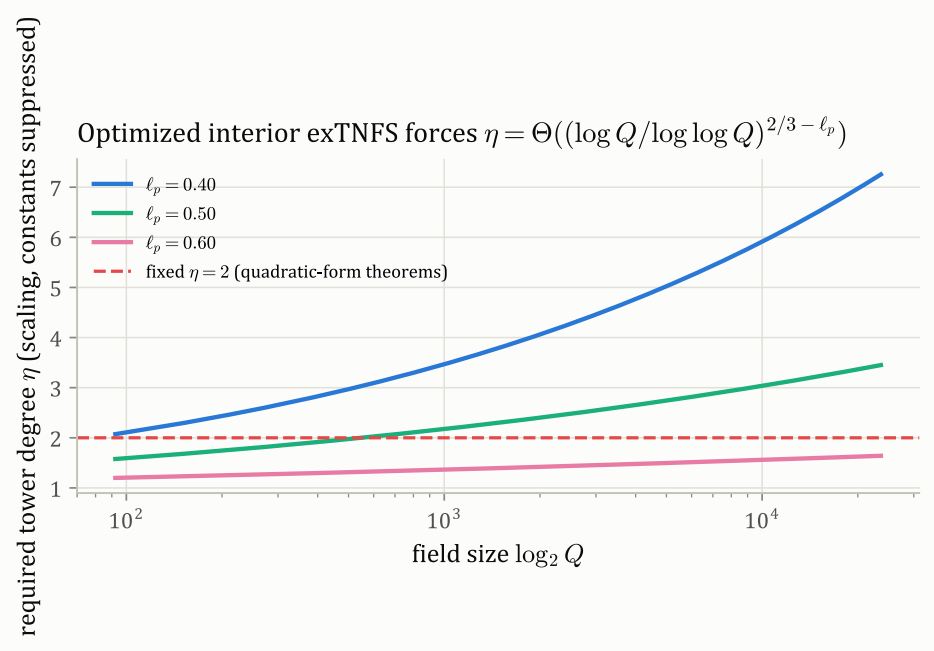

Figure 2: PROVED The barrier of Theorem 6, drawn as pure scaling (constants suppressed): the tower degree η required to keep relation norms on the optimized $L_Q(2/3)$ scale grows without bound for every fixed interior ℓ_p, while every audited smooth-form theorem lives on the dashed fixed-η line. Only the boundary $\ell_p = 2/3$ (exponent 0) escapes.

6 What can be proved on the boundary: a relation-supply theorem

At the boundary $\ell_p = 2/3$ the barrier exponent vanishes, and the record's attempt A005 converts the algebraic opening into an unconditional — but sharply restricted — theorem. Its ingredients, each proved in the record: a bounded-fiber lemma for the primitive map $(u, v) \mapsto Au + Bv$ on a fixed imaginary quadratic order (Bézout solution plus kernel-lattice reduction); a count of monic irreducibles with prescribed nonzero evaluation over a residue field,

$$I_d(r, c) \geq \frac{1}{d}\left(\frac{s^d - 1}{s - 1} - ds^{d/2}\right) = \Theta_s\left(s^{d-1}/d\right), \tag{12}$$

via norm fibers and subfield exclusion; residue-class equidistribution of the kernel fibers; and the fixed-form smoothness input for the Gaussian norm form of discriminant -4 [8] with CEP density [5].

Theorem 7 (boundary relation supply, unconditional). PROVED Consider inputs $Q = p^n$ with $p = L_Q(2/3, c_p)$, $p \equiv 3 \bmod 4$, $n = 2\kappa$, tower $R = \mathbb{Z}[i]$ (so $h = T^2 + 1$ stays irreducible modulo p). Lift a random irreducible degree-κ factor to $K \in R[x]$ and randomize both sides by the kernel families

$$f = K(1 + v_f) + pU_f, \quad g = K(1 + v_g) + pU_g. \tag{13}$$

Put $A = L_Q(1/3, a)$, $B = L_Q(1/3, \beta)$, and $m = c_p + a/(2c_p)$. For any β with $\beta c_p > 1/3$ and any

$$a > \frac{\beta + \frac{8c_p}{3\beta}}{4 - \frac{4}{3\beta c_p}}, \tag{14}$$

the randomized collector finds an irreducible pair (f, g) together with at least $L_Q(1/3, \beta)$ jointly B-smooth relation candidates in expected time

$$L_Q\left(1/3,\, 4a + \frac{8m}{3\beta} + o(1)\right). \tag{15}$$

The constant is explicit and **nonoptimized**; the theorem's value is existence, not efficiency. Its proof multiplies the counts of admissible f- and g-coefficients for each **fixed** candidate — so it never assumes independence of the two norm values — and finishes with a Markov/averaging argument over polynomial pairs.

EMPIRICAL: 161 finite-field cases; 121 quadratic-order targets The two computational lemmas were validated exactly: Möbius/norm-fiber counts satisfied the irreducible-count lower bound in all 161 tested instances (fields $s \in \{2, 3, 5, 7, 11\}$, degrees $2 \le d \le 8$, minimum slack 1.5), and the toy Bézout identity $5(-1) + (1 - \iota)(2 + 2\iota) = 1$ in $\mathbb{Z}[\iota]/(\iota^2 + 2)$ represented all 121 targets in $[-5, 5]^2$ with coefficients bounded by 30.

Exact scope — read before citing. PROVED Theorem 7 is a **relation-supply** theorem for one boundary subfamily. It does not prove: (1) that accepted rows span the relation lattice (the (R2) gap of §4); (2) that arbitrary targets split into covered objects (S-04–S-06); (3) that special-q descent succeeds with uniform probability and cost (S-07–S-08); (4) anything for fixed interior $\ell_p < 2/3$, where Theorem 6 forces $\eta \to \infty$ and the only smooth-value input used — one fixed quadratic form — is unavailable. Do not promote it to an exTNFS or DLP complexity theorem.

7 Why the remaining stages do not linearize

Two proved barriers explain why the boundary construction stalls exactly at target splitting and descent.

Proposition 8 (targets do not land in low-degree subspaces). PROVED Write $\mathbb{F}_Q = \mathbb{F}_{p^\eta}[\theta]/(k(\theta))$ and let $W_m = \left\{\sum_{j \le m} c_j \theta^j\right\}$, so $\#W_m = p^{\eta(m+1)}$. If s generates a subgroup of order ℓ and e is uniform modulo ℓ, then

$$\Pr[s^e \in W_m] \le \frac{p^{\eta(m+1)}}{\ell}, \tag{16}$$

which for a full-size subgroup ($\ell = Q^{1-o(1)}$) is $p^{-\eta(\kappa - m - 1) + o(n)}$. Dropping even **one** tower coefficient from a full-degree representative costs a factor $p^{-\eta}$, whose inverse has logarithm $\Theta\big((\log Q)^{2/3}(\log\log Q)^{1/3}\big)$ — an $L_Q(2/3, O(1))$ wait, far beyond any $L_Q(1/3)$ budget. Exponent randomization therefore cannot make a generic full-group target linear in θ; a nonlinear representation or a genuinely new theorem about full-degree lift norms is required. (For much smaller subgroups the bound is vacuous, and the needed subgroup–subspace intersection statement is open.)

Proposition 9 (resultant bidegree barrier). PROVED $\mathrm{Res}_x(P, f)$ is homogeneous of degree $\deg f$ in the coefficients of P and of degree $\deg P$ in the coefficients of f. The A005 mechanism works **only** because a relation polynomial $a - bx$ has degree one, making the resultant **affine** in the randomized coefficients of f — indeed $\mathrm{Res}_x(a - bx, f) = b^{\deg f} f(a/b)$. A generic target lift of degree $\kappa - 1$ makes the randomized outer form have total degree $\eta\kappa = n + O(\eta)$: exactly the growing-degree regime with no known smooth-value theorem. Factoring the lift as $L \cdot S$ with only a linear factor randomized just relocates the problem into the fixed factor S, whose norm must **also** be smooth.

PROVED Special-q descent inherits both horns of the dilemma: re-randomizing the polynomial changes the number field and destroys the identification of $\mathfrak{q}$, its lattice, and every factor-base virtual logarithm (cross-field relations are themselves target-lifting problems); keeping the field fixed removes the A005 randomness and returns exactly to the open (SQ). These are barriers to **this** modification family, not impossibility theorems — the record states them as such.

8 The exact toy experiment

The record's SG-04 experiment instantiates a genuine exTNFS tower — not a proxy form — at the smallest interesting size:

$$p = 5, \quad h = t^2 + 2, \quad f = x^3 + x + 1, \quad g = x^3 + x - 4, \quad Q = 5^6, \tag{17}$$

where h is irreducible modulo 5 and $f \equiv g \bmod 5$ is an irreducible cubic ($\eta = 2, \kappa = 3$). PROVED For $a = a_0 + a_1\iota, b = b_0 + b_1\iota$ with $\iota^2 = -2$, the inner resultant for $x^3 + x + c$ is $a^3 + ab^2 + cb^3$ and the outer norm of $u_0 + u_1\iota$ is $u_0^2 + 2u_1^2$. The program exhausts all 5,856 primitive vectors in $[-4, 4]^4$, **fully factors** both norms of every vector by deterministic trial division, and draws one independent uniform integer from each norm's dyadic interval as a size-matched baseline (seed 4304). Every stored factorization (23,424 of them) and 128 nested resultants were reproduced by independent SymPy checks.

Table 2: EMPIRICAL: p=5, eta=2, kappa=3, A=4, exhaustive Joint B-smoothness of the two tower norms versus the matched random-integer baseline. The preregistered divergence criterion (outside the baseline's 95% interval by a factor ≥ 1.5 at two adjacent bounds) is met at all four bounds, and the preregistered positive-dependence prediction is not refuted at any of them.

B	actual joint rate	baseline rate (95% Wilson)	ratio	joint / marginals
7	0.023224	0.002732 (0.001683–0.004434)	8.50	7.73
13	0.043716	0.005635 (0.004016–0.007903)	7.76	5.14
31	0.073087	0.019980 (0.016698–0.023891)	3.66	3.81
61	0.114071	0.037398 (0.032833–0.042568)	3.05	2.82

Figure 3: EMPIRICAL: 5,856 primitive vectors, seed 4304 Actual joint smooth rates (exact population values) against the dyadic baseline with Wilson intervals, log scale. Annotations give the actual/baseline ratio. Data: `measure_norm_smoothness_p5_A4_B7-13-31-61_all_s4304_20260702_summary.csv`.

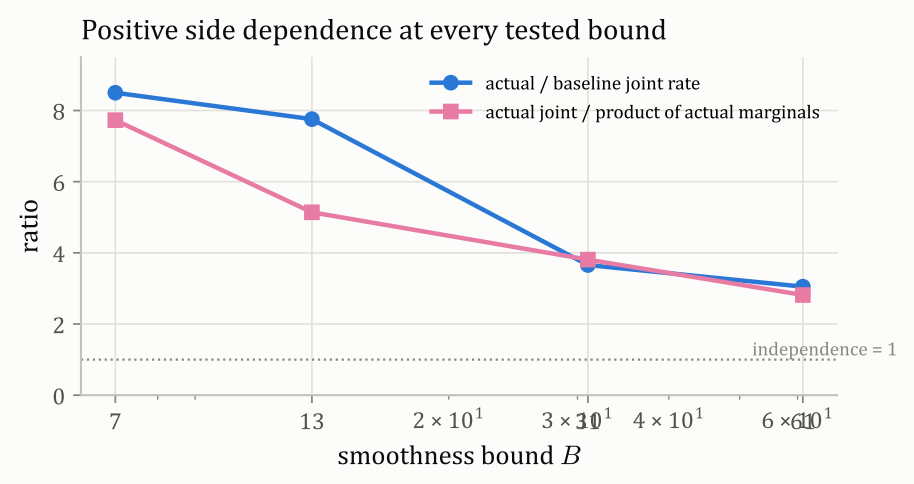

Figure 4: EMPIRICAL: same run Two dependence diagnostics versus B: the actual/baseline joint-rate ratio and the actual joint rate divided by the product of the actual marginals. Both decrease toward — but stay well above — the independence line at these parameters.

PROVED The honest reading: this finite example **falsifies literal equality** with the chosen random-integer model at these parameters and shows the two sides are positively dependent (they share the term $a^3 + ab^2$), but it neither supports nor refutes the asymptotic L-exponent. It is evidence about model structure, not about complexity.

9 Rigorous alternatives, and their price

CITED Bender–Pomerance give a rigorous index-calculus algorithm for **every** finite field via smooth polynomials, with growing-extension costs $L_Q\left(1/2, \sqrt{2} + o(1)\right)$ for $p \le n^{o(n)}$, an intermediate $L_Q\left(1/2, \sqrt{8/3} + o(1)\right)$ range, and $p^{2+o(1)}$ for $p > n^{4n/3}$ [10], building on the randomization precedents of [9]. PROVED In medium-characteristic notation this fallback costs

$L_Q\big(1/2, \sqrt{2} + o(1)\big)$ for $\ell_p < 1/2$ and $L_Q\big(\ell_p, 2c_p + o(1)\big)$ for $\ell_p > 1/2$ — an unconditional theorem for the same DLP, with exponent $1/2$ or ℓ_p in place of $1/3$. It is the cleanest currently identified "modified algorithm with worse complexity", and it is not an analysis of exTNFS.

CITED Lee–Venkatesan's randomized ordinary NFS is the closest rigorous relative at exponent $1/3$: randomizing $f = \hat{f} + (x - my)R$ makes evaluations vary in an arithmetic progression, where smooth-number theorems apply (their factoring conclusion still cites a separate character conjecture) [4]. PROVED The literal tower analogue fails: a kernel-randomized relative resultant varies in $R = \mathbb{Z}[t]/(h)$, and the outer norm $N_{R/\mathbb{Q}}$ is a **nonlinear** degree-η form of it — not an integer progression — so the progression theorem proves none of S-01–S-03 and gives no special-q statement. Sampling smooth **ideals** first [8] is a one-side idea only: the same generator must still be smooth on the other side, so (RC) and (SQ) survive intact.

10 Why the formal target failed

PROVED Assembling the audit: the smallest missing relation theorem is (RC) — not two marginal Dickman statements — and the strongest missing theorem is (SQ) under S-08′s adaptive uniformity; even both, plus S-04–S-06, would leave polynomial-selection accounting and the rank bound (R2). The boundary theorem (Theorem 7) closes relation supply for one restricted family and provably does not extend inward (Theorem 6). No statement in the audited literature has the exTNFS quantifiers. The record therefore closes P4.3 with its target unmet, and lists the resume conditions verbatim: a uniform lower bound for smooth values of growing-degree tower norm forms; or character cancellation / hyperplane escape after conditioning on simultaneous smoothness; or a provable target-splitting distribution compatible with the restricted generator; or a rigorous special-q descent whose degrees and coefficient sizes stay within the claimed complexity. Any one of these would reopen the problem with a real chance of progress.

11 Conclusion

The extended tower number field sieve's $L_Q(1/3, c)$ complexity remains, today, a heuristic statement. This report replaces the folklore version of that sentence with an exact one: the heuristic content is precisely S-01 through S-08 plus polynomial-selection accounting and the rank escape (R2); the two rigorous ingredients (S-09, S-10) are cost and benchmark, not supply; fixed-form theorems are structurally blocked from the optimized interior by the degree-uniformity barrier; the boundary family $\ell_p = 2/3$, $\eta = 2$ admits an unconditional relation-supply theorem with explicit constants; and the toy data show the random-integer model is already false in the small, in the direction (positive dependence) that at least does not starve relation collection. We did not prove the theorem we wanted. What a future proof must contain is now, we believe, stated precisely enough to be attacked.

Reproducibility

The full formal statements of S-01–S-10, the status table, and the modified algorithm audits are the research file `SMOOTHNESS_ASSUMPTIONS.md`; the partial theorems are attempts A003 (quadratic identities, kernel surjectivity, degree barrier), A004 (rank escape), A005 (boundary relation-supply theorem and its four lemmas), A006 (subspace counting), and A007 (resultant bidegree), all in `attempts/`. The toy experiment is `code/`

`measure_norm_smoothness.py` (deterministic, seed 4304, exhaustive $A = 4$ box; independent SymPy replay of all 23,424 factorizations and 128 nested resultants), with dated raw and summary CSVs in `data/`; the figures here are regenerated from the summary CSV by `papers/figures/P4.3/make.py`, and the audit map and barrier plot are drawn from the proved statements cited in their captions. The task's terminal status (abandoned, formal target failed) is recorded in `STATE.md` and `HANDOFF.md`. Every mathematical claim above carries the epistemic tag of the research record; untagged sentences are exposition, not claims.

References

[1] T. Kim and R. Barbulescu, "Extended tower number field sieve: A new complexity for the medium prime case," in *CRYPTO 2016, Part I*, in LNCS, vol. 9814. 2016, pp. 543–571.

[2] O. Schirokauer, "The number field sieve and the discrete logarithm problem," in *Surveys in Algorithmic Number Theory*, vol. 44, in MSRI Publications, vol. 44., Cambridge University Press, 2008, pp. 397–420.

[3] J. Hendrik W. Lenstra, J. Pila, and C. Pomerance, "A hyperelliptic smoothness test, I," *Philosophical Transactions of the Royal Society A*, vol. 345, pp. 397–408, 1993.

[4] J. D. Lee and R. Venkatesan, "Rigorous analysis of a randomised number field sieve," *Journal of Number Theory*, vol. 187, pp. 92–159, 2018.

[5] E. R. Canfield, P. Erdős, and C. Pomerance, "On a problem of Oppenheim concerning ``factorisatio numerorum''," *Journal of Number Theory*, vol. 17, no. 1, pp. 1–28, 1983.

[6] J. A. Buchmann and C. S. Hollinger, "On smooth ideals in number fields," *Journal of Number Theory*, vol. 59, no. 1, pp. 82–87, 1996.

[7] A. Balog, V. Blomer, C. Dartyge, and G. Tenenbaum, "Friable values of binary forms," *Commentarii Mathematici Helvetici*, vol. 87, pp. 639–667, 2012.

[8] R. Barbulescu and A. Lachand, "Some mathematical remarks on the polynomial selection in NFS," *Mathematics of Computation*, vol. 86, pp. 397–418, 2017.

[9] C. Pomerance, "Fast, rigorous factorization and discrete logarithm algorithms," in *Discrete Algorithms and Complexity*, Academic Press, 1987, pp. 119–143.

[10] R. L. Bender and C. Pomerance, "Rigorous discrete logarithm computations in finite fields via smooth polynomials," in *Computational Perspectives on Number Theory*, vol. 7, in AMS/IP Studies in Advanced Mathematics, vol. 7., 1998, pp. 221–232.