ECC RESEARCH PROGRAM · P4.1

Minimum-ρ Pairing-Friendly Families Under a Transparent exTNFS Cost Model

Exact nested-resultant regressions of published SexTNFS estimates, a KSS16 toy-ceiling certificate, and why a ρ-only optimum selects BN

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

with GPT 5.6 Sol · Claude Fable 5

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

ABSTRACT

Choosing a pairing-friendly curve family for a target security level λ requires balancing two attacks: Pollard rho on the order-r subgroup and the (special, extended) tower number field sieve on the embedding field $\mathbb{F}_{p^k}$. The problem posed as P4.1 asks for a reproducible, explicitly parameterized search tool that minimizes $\rho = \log p / \log r$ over the BN, BLS12, BLS24, and KSS16 families subject to an exTNFS cost model, for $\lambda \in \{128, 192, 256\}$. We build such a tool around a fully transparent finite-size cost model in the style of Barbulescu–Duquesne, with every modelling constant a named input and a numerically stable 80-digit interval evaluator for Dickman's ρ-function. Calibration is validated, not fitted: the model reproduces the published BN 99.69-bit worked example within 0.117 bits, and — with **no refitted prefactor** — exact integer nested-resultant sampling of the published polynomial constructions reproduces the printed BLS12 row within 0.011 bits, BN within 0.04 bits, and KSS16 within 0.13 bits once we adopt the coefficient-domain convention of the authors' surviving public code, which provably differs from the domain stated in their paper. The printed BLS24 row does not reproduce under either convention (−1.9 bits); a preregistered sensitivity run shows a draw bound of 10 instead of 9 recovers both printed norms within one bit, isolating an unrecorded historical setting rather than tuning it away. An exhaustive toy sweep of all seeds $|u| \leq 10^4$ under the repository ceiling $p < 2^{60}$ accepts 302 candidates (273 BN, 24 BLS12, 5 BLS24, 0 KSS16), and we prove that **no** KSS16 prime pair exists below that ceiling. In the resulting leading-term optimization BN is provably the unique minimum-ρ family at all three targets and under all three cost scenarios — but we show this label stability is a property of the ρ-only objective, not of stable security estimation, and that it is consistent with practical recommendations favouring BLS12/BLS24. All extrapolated production rows remain explicitly labelled model extrapolations; no concrete prime seed is claimed.

Keywords. pairing-friendly curves · exTNFS · SexTNFS · key-size estimation · Dickman function · BN · BLS12 · BLS24 · KSS16

1 Introduction

A pairing-friendly elliptic curve $E/\mathbb{F}_p$ with subgroup order r and embedding degree k exposes two attack surfaces. The order-r subgroup admits generic attacks costing $\Theta(\sqrt{r})$ group operations,

and the embedding field $\mathbb{F}_{p^k}$ admits variants of the number field sieve, whose complexity dropped sharply when Kim and Barbulescu introduced the extended tower number field sieve (exTNFS) for composite k and its special-form variant (SexTNFS) for polynomially parameterized p [1]. The parameter $\rho = \log p / \log r \geq 1$ measures how much base-field size a family wastes relative to its subgroup size; the classical design goal is ρ as close to 1 as possible, subject to both attacks costing at least 2^λ.

Problem P4.1 asks for this trade-off to be made **transparent**: implement and calibrate an exTNFS cost model whose every constant is a named, inspectable input; implement the BN, BLS12, BLS24, and KSS16 family generators with independent pairing-friendliness verification; sweep an explicitly stated seed space; and report the minimum-ρ optimum at 128, 192, and 256 bits together with its sensitivity to cost-model assumptions. This paper reports the completed pipeline and its findings, including two reproducibility results about the published record that we believe are of independent interest.

Main findings. **(i) Validation without refitting.** EMPIRICAL: BD19 fixtures A single fixed finite-size model reproduces the published BN254 SexTNFS estimate within 0.117 bits, and exact integer nested-resultant sampling of the published constructions then reproduces the BLS12, BN, and KSS16 printed rows within 0.011, 0.04, and 0.13 cost bits respectively — with no prefactor or overhead refitted per family.

(ii) A documented reproducibility gap. PROVED The coefficient domain stated in the Barbulescu–Duquesne paper ($[-A, A]$) and the domain drawn by their surviving public code ($[-A, A+1]$, via inclusive `randint`) are provably different discrete distributions; the printed BLS24 row reproduces under **neither**, while a preregistered draw-bound-10 sensitivity run recovers both printed norms within one bit.

(iii) A complete toy-ceiling certificate. PROVED No KSS16 seed whose field satisfies the repository ceiling $p < 2^{60}$ yields both p and r prime; the proof reduces to 511 explicit seeds and exact factorization.

(iv) The optimum and its meaning. PROVED Within the four-family leading-term model, BN is the **unique** minimum-ρ family at 128, 192, and 256 bits under all three cost scenarios — yet the required BN field size swings by 45–47% between scenarios. Label stability is caused by the formal ρ-only objective, which rewards BN's equal p- and r-sizes even when its base field is far larger than the alternatives'.

Throughout, claims carry the epistemic tags of the underlying research record: PROVED for machine-checked or elementary proved statements, CITED for statements taken from the literature, EMPIRICAL: scope for measured results with their exact scope, and HEURISTIC for model extrapolations that must not be read as security claims.

2 Setting and notation

Fix a pairing-friendly family with polynomial parameterization $(p(u), r(u), t(u))$: for suitable integer seeds u, the curve $E/\mathbb{F}_{p(u)}$ has trace $t(u)$, a subgroup of order $r(u)$ with $r(u) \mid p(u) + 1 - t(u)$, and embedding degree k, the least integer with $r \mid p^k - 1$. We use the four families as restated by Barbulescu and Duquesne [2]: BN ($k = 12$, $\deg p = \deg r = 4$), BLS12 ($k = 12$, $\deg p =$

6, $\deg r = 4$), BLS24 ($k = 24$, $\deg p = 10$, $\deg r = 8$), and KSS16 ($k = 16$, $\deg p = 10$, $\deg r = 8$). The leading-term limits of $\rho = \log p / \log r$ are 1, 3/2, 5/4, and 5/4 respectively.

Security is assessed on two axes. Subgroup security is the generic bound $\approx (\log_2 r)/2$ bits. Field security is the cost of a (S)exTNFS computation in $\mathbb{F}_{p^k}$, written in the L-notation

$$L_Q(\alpha, c) = \exp((c + o(1))(\log Q)^{\alpha}(\log\log Q)^{1-\alpha}), \quad Q = p^k. \tag{1}$$

Kim–Barbulescu give $\alpha = 1/3$ with constants such as $(48/9)^{1/3}$ for composite-degree exTNFS and $(32/9)^{1/3}$ for the special variant relevant to polynomial families [1]; these asymptotics suppress an $o(1)$ and do not by themselves determine finite key sizes, which is why the finite-size machinery of §3 exists [2].

3 The transparent finite-size cost model

3.1 Named inputs

PROVED The shared implementation `lib/tnfs_cost.py` exposes the Barbulescu–Duquesne finite-size decomposition with **every** modelling choice a named input: the coefficient bound A, the smoothness-bound bit size $\log_2 B$, the tower dimension η, the roots-of-unity quotient, the relation- and linear-algebra Galois-automorphism quotients, the factor-base and matrix numerators, the per-row weight, and additive cost overheads. Nothing is hidden in code constants, so a sensitivity scenario is a parameter change, not a code change.

> **Definition 1** (finite-size SexTNFS cost). Given mean norm bit sizes (ν_f, ν_g) for the two sides and a smoothness bound of $b = \log_2 B$ bits, the model evaluates
>
> $$P_s = \rho_D(\nu_s / b), \quad s \in \{f, g\}, \tag{2}$$
>
> with ρ_D Dickman's function; the factor base has $\mathcal{F} = 2B/\ln B$ elements, reduced by the automorphism quotient to the relation target $\mathcal{F}_{\text{red}}$; relation collection costs $\mathcal{F}_{\text{red}}/(P_f P_g)$; the filtered matrix has dimension $D = (B/\ln B)/(b \cdot a_{\text{lin}})$ and sparse linear algebra costs $128 \cdot D^2$; the total is the sum of the two phase costs, reported in bits [2].

The sieve space $((2A+1)^{2\eta})/2$ divided by the roots-of-unity index is checked against the relation target, so a parameter set that cannot supply enough relations is flagged rather than silently extrapolated.

3.2 A numerically stable Dickman evaluator

PROVED Smoothness probabilities are computed from Dickman's delay equation $u\rho_D'(u) + \rho_D(u-1) = 0$, $\rho_D \equiv 1$ on $[0, 1]$, integrated separately on each unit interval from an 80-decimal-digit, degree-30 Chebyshev fit stored and evaluated in the power basis. This construction exists because the obvious one fails: a binary64 cumulative trapezoidal table suffers catastrophic cancellation and returns **non-positive** values near $\rho_D(10.76)$ — exactly the argument the published BLS12 row needs ($791.2/73.5 \approx 10.76$). The interval evaluator retains $\rho_D(2) = 1 - \ln 2$ to displayed precision and is validated through $u = 16$.

EMPIRICAL: published BLS12 norm inputs At the published BLS12 inputs the evaluator returns smoothness log-probabilities -39.171 and -24.671, matching the printed -39.17 and -24.67.

3.3 Calibration against the BN254 anchor

CITED Barbulescu–Duquesne's worked BN instance at $u = -2^{62} - 2^{55} - 1$ reports average norm sizes 414.7 and 460.8 bits, smoothness-bound size 57 bits, and total cost 99.69 bits [2].

EMPIRICAL: BD19 BN fixture Fed the published norm inputs, the model obtains smoothness log-probabilities -21.410 and -25.295 and returns 99.573 bits $-$ 0.117 bits below the published 99.69 and inside the declared 0.2-bit tolerance. No prefactor was fitted to achieve this.

4 Exact nested-resultant norm regressions

The finite-size equation consumes **norm sizes**, which depend on the family's polynomial selection, not only on (p, k). The published rows estimate them by Monte Carlo over coefficient boxes. We re-derive them by **exact integer** computation: for each family we implement the published tower construction and compute both norms as the nested resultants

$$N = |\mathrm{Res}_t(\mathrm{Res}_x(a(t) - b(t)x,\, f(t,x)),\, h(t))|, \tag{3}$$

over $\mathbb{Z}$, with every sampled coefficient recorded [2].

4.1 The BLS12 row, reproduced to 0.011 bits

CITED The published BLS12-128 SexTNFS row uses $h = t^6 - t - 1$, $f = P(x^2 + t + t^2 + t^4 + 1)$ with P the BLS12 field polynomial, $g = x^2 + t + t^2 + t^4 + 1 - u$ at the seed $u = -2^{77} + 2^{50} + 2^{33}$, coefficient bound $A = 1169$, and $\log_2 B = 73.5$; it prints norm sizes 791.2/584.8 and security 131.8 bits [2].

PROVED Our sampler draws all twelve coefficients of $a(t), b(t)$ independently and uniformly from the inclusive integer interval $[-A, A]$, computes both nested resultants exactly over the integers, rejects zero norms, and records every coefficient and every exact norm.

EMPIRICAL: 1,024 exact samples, RNG seed 20260722 Mean integer bit lengths are 791.083 (f-side) and 584.756 (g-side), with normal 95% intervals $[789.720, 792.446]$ and $[584.369, 585.143]$ — both containing the printed 791.2/584.8. Figure 1 shows the two sampled distributions.

EMPIRICAL: same sample Mean **real logarithms** are 790.593 and 584.253, about half a bit below the bit-length means. The printed values therefore follow the integer-bit-length convention although the surrounding prose describes average logarithms; both metrics are preserved in the archived JSON.

EMPIRICAL: same sample, fixed BD19 parameters The sampled means imply smoothness log-probabilities $-39.162/-24.668$, an expected relation supply of 69.470 bits, and a total cost of 131.789 bits; the published norm inputs independently give $-39.171/-24.671$, 69.459, and 131.800. The regression difference is 0.011 bits, and PROVED no prefactor or overhead was refitted — the equation and constants are the same named implementation used by the BN anchor test.

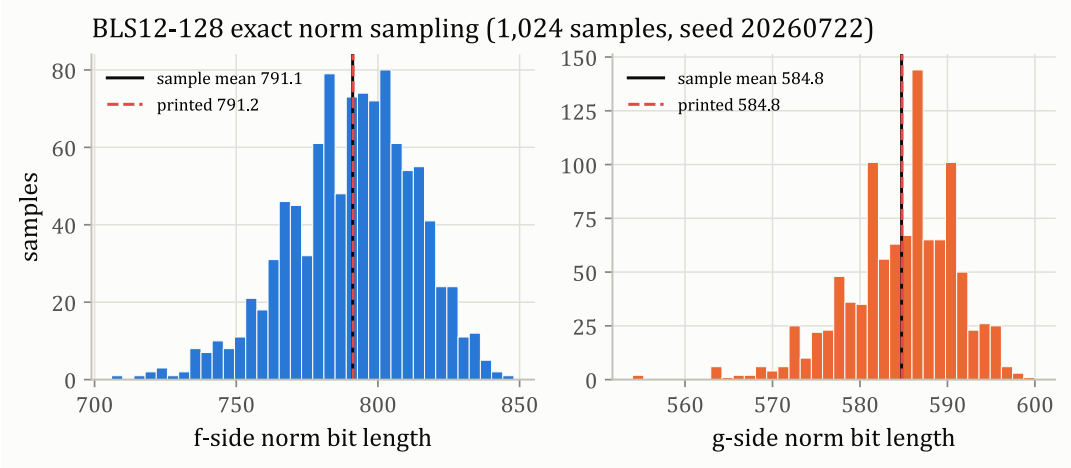

Figure 1: EMPIRICAL: 1,024 exact samples, seed 20260722 Exact integer bit lengths of the two BLS12-128 nested-resultant norms across all 1,024 samples. Solid lines are sample means (791.083/584.756); dashed lines are the printed values (791.2/584.8), both inside the 95% intervals. Data: `bls12_norms_n1024_a1169_s20260722_20260704.csv`.

4.2 A provable paper/code domain discrepancy

CITED The paper states that coefficients are drawn uniformly from $[-A, A]$ (Section 4.2), while the authors' surviving public sampler calls Python `randint(-A, A+1)` — and Python's `randint` includes **both** endpoints, so the code draws from $[-A, A+1]$ [2].

> **Proposition 2** (the two conventions are distinct distributions). PROVED The printed domain $[-A, A]$ and the public-code domain $[-A, A+1]$ cannot denote the same discrete distribution: the former has $2A+1$ equiprobable values and the latter $2A+2$, and the value $A+1$ has probability 0 under one and $1/(2A+2)$ under the other.

Rather than silently merging them, every cross-family experiment below is run **separately** under both conventions, as named, deterministic runs.

4.3 Cross-family audit: BN, KSS16, BLS24

PROVED We implement the exact published nested-resultant constructions for the BN (Section 7.1.1), KSS16 (Section 7.1.3), and BLS24 (Section 7.2.2) rows of [2] and combine all runs into a single audit without changing any sampled result.

Table 1: Cross-family exact-norm audit against the printed Barbulescu–Duquesne rows. All runs use checked-in RNG seeds (20260722; the $A = 10$ sensitivity uses preregistered seed 20260724). Costs in bits.

Profile, draw convention	n	mean norms f/g	printed norms	sampled cost	printed
BLS12-128, paper domain	1,024	791.08 / 584.76	791.2 / 584.8	131.789	131.8
BN-128, paper domain	512	554.37 / 809.40	557.0 / 808.9	131.326	131.3
BN-128, public-code bound	512	554.32 / 809.63	557.0 / 808.9	131.335	131.3
KSS16-128, paper domain	512	911.50 / 628.46	920.4 / 628.9	138.153	139.0
KSS16-128, public-code bound	512	921.82 / 629.65	920.4 / 628.9	138.880	139.0
BLS24-192, paper domain ($A = 9$)	512	1241.57 / 1455.86	1295 / 1460	201.255	203.72
BLS24-192, public-code ($A = 9$)	512	1262.32 / 1457.42	1295 / 1460	201.803	203.72
BLS24-192, public-code ($A = 10$)	128	1295.87 / 1461.00	1295 / 1460	203.171	203.72

EMPIRICAL: checked-in seeded samples As Table 1 and Figure 2 show, BLS12 and BN sampled costs agree with their printed rows within 0.04 bits under either convention. For KSS16 the public-code convention reduces the discrepancy from −0.85 to −0.12 bits — evidence that the historical runs used the code's inclusive bound. An independent-seed confirmation run (seed 20260723, 128 samples, prediction preregistered) kept KSS16 within 1.0 bit and BLS24 more than 1.0 bit low, as predicted.

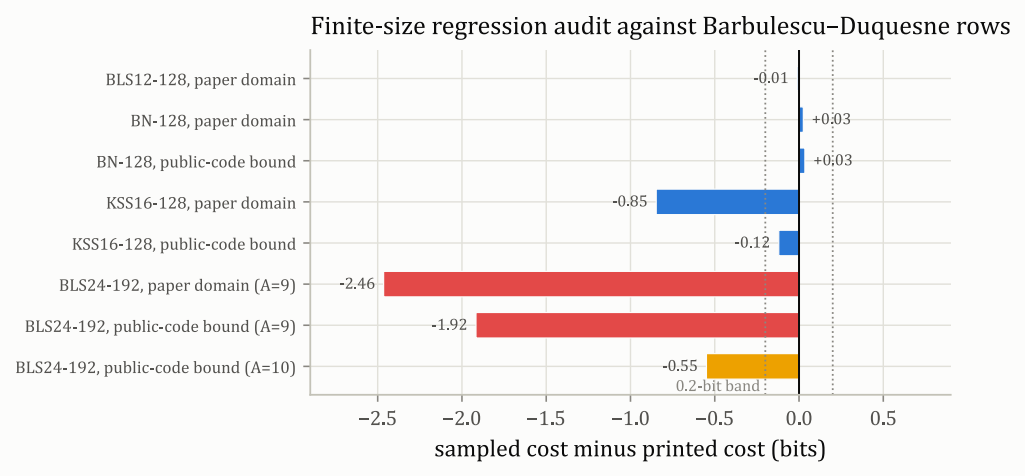

Figure 2: EMPIRICAL: eight audit rows Sampled-minus-printed total cost for every audit row. Blue: within the 0.2-bit calibration band or explained by the public-code convention; red: the irreproducible printed BLS24 $A = 9$ row under both conventions; yellow: the preregistered draw-bound-10 sensitivity, which lands within 0.55 bits. Data: `published_norm_regressions_20260715.csv`.

4.4 The BLS24 row and the isolated historical ambiguity

EMPIRICAL: 512 samples per convention, seed 20260722 With the printed bound $A = 9$, the BLS24-192 row gives 201.255 bits on the paper domain and 201.803 bits under the public-code convention — 2.465 and 1.917 bits below the printed 203.72, far outside every other row's residual. Both printed norms (1295/1460) lie **outside** the sampled 95% intervals.

EMPIRICAL: preregistered, 128 samples, seed 20260724 Changing only the draw bound from 9 to 10 (public-code convention) places **both** printed norm values inside the 95% intervals (means 1295.867/1461.000) and reduces the cost discrepancy to 0.549 bits. The preregistered acceptance threshold (mean f-norm within 10 bits of 1295) was met with 0.87 bits to spare.

HEURISTIC The strongest surviving explanation is an unrecorded internal rounding of the coefficient bound near 10 before the public code applied its extra endpoint. This supports, but does not prove, the inference; the unavailable historical setting is recorded as the non-blocking open question Q027 rather than tuned away. The audit therefore establishes both a positive result — the finite-size equation transfers across families without recalibration — and a negative one: printed norm rows are not always reproducible from the surviving public artifacts, and coefficient-domain conventions alone can move a printed estimate by multiple bits.

5 Family generators, validation, and the bounded toy search

5.1 Independent pairing-friendliness checks

CITED The BN, BLS12, BLS24, and KSS16 evaluators implement the parameterizations restated in Sections 2.1–2.3 of [2]. Every generated candidate is checked independently: deterministic primality of p and r, the divisibility $r \mid p + 1 - t$, and the **least** embedding degree — divisibility $r \mid p^k - 1$ at the claimed k together with non-divisibility at every smaller positive degree.

EMPIRICAL: toy fixtures u=-2,-2,-5 The prime BN, BLS12, and BLS24 toy fixtures satisfy all checks with least embedding degrees 12, 12, and 24.

EMPIRICAL: published BLS12-381 constants Evaluating the BLS12 polynomials at the published seed $u = -0\text{xd201000000010000}$ reproduces the 381-bit p and 255-bit r of BLS12-381 **exactly**, with least embedding degree 12 [3]. This is the independent known-answer fixture for the generators.

5.2 The exhaustive bounded sweep

EMPIRICAL: every seed -10^4 <= u <= 10^4, p < 2^60 The deterministic sweep of all 20,001 integer seeds per family under the repository toy ceiling accepts 302 candidates: 273 BN, 24 BLS12, 5 BLS24, and 0 KSS16. Figure 3 shows all of them. No accepted row approaches a production target: the largest generic-group estimate in the ledger is 29.153 bits.

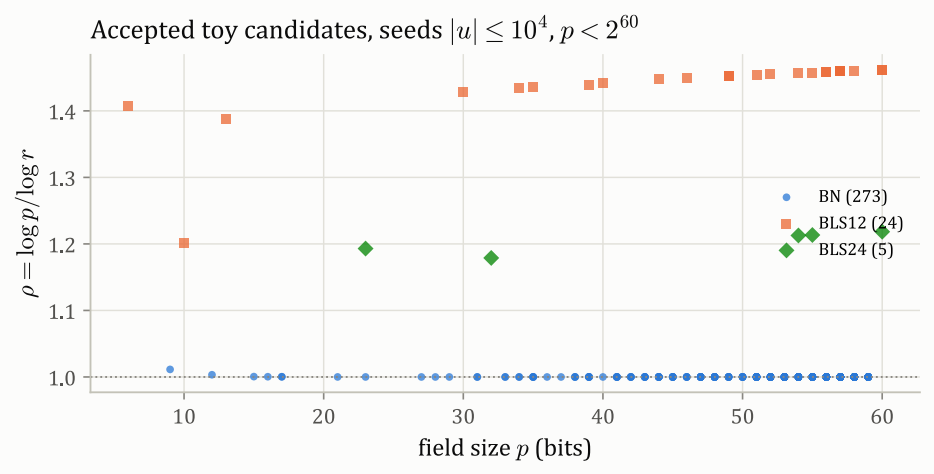

Figure 3: EMPIRICAL: 302 accepted rows The complete accepted ledger of the bounded toy sweep: field size versus ρ, by family. BN dominates numerically (273 rows) and clusters toward $\rho = 1$; BLS12 sits near its leading-term limit $3/2$ at these sizes; the five BLS24 rows sit near 1.25. KSS16 contributes no row — provably (Theorem 3). Data: `search_families_m10000_10000_20260626.csv`.

5.3 No KSS16 fixture exists below the ceiling

The empty KSS16 column is not a search artifact but a theorem.

Theorem 3 (KSS16 toy-ceiling certificate). PROVED Let $N(u) = 980\, p(u)$ be the numerator of the KSS16 field polynomial and $x = |u|$. No integer seed u gives both $p(u)$ and $r(u)$ prime with $p(u) < 2^{60}$.

Proof. For $x \ge 256$ the triangle inequality applied to the ten-term numerator gives

$$N(u) \ge x^6(x^4 - 2x^3 - 5x^2 - 6588), \tag{4}$$

since the omitted lower-order coefficients sum in absolute value to 6588. The bracketed quartic is positive and increasing for $x \ge 256$, and the right-hand side evaluated at $x = 256$ exceeds $980 \cdot 2^{60}$. Hence every seed with $p = N(u)/980 < 2^{60}$ lies in $[-255, 255]$. Deterministic exhaustion of these 511 seeds shows exactly eight produce positive integral family parameters below the ceiling, namely $u \in \{-115, -95, -45, -25, 25, 45, 95, 115\}$, and exact factorization of each candidate pair shows none has both p and r prime. The complete certificate, including the factorizations, is archived as `kss16_p_lt_2pow60_certificate_20260715.json`. □

6 The optimization tables and the meaning of the optimum

6.1 Leading-term extrapolation

Because no toy candidate reaches production scale, the optimization tables are **extrapolations** using only family-polynomial degrees and leading coefficients. HEURISTIC In every row below, "minimum $\log_2|u|$" is a continuous size threshold, not a concrete seed; primality at that size is neither tested nor claimed.

The default field-security scenario is the calibrated special-form model $L_Q\left(1/3, (32/9)^{1/3}\right)$ with a -10.425-bit prefactor calibrated to the single BN anchor of §3.3; subgroup security is $(\log_2 r)/2$.

For each family and target λ the tool returns the least seed size meeting both constraints and the implied (p, r, ρ).

Table 2: HEURISTIC Primary extrapolated optimization table (calibrated SexTNFS scenario). Bold rows are the ρ optima. The binding constraint is the field for BN, BLS12, KSS16 and the subgroup for BLS24 at 128 bits. Full 36-row data: `optimization_extrapolated_20260715.csv`.

λ	family	min $\log_2\|u\|$	p bits	r bits	ρ est.	field / Pollard bits
128	**BN**	109.44	442.9	442.9	**1.0000**	128.0 / 221.5
128	BLS12	74.09	442.9	296.3	1.4947	128.0 / 148.2
128	BLS24	32.00	318.4	256.0	1.2438	150.4 / 128.0
128	KSS16	34.21	332.2	257.8	1.2886	128.0 / 128.9
192	**BN**	278.33	1118.5	1118.5	**1.0000**	192.0 / 559.3
192	BLS12	186.68	1118.5	746.7	1.4979	192.0 / 373.4
192	BLS24	56.08	559.3	448.7	1.2465	192.0 / 224.3
192	KSS16	84.88	838.9	663.2	1.2650	192.0 / 331.6
256	**BN**	550.94	2208.9	2208.9	**1.0000**	256.0 / 1104.5
256	BLS12	368.42	2208.9	1473.7	1.4989	256.0 / 736.8
256	BLS24	110.61	1104.5	884.8	1.2482	256.0 / 442.4
256	KSS16	166.66	1656.7	1317.4	1.2575	256.0 / 658.7

Proposition 4 (BN is the unique minimum-ρ row). PROVED Within the four-family leading-term model, BN uniquely minimizes estimated ρ at each of the three targets and under each of the three cost scenarios: its estimated p- and r-sizes are equal (so $\rho = 1$, the global minimum of the objective), while every competing row has estimated p-size strictly larger than r-size, hence $\rho > 1$.

6.2 Sensitivity: the label is stable, the sizes are not

EMPIRICAL: three named scenarios, three targets Varying the cost model across the calibrated special-form scenario, the same scenario with the $o(1)$-suppressing prefactor removed, and a composite-exTNFS scenario with L-constant 48, the selected family never moves — but the **required BN field size** changes by roughly 45–47%:

Table 3: Required BN embedding-field size (bits) under the three cost scenarios. The optimum's **label** is scenario-independent; its **size** is not.

λ	calibrated $c = 32$	$o(1)$-less $c = 32$	composite $c = 48$	ρ optimum
128	5,315	5,004	3,618	BN
192	13,422	12,870	9,241	BN
256	26,507	25,666	18,348	BN

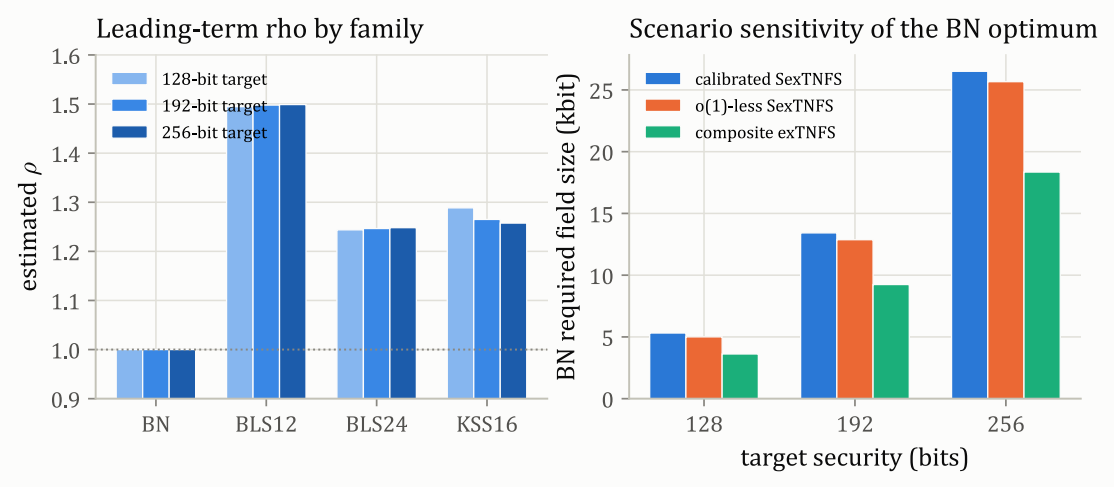

Figure 4: Left: HEURISTIC estimated ρ by family and target under the calibrated scenario — BN sits at the objective's global minimum $\rho = 1$ at every target. Right: EMPIRICAL: three scenarios the scenario dependence of the BN required field size from Table 3 the 45–47% swing shows that stable family selection does **not** mean stable security estimation. Data: `optimization_extrapolated_20260715.csv`.

PROVED This label stability is caused by the formal minimum-ρ objective, not by stable security estimation: minimizing ρ alone rewards BN even when its required base field and pairing-cost proxy are much larger than the competitors'. At 128 bits the BN row needs a 443-bit p (field 5,315 bits) while the BLS24 row needs a 318-bit p — the ρ-optimal family is simultaneously the most expensive row of Table 2 by the deterministic Miller-scalar proxy (1,344 versus 768 and 547).

6.3 Consistency with practical recommendations

CITED Practical 128-bit recommendations refine special-TNFS polynomial selection over a wider family space and favour BLS12 or related $k = 12$ families over roughly 440–448-bit base fields [4]; practical 192-bit shortlists evaluate embedding degrees 15 through 28 alongside BN, BLS, and KSS baselines with implementation-oriented operation counts, listing e.g. 1022-bit BN versus 509-bit BLS24 base fields [5].

PROVED Those recommendations do not contradict the tables above: they optimize a different objective. The P4.1 primary table minimizes ρ subject to the two security constraints and uses the deterministic pairing proxy only as a tie-breaker; the practical literature optimizes operation cost across a larger design space. The correct reading of Proposition 4 is therefore a caution: **a ρ-only ranking is formally well-posed but operationally misleading**, and any deployment-facing use of this pipeline should promote operation cost from tie-breaker to objective.

7 Limitations and open questions

PROVED The delivered optimum is a theorem only inside the explicitly stated four-family leading-term model, for the three named cost scenarios. It is not a statement over the full Freeman–Scott–Teske taxonomy [6]: adding MNT, KSS18, Cocks–Pinch, or Brezing–Weng families, or implementation-level operation costs, defines a strictly larger optimization problem that this tool does not claim to solve.

HEURISTIC The production-size rows of Table 2 remain model extrapolations. Exact finite-size regressions now cover BN, BLS12, KSS16, and BLS24 — but the BLS24 audit demonstrates that a coefficient-bound convention alone can move a printed estimate by nearly two bits, so finite-size numbers should always be read together with their sampling conventions.

Open items, in the repository's numbering: **Q027** — recovering the unavailable historical BLS24 sampling bound (non-blocking; the preregistered $A = 10$ sensitivity supports but cannot prove the internal-rounding explanation). Larger-taxonomy search and concrete production prime-seed generation are recorded as new scope, not unfinished work: the sweep's toy ceiling $p < 2^{60}$ is a repository-wide ground rule, and every number that crosses it in this paper is tagged as an extrapolation.

8 Conclusion

P4.1 asked for a reproducible, explicitly parameterized minimum-ρ optimization over four pairing-friendly families under a transparent exTNFS cost model. The delivered pipeline validates at every level the record permits: the fixed finite-size equation reproduces the published BN anchor within 0.117 bits; exact nested-resultant sampling reproduces the BLS12 row within 0.011 bits and the BN and KSS16 rows within 0.04 and 0.13 bits under the (provably distinct) public-code coefficient convention; the BLS24 row is irreproducible as printed, and the discrepancy is isolated to a one-unit draw-bound ambiguity by a preregistered sensitivity run rather than absorbed into a refit. Below the toy ceiling the search is exhaustive and the KSS16 emptiness is a certificate, not an absence. The optimization itself returns a provably unique answer — BN, at every target and under every scenario — and the analysis shows exactly why that answer should be trusted as a statement about the objective and distrusted as a deployment recommendation: the ρ-only optimum is scenario-stable while the underlying security estimate moves by 45–47%, and the same tables' own cost proxy ranks BN last. Both the theorem and its limits are, we believe, the honest deliverable.

Reproducibility

All numbers trace to checked-in artifacts of the research record: the audit rows to `published_norm_regressions_20260715.csv` and its per-run JSON files (RNG seeds 20260722, 20260723, 20260724 as stated inline), the sweep to `search_families_m10000_10000_20260626.csv`, the optimization and sensitivity tables to `optimization_extrapolated_20260715.csv`, and the KSS16 certificate to `kss16_p_lt_2pow60_certificate_20260715.json`. The model, family evaluators, samplers, and certificate script are `lib/tnfs_cost.py`, `lib/curves.py`, `code/sample_bls12_norms.py`, `code/sample_family_norms.py`, `code/compare_norm_regressions.py`, and `code/certify_kss16_ceiling.py`; the figures in this paper are regenerated from those CSVs by `papers/figures/P4.1/make.py`. Every mathematical claim carries the epistemic tag of the research log; 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] R. Barbulescu and S. Duquesne, "Updating key size estimations for pairings," *Journal of Cryptology*, vol. 32, no. 4, pp. 1298–1336, 2019.

[3] S. Bowe, "BLS12-381: New zk-SNARK elliptic curve construction." 2017.

[4] A. Guillevic, "A short-list of pairing-friendly curves resistant to special TNFS at the 128-bit security level," in *PKC 2020*, 2020.

[5] D. F. Aranha, G. Fotiadis, and A. Guillevic, "A short-list of pairing-friendly curves resistant to the special TNFS algorithm at the 192-bit security level," *IACR Communications in Cryptology*, vol. 1, no. 3, 2024.

[6] D. Freeman, M. Scott, and E. Teske, "A taxonomy of pairing-friendly elliptic curves," *Journal of Cryptology*, vol. 23, pp. 224–280, 2010.