ECC RESEARCH PROGRAM · P1.6

The Height Obstruction That Isn't: Logarithmic-Height Lifts and the Conditional Anatomy of the Xedni Failure

A refutation of the uniform lift-height lower bound, explicit simultaneous lifts for up to four points, a measured growth law, and an exact bounded-relation audit

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

with GPT 5.6 Sol · Claude Fable 5

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

ABSTRACT

A recurring intuition about why lifting attacks on the elliptic-curve discrete logarithm problem fail is that any rational lift $\tilde{P}$ of a finite point $P \in E(\mathbb{F}_p)$ that satisfies the constraints of a xedni-calculus attack must have canonical height $\hat{h}(\tilde{P}) = \Omega(p^c)$ for some explicit $c > 0$. Problem P1.6 asks whether that literal uniform statement is tenable. We prove it is not: every single affine target point admits a short-Weierstrass lift with $\hat{h}(\tilde{P}) = O(\log p)$, and every k-point tuple ($k \le 4$) whose five-column coordinate constraint matrix has row rank k modulo p admits a simultaneous generalized-Weierstrass lift with all selected heights $O_k(\log p)$. Fixed control curves (LMFDB 37.a1 and 11.a2) give even simpler p-independent counterexamples. We then show, against the primary source, that the Jacobson–Koblitz–Silverman–Stein–Teske failure analysis never uses such a lower bound: its mechanism is a **conditional** absolute bound on rational relation coefficients — under Lang's height conjecture and a discriminant-to-height comparison — followed by finite-group counting. We quantify the constructive picture at toy scale: across six primes from 31 to 32719, least-norm lift heights are fitted by $\alpha_k \log p + \beta_k$ with α_k from 4.5 to 10.3, and logarithmic fits beat power fits on the height scale for $k = 2, 3, 4$; an exact bounded-relation audit of 144 lifted curve variants finds 99 finite-field relations but only two rational relations, both two-torsion artifacts. The attack-relevant refinement — forcing rational dependence while the finite inputs exclude small relations — remains open, and we state precisely which hypotheses a theorem there would need. A faithful reproduction of the published $p = 17$ dependence rate was attempted and is reported honestly as a failure.

Keywords. ECDLP · xedni calculus · canonical height · lifting · Mordell–Weil rank · Lang's height conjecture

1 Introduction

The xedni calculus of Silverman [1] was the last serious proposal for a subexponential-style attack on the elliptic-curve discrete logarithm problem (ECDLP) that works by **lifting**: given an instance on $E/\mathbb{F}_p$, one chooses a handful of points derived from the instance, lifts them to rational

points on an elliptic curve $\tilde{E}/\mathbb{Q}$ with good reduction at p, and wins precisely when the lifted points are linearly **dependent** in $\tilde{E}(\mathbb{Q})$. The attack was analyzed and dispatched by Jacobson, Koblitz, Silverman, Stein, and Teske [2], and the episode left behind a widely repeated intuition: lifting fails because rational lifts of finite-field points are forced to have enormous canonical height — the Mordell–Weil group is too sparse at low height to contain the lifts an attack needs.

Problem P1.6 asks for the sharpest version of that intuition. Let $E/\mathbb{F}_p$, let $\tilde{E}/\mathbb{Q}$ have good reduction E at p, and let $\tilde{P} \in \tilde{E}(\mathbb{Q})$ reduce to $P \in E(\mathbb{F}_p)$. The proposed target is a lower bound

$$\hat{h}(\tilde{P}) = \Omega(p^c), \quad c > 0 \text{ explicit}, \tag{1}$$

uniform over the lifts satisfying the constraints a xedni attack needs. The operational brief is to decide whether the literal uniform statement is tenable, to measure one- and multi-point lifts at toy primes under a stated height convention and lift-sampling rule, to fit a growth law with stored residuals, and to state exactly what a proof of an attack-relevant theorem would require.

This paper answers the first question negatively, in constructive and quantitative form, and reorganizes what remains.

Main findings. **(1)** PROVED The literal uniform bound is false for every $c > 0$: every single affine target point has a short-Weierstrass lift with $\hat{h}(\tilde{P}) = O(\log p)$, and for $1 \le k \le 4$ every point tuple whose five-column lift matrix has row rank k over $\mathbb{F}_p$ has a simultaneous generalized-Weierstrass lift in which every selected point has canonical height $O_k(\log p)$ (Theorems 2 and 4).

(2) PROVED Neither construction controls Mordell–Weil rank or forces the selected rational points to be dependent, so neither is a xedni attack; the obstruction that matters lives in the dependence condition, not in height growth (Proposition 5).

(3) CITED The primary 2000 failure analysis derives asymptotic xedni failure from a **conditional** absolute bound on relation coefficients — under Lang's height conjecture and a discriminant-to-point-height comparison — together with finite-group counting. It neither proves nor uses a p^c height lower bound for the selected lifts [2].

(4) EMPIRICAL: p in {31,...,32719}, 18 inputs per k On the observed height scale the least-norm simultaneous-lift heights are better described by an affine function of $\log p$ than by the fitted power law for $k = 2, 3, 4$; the fitted logarithmic slopes are $4.543, 6.983, 8.643, 10.308$ for $k = 1, 2, 3, 4$.

1.1 Contributions and honest scope

We contribute (i) the two refutation theorems with complete elementary proofs (§3–4); (ii) fixed-rank control families that defeat any literal uniform statement without excluding torsion and bounded subgroups (§5); (iii) a primary-source reading of [2] identifying its actual logical inputs, with an exact reproduction of its $p = 257$ small-relation probability (§6); (iv) a seeded, residual-storing measurement of height growth across six primes with a falsifiable $\Theta_k(\log p)$ conjecture (§7); (v) an exact bounded-relation audit of all 144 stored lift variants (§8); and (vi) an honest post-mortem of a failed attempt to reproduce the published $p = 17$ dependence count (§9), with the precise reopening condition.

The scope is stated plainly. The refutation concerns the **literal uniform** height statement; it does not revive xedni calculus, because the missing attack ingredient is rational dependence compatible with finite inputs that exclude useful small relations, and our constructions supply no such dependence (§10). The empirical claims are toy-scale, tied to an explicit sampling rule, and carry their ranges in their tags. Mordell–Weil ranks of the random lifted curves were not computed — the local environment had no working rank routine — and every affected row says so rather than guessing (§11).

2 Setting: heights, lifts, and xedni calculus

Fix a prime $p > 3$ and a nonsingular curve $E/\mathbb{F}_p$. A **lift** of a point $P \in E(\mathbb{F}_p)$ is a pair $(\tilde{E}, \tilde{P})$ where $\tilde{E}/\mathbb{Q}$ has good reduction at p with reduction E, and $\tilde{P} \in \tilde{E}(\mathbb{Q})$ reduces to P. Throughout, heights use the LMFDB/Sage **non-normalized** convention [3], [4]:

Definition 1 (canonical height, LMFDB convention). CITED For $\tilde{P} \in \tilde{E}(\mathbb{Q})$ with x-coordinate written in lowest terms,

$$\hat{h}(\tilde{P}) = \lim_{n\to\infty} n^{-2} h_x([n]\tilde{P}), \quad h_x(a/b) = \log\max(|a|, |b|), \tag{2}$$

and $\hat{h} = 0$ exactly on torsion points. Sage's rational-point height uses the same numerical convention. A convention with prefactor $1/2$ halves every height reported here; PROVED this changes constants but no growth conclusion.

The doubling subsequence gives the equivalent limit $\hat{h}(\tilde{P}) = \lim_n 4^{-n} h_x([2^n]\tilde{P})$, which is what the exact-arithmetic estimator below iterates. Standard properties — quadraticity $\hat{h}([m]P) = m^2\hat{h}(P)$ and vanishing exactly on torsion — are as in [5].

Xedni calculus. CITED The attack of [1], in the form analyzed by [2], samples r random linear combinations of the ECDLP pair — usually $4 \le r \le 6$ — lifts all r finite points to rational points on a single rational elliptic curve, and succeeds only when those rational points are dependent in $\tilde{E}(\mathbb{Q})$: a dependence relation reduces modulo p to a linear relation among the finite combinations, which reveals the discrete logarithm with noticeable probability. The attack's own preprocessing deliberately excludes inputs with **small** finite-field relations, since those yield relations already known before lifting. An attack-relevant obstruction must therefore engage three constraints at once: the reductions are prescribed, the lifted points must be dependent (equivalently, the span of the r lifted points has rank $< r$), and the finite inputs carry no useful small relation. The literal P1.6 target quantifies only the first constraint. The next two sections show that the first constraint alone is cheap.

3 The single-point logarithmic lift

Theorem 2 (direct short lift). PROVED Let $E : y^2 = x^3 + \bar{a}x + \bar{b}$ over $\mathbb{F}_p$ and let $P = (\bar{x}, \bar{y}) \in E(\mathbb{F}_p)$. Then there is a lift $(\tilde{E}, \tilde{P})$ of P with integral coordinates and

$$\hat{h}(\tilde{P}) = O(\log p), \tag{3}$$

with an absolute implied constant.

Proof. Choose the balanced representatives $A, x, y \in [-(p-1)/2, (p-1)/2]$ of $\bar{a}, \bar{x}, \bar{y}$ and set

$$B = y^2 - x^3 - Ax. \tag{4}$$

Since P lies on E, reduction gives $B \equiv \bar{b} \bmod p$, so $\tilde{E} : y^2 = x^3 + Ax + B$ reduces to E, the integer point $\tilde{P} = (x, y)$ lies on $\tilde{E}$ by construction and reduces to P, and the reduction is good because the discriminant of $\tilde{E}$ reduces to the nonzero discriminant of E. The construction satisfies $|x|, |y|, |A| \le p/2$ and $|B| = O(p^3)$. On the fixed lifted curve the duplication map is

$$x([2]Q) = \frac{x(Q)^4 - 2Ax(Q)^2 - 8Bx(Q) + A^2}{4(x(Q)^3 + Ax(Q) + B)}. \tag{5}$$

Homogenizing numerator and denominator to degree four, both have integer coefficients of logarithmic height $O(\log p)$, so evaluating at $x(Q) = u/v$ in lowest terms gives $h_x([2]Q) \le 4h_x(Q) + C \log p$ for an absolute constant C. Iterating n times yields $h_x\big([2^n]\tilde{P}\big) \le 4^n h_x\big(\tilde{P}\big) + \frac{4^n - 1}{3} \cdot C \log p$; dividing by 4^n and passing to the limit,

$$\hat{h}\big(\tilde{P}\big) \le h_x\big(\tilde{P}\big) + (C/3) \log p = O(\log p), \tag{6}$$

because $h_x\big(\tilde{P}\big) = \log|x| \le \log p$. If some iterate is the identity, $\tilde{P}$ is torsion, the height is zero, and the same conclusion holds. □

Corollary 3. PROVED Since $\log p = o(p^c)$ for every fixed $c > 0$, no positive lower bound $\hat{h}\big(\tilde{P}\big) = \Omega(p^c)$ can hold uniformly over all single-point lifts. The literal P1.6 target is false.

The construction is not exotic: it is the **first** lift anyone would try — take the balanced coordinates, adjust the constant term. The intuition that lifts are forced high thus fails at the very first fence. What survives, if anything, must be a statement about **simultaneous** lifts of several points with prescribed structure. That is the next fence, and for up to four points it also falls.

4 Simultaneous lifts for $k \le 4$

For several points on one curve the direct trick no longer has enough freedom: a short Weierstrass model has two coefficients, and k prescribed points impose k conditions. The generalized Weierstrass model has five.

Theorem 4 (simultaneous lifts under a row-rank condition). PROVED Let $E : y^2 = x^3 + \bar{a}x + \bar{b}$ over $\mathbb{F}_p$ and let $P_1, \dots, P_k \in E(\mathbb{F}_p)$ with $1 \le k \le 4$. Write balanced lifts (x_i, y_i) of the points and balanced representatives a, b of $\bar{a}, \bar{b}$. If the $k \times 5$ matrix with rows $(x_i y_i, -x_i^2, y_i, -x_i, -1)$ has row rank k modulo p, then there is a generalized Weierstrass curve $\tilde{E}/\mathbb{Q}$ with good reduction E at p on which every (x_i, y_i) is a rational point, and

$$\hat{h}\big(\tilde{P}_i\big) = O_k(\log p) \quad (1 \le i \le k). \tag{7}$$

Proof. Seek the curve in the five-parameter family

$$\begin{aligned} \tilde{E} &: y^2 + a_1 xy + a_3 y = x^3 + a_2 x^2 + a_4 x + a_6, \\ (a_1, a_2, a_3, a_4, a_6) &= (0, 0, 0, a, b) + p(z_1, z_2, z_3, z_4, z_6). \end{aligned} \tag{8}$$

The condition that (x_i, y_i) lie on $\tilde{E}$ is exactly the linear equation

$$(x_i y_i, -x_i^2, y_i, -x_i, -1) \cdot \boldsymbol{z} = \frac{x_i^3 + a x_i + b - y_i^2}{p}, \tag{9}$$

whose right side is an integer because P_i lies on E modulo p. If the $k \times 5$ coefficient matrix has row rank k modulo p, choose a $k \times k$ minor that is nonsingular modulo p, set the variables outside it to zero, and solve. Cramer's rule gives rational z_j whose denominators (the minor determinant) are prime to p and whose numerator and denominator sizes are polynomial in p for fixed k, since every matrix entry is $O(p^2)$. Consequently the lifted coefficients are p-integral, reduce to $(0, 0, 0, \bar{a}, \bar{b})$, have logarithmic height $O_k(\log p)$, and define a nonsingular rational curve with good reduction at p.

For the heights, complete the square and remove the quadratic term: the standard substitution sends the general model to

$$Y^2 = X^3 - 27c_4X - 54c_6, \quad X = 36x + 3b_2, \tag{10}$$

where b_2, c_4, c_6 are the usual covariants [5]. The transformed coefficients and the images of the selected points are polynomial in the original data, so they still have logarithmic height $O_k(\log p)$. The duplication-map argument of Theorem 2, run on this short model with coefficient heights $O_k(\log p)$, bounds every $\hat{h}$ of an image point by $O_k(\log p)$. Finally, the affine x-coordinate change alters naive heights by at most a curve-dependent constant, which vanishes after division by 4^n in the canonical-height limit; hence the conclusion transfers back to the original generalized model.□

Proposition 5 (this is not an attack). PROVED The hypothesis of Theorem 4 is a row-rank condition on the coordinate constraint matrix, not a Mordell–Weil condition. The construction controls neither the rank of $\tilde{E}(\mathbb{Q})$ nor any dependence among $\tilde{P}_1, \ldots, \tilde{P}_k$; a xedni attack additionally needs the lifted points to be rationally dependent while their finite-field reductions have no useful small relation, and the construction supplies no such dependence guarantee.

Proposition 5 is the pivot of the paper. The height obstruction, read literally, asked the wrong question: cheap simultaneous lifts exist, so sparsity-of-low-height-points cannot be the mechanism that kills lifting attacks. Sections 6–8 examine the mechanism that actually appears in the literature, and what our exact data say about it.

5 Fixed-rank controls

Random constructions aside, fixed rational curves already refute any literal uniform statement that does not exclude torsion and bounded subgroups.

Proposition 6 (rank-one control, 37.a1). CITED The curve 37.a1, $y^2 + y = x^3 - x$, has Mordell–Weil rank one, generator $P = (0, 0)$, and $\hat{h}(P) = 0.0511114082399688$ in the LMFDB/Sage convention [3], [4]. PROVED For every good prime p and $1 \le i \le 4$, the rational points $[i]P$ simultaneously lift their reductions and satisfy $\hat{h}([i]P) = i^2\hat{h}(P)$, so their maximum height is $16\hat{h}(P) \approx 0.818$, independent of p, even though the rational curve has rank one.

Proposition 7 (rank-zero control, 11.a2). CITED The curve 11.a2 has rank zero and torsion group $\mathbb{Z}/5\mathbb{Z}$ with generator $(5, 5)$ [3]. PROVED Its first four nonzero torsion multiples have canonical height exactly zero at every good prime, a still simpler counterexample to any literal uniform statement.

PROVED Both controls, however, have small rational relations known **before** the finite-field problem is posed — $[5](5, 5) = O$ on 11.a2, and the tautological relations among $[i]P$ on 37.a1 — so

they violate the no-small-relation condition needed for a meaningful ECDLP attack. They refute the literal statement while illustrating exactly which hypothesis the literal statement forgot.

6 What the 2000 failure analysis actually uses

The natural historical question is whether [2] proves — or needs — the height lower bound that Sections 3–5 refuted. It does not.

> **Proposition 8** (the conditional mechanism of [2]). CITED The 2000 failure analysis samples finite-field points with deliberately excluded small relations, lifts them, and asks whether the rational lifts are dependent. Under Lang's height conjecture [6] — $\hat{h}(Q) \geq c \log|\Delta|$ for non-torsion Q — together with a discriminant-to-point-height comparison $\log|\Delta| \geq C \max_i \hat{h}(\tilde{P}_i)$ for the constructed lifts, its Theorem 4.1 bounds the coefficients of any rational dependence relation by an absolute constant, and then bounds the success probability by C_0/p by counting, in the finite group, the chance that random reductions satisfy a relation with such bounded coefficients.

Three features deserve emphasis. First, the argument is **conditional**: Theorem 4.1 of [2] is stated under plausible assumptions, and the general unconditional Lang-type estimate the route would need was not available there CITED. Second, the mechanism is about **rare dependence with bounded coefficients**, not about lifts being forced to great height — indeed the attack's own lifting step chooses small-coefficient curves, pulling in the opposite direction from the P1.6 target. Third, the finite-group counting step is exactly checkable, and we checked it.

EMPIRICAL: exact enumeration at p=257 Section 5.4.1 of [2] reports, for Experiment C on $E/\mathbb{F}_{257}$: $y^2 = x^3 + 88x - 41$ with generator $(2, 20)$ and group order 263, that after excluding the identity and $\pm P_1$ the probability of a relation with coefficients at most two is $4/(263 - 3) = 1/65$. Our exact enumeration reproduces every number.

Table 1: Exact reproduction of the Experiment C small-relation probability of [2]. Data: `reproduce_xedni_probability_p257_20260703.csv`.

Quantity	order	eligible P_2	favorable	measured	published
$E/\mathbb{F}_{257} : y^2 = x^3 + 88x - 41$, gen. $(2, 20)$	263	260	4	0.015385	$1/65$

7 Measured height growth

Sections 3–4 give upper bounds; this section measures what an explicit, stated lifting procedure actually produces, with fits and stored residuals as the problem brief requires.

7.1 Design

EMPIRICAL: seed 16062026, recorded in the CSVs The experiment uses the six primes $p \in \{31, 127, 503, 2039, 8191, 32719\}$ — for each $b \in \{5, 7, 9, 11, 13, 15\}$ the largest prime below 2^b congruent to 3 modulo 4 — with three random inputs per prime. Each input is a nonsingular short curve over $\mathbb{F}_p$ and four points sampled so that every prefix of the constraint matrix of Theorem 4 has full row rank modulo p; coordinates are balanced representatives. For each $k = 1, \ldots, 4$ the five-coefficient correction $\boldsymbol{z}$ is chosen as the **least-Euclidean-norm** solution $M^T(MM^T)^{-1}\boldsymbol{r}$ of the constraint system, and two further variants add seeded integer nullspace offsets with parameters in $[-2, 2]$; the direct short lift of Theorem 2 is measured alongside. Every lifted curve and point is

re-verified exactly: correct reduction of all five coefficients, good reduction, containment, and p-integrality are checked and non-p-integral solutions are refused. Canonical heights are computed by exact rational arithmetic along five doublings, $\hat{h} \approx 4^{-5} h_x\big([2^5]\tilde{P}\big)$, with the last-iteration change stored per point.

EMPIRICAL: three LMFDB values, nine exact doublings The height routine was validated against the database before use: it reproduces the LMFDB generator height of 37.a1 (0.0511114082399688) and both generator heights of 389.a1 (0.32700077365160495 and 0.47671165934373954) within $2 \cdot 10^{-6}$ [3], [4]. EMPIRICAL: 216 random general-model variants, five doublings In the experiment itself the last-iteration change was at most 0.4723, at most 0.362% of the reported positive group maximum, with median relative change 0.317%.

7.2 Results

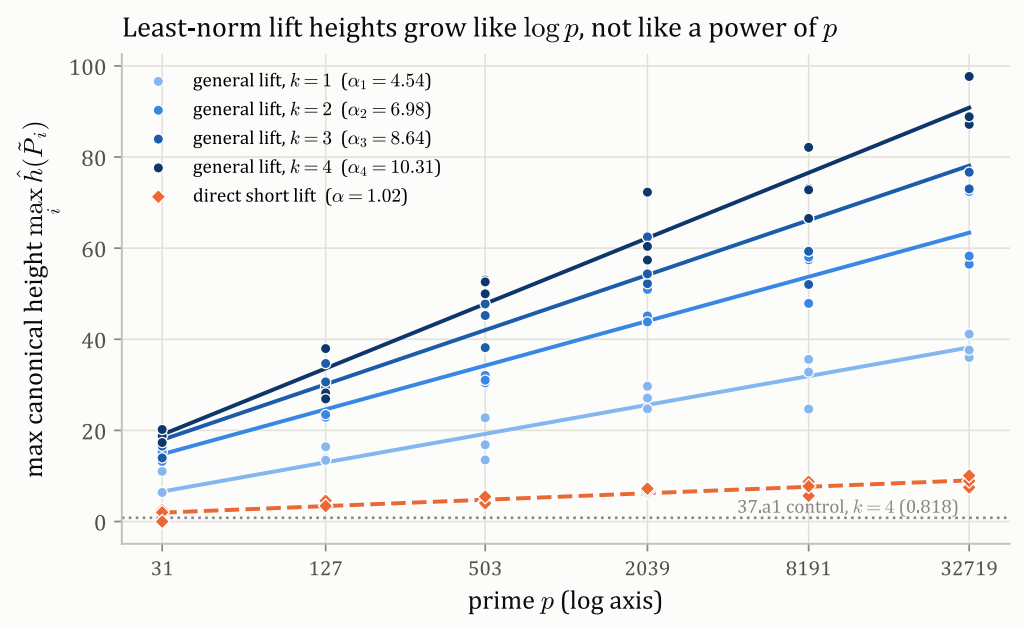

Figure 1: Per-trial maximum canonical heights of the least-norm simultaneous lifts ($k = 1, \ldots, 4$) and the direct short lift, against $\log p$, with the stored OLS logarithmic fits. The 37.a1 control sits flat at 0.818 and the 11.a2 torsion control at zero. Data: `measure_height_growth_b5-7-9-11-13-15_t3_v3_i5_s16062026_20260703_{groups,fits}.csv`.

EMPIRICAL: p in {31,...,32719}, 18 rows per k Regressing the maximum selected height on $\log p$ gives the ordinary-least-squares results of Table 2 the intervals are unadjusted 95% Student intervals over the 18 rows. The direct single-point lift has fitted slope 1.018 with interval $[0.803, 1.233]$ and height-scale RMSE 0.966 — consistent with its proof, whose dominant term is $h_x\big(\tilde{P}\big) \le \log p$. The general-model slopes grow with k, reflecting the O_k constants of Theorem 4.

Table 2: Least-norm general lifts: OLS fits of the maximum canonical height as $\alpha_k \log p + \beta_k$, height-scale RMSE of the logarithmic and power models, and the fitted power exponents. [a]On the 17 rows with positive height. Data: `measure_height_growth_b5-7-9-11-13-15_t3_v3_i5_s16062026_20260703_fits.csv`.

k	slope α_k	95% CI	RMSE (log fit)	RMSE (power fit)	power exponent
1	4.543	$[3.752, 5.334]$	3.548	3.687[a]	0.205
2	6.983	$[6.033, 7.933]$	4.259	5.849	0.201
3	8.643	$[7.346, 9.940]$	5.818	8.275	0.212
4	10.308	$[9.144, 11.472]$	5.219	8.828	0.220

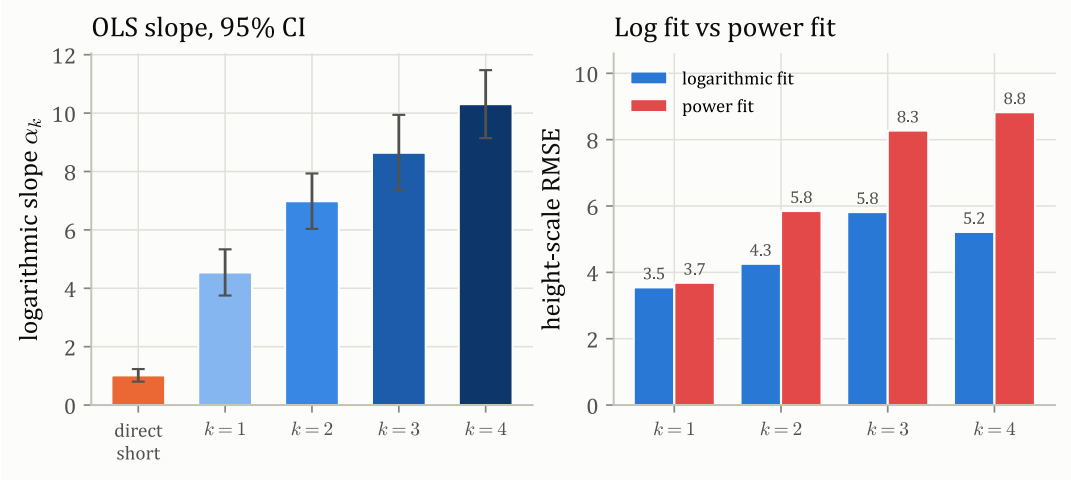

Figure 2: Left: fitted logarithmic slopes with 95% Student intervals for the direct short lift and the least-norm general lifts. Right: height-scale RMSE of the logarithmic versus power fits; the logarithmic model wins for $k = 2, 3, 4$. Data: same fits CSV as Table 2.

Three empirical statements, tagged as in the research notes, summarize the measurement. EMPIRICAL: 18 rows per k On the observed height scale the least-norm experiment is better described by an affine function of $\log p$ than by the fitted power law for $k = 2, 3, 4$ (Table 2, RMSE columns). EMPIRICAL: same range The fitted power exponents $0.205, 0.201, 0.212, 0.220$ drift over a short range and have worse original-scale residuals for $k = 2, 3, 4$; they do **not** constitute evidence for a positive asymptotic exponent. EMPIRICAL: 72 comparisons The least-norm correction was the smallest-height choice among itself and the two sampled nullspace offsets in all 72 comparisons — 18 for each k — so the reported growth is the best of the three sampled variants, not an average-case claim.

7.3 Conjecture and falsifier

CONJECTURE For each fixed $1 \le k \le 4$, under the script's random input distribution and least-Euclidean-norm lift rule, the maximum canonical height is $\Theta_k(\log p)$ in probability. Operationally, refute this by running at bits $17, 19, 21, 23, 25$ with at least 30 independent trials per size and showing either (i) that the 95% interval for the logarithmic slope contains zero, (ii) that the median of height$/\log p$ changes by more than a factor of two between bits 17 and 25, or (iii) that a power fit improves height-scale AIC by at least 10.

8 The exact bounded-relation audit

Proposition 5 says the missing attack ingredient is dependence. The audit quantifies, by exact arithmetic, how often the **measured** lifts exhibit bounded dependence — finite or rational — and what the answer is worth.

8.1 Method

EMPIRICAL: 144 variants, bound 8, no censored rows All stored lift variants at bits $5, 7, 9, 11$ — twelve inputs at each $k = 1, \ldots, 4$, times three lift variants — were searched for a nonzero relation $\sum_i c_i \tilde{P}_i = O$ with $\|\boldsymbol{c}\|_\infty \le 8$, by an exact meet-in-the-middle search that returns a minimum-ℓ^∞ witness; every witness is re-evaluated exactly to the identity. The finite-field tuple is searched first: any rational relation reduces modulo p to a finite relation with the same coefficients, so a finite search through the bound is a necessary filter, and rational rows without a finite relation are recorded as `skipped-no-finite-relation`. Exact rational rows had a 60-second ceiling; no row hit it.

8.2 Results

EMPIRICAL: 144 variants at bits 5-11, bound 8 The audit classified the 144 curve variants as follows: 99 had a finite-field relation through the bound and 45 did not; on the rational side, 97 had no rational relation through the bound after passing the finite filter, 45 were skipped by the filter, and **two** had a rational relation. EMPIRICAL: the two witnesses Both rational witnesses have coefficient -2, canonical height zero, and arise from the same $p = 31$, $k = 1$ point with $y = 0$ — a two-torsion input, $[2]\tilde{P} = O$ — under the least-norm and one nullspace-offset lift. They are torsion controls in the wild, not evidence of useful multi-point dependence. EMPIRICAL: 108 variants, k in {2,3,4} No rational relation through the bound occurred for any multi-point variant.

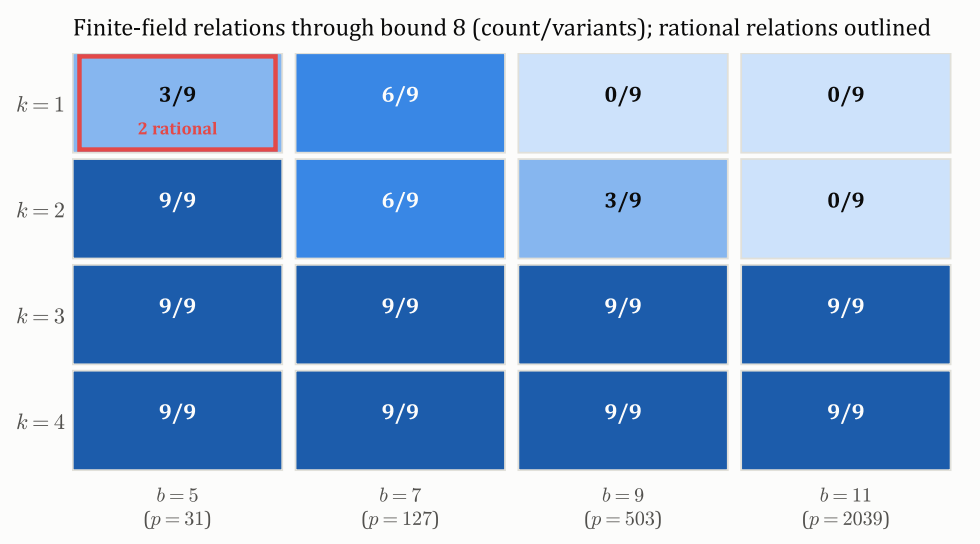

Figure 3: The audit by (k, bits) cell: finite-field relations through ℓ^∞-bound 8 out of nine variants per cell, with the single cell containing rational relations outlined — both witnesses are the same two-torsion input under two lift variants. Every tested $k = 3, 4$ tuple had a finite relation; none had a rational one. Data: `analyze_lift_relations_b5-7-9-11_B8_allv_20260714_summary.csv`.

Table 3: SG-08 audit totals by k (bound 8, bits 5–11, exact arithmetic, zero censored rows). "Skipped" rows failed the necessary finite-field filter. Data: `analyze_lift_relations_b5-7-9-11_B8_allv_20260714_rows.csv`.

k	finite relation / variants	rational relation	rational status of the rest
1	9 / 36	2 (two-torsion, $p = 31$)	7 no-relation, 27 skipped
2	18 / 36	0	18 no-relation, 18 skipped
3	36 / 36	0	36 no-relation
4	36 / 36	0	36 no-relation

> **Remark 9** (what a negative row means). PROVED A row labeled `no-relation-through-bound` is only a bounded exact statement: it certifies that no relation with $\|c\|_\infty \le 8$ exists, and is **not** a certificate of Mordell–Weil independence. We never write "independent" for these rows.

EMPIRICAL: this bounded audit The audit reversed a pre-registered expectation: small finite-field relations turned out to be **common** — every tested $k = 3, 4$ tuple has one through bound eight, an unsurprising fact in hindsight for four points in a group of order $\approx p$ — but they almost never lift to equally small rational relations. Finite relation incidence alone is therefore a weak proxy for useful lifted dependence, which sharpens Proposition 5: the object an attack must manufacture, and an obstruction theorem must exclude, is a rational relation whose coefficients stay bounded while the finite inputs pass the no-small-relation filter.

9 The failed $p = 17$ reproduction

The strongest published quantitative datum on xedni dependence is experimental: CITED Table 3 of [2] reports 317 dependent cases among 100,000 executions of Experiment A at $p = 17$ — on $E/\mathbb{F}_{17} : y^2 = x^3 + 2x + 2$ of order 19 with generator $(3, 1)$ and the restriction $P_1 \neq \pm P_2$ — using projective lattice lifts, small-coefficient nearby curves, and a 2-descent dependence test, computed with LiDIA and SIMATH. Attempt A002 tried to reproduce that number and failed. We record the failure and its boundary precisely, because the temptation to compare an **unfaithful** reimplementation against $317/100000$ is exactly the kind of soft evidence this problem's ground rules exclude.

PROVED The available source does not fix a probability distribution or tie-breaking rule over the many short projective vectors that lift a point and the nearby coefficient-lattice vectors that select a small-discriminant curve; those choices directly change the resulting curve and discriminant distributions, and therefore the dependence rate being estimated. EMPIRICAL: local environment, 2026-07-03 The original LiDIA/SIMATH pipeline and the paper's 2-descent dependence test were unavailable; a coefficient-bound-eight relation search can certify a found relation but cannot certify rational independence, so it is not an equivalent replacement. The local prototype (`reproduce_xedni_p17.py`) reconstructs the finite input, computes integer solution lattices through Smith decomposition, converts valid projective models to standard Weierstrass form, and rechecks containment and reduction exactly — three smoke rows only. No dependency rate produced by it is accepted as evidence or compared with the published count.

CONDITIONAL: original sampling code or a complete specification plus an equivalent 2-descent implementation The validated finite parameters, projective containment equations, Smith-lattice construction, and standard-model conversion can be reused in a faithful retry; this reopening condition is recorded as Q018 in the repository's open questions.

10 What an attack-relevant theorem would require

The refutation and the audit together locate the real open problem. We state its required inputs exactly as the research notes do.

> **Proposition 10** (the shape of a real obstruction theorem). PROVED A height lower bound for every selected lift is neither true (Theorems 2 and 4) nor the logical input used in the 2000 analysis (Proposition 8). A structural **unconditional** theorem that xedni-style lifting fails needs, at minimum: an explicit input distribution over finite-field tuples; a **growing** lower bound on the coefficients of any finite-field relation the inputs admit; and either an unconditional uniform Lang-type lower bound for the smallest non-torsion rational height in the relevant curve family, or a direct unconditional bound on rational relation coefficients for this lift family.

CITED The general unconditional Lang-type estimate required in that route was not available in [2] — their Theorem 4.1 is stated under plausible assumptions rather than as an unconditional result — and Lang's conjecture itself [6] remains open in the required uniformity. The fixed-rank controls of §5 show the input distribution cannot be an afterthought: any formulation must exclude the explicit $O_k(\log p)$ lift family of Theorem 4, and it can only do so through a precise rank, dependence, and finite-field no-small-relation condition, not through height.

11 Limitations

We list the boundaries of the evidence, all recorded row-level in the data.

EMPIRICAL: local environment, 2026-07-03 Sage and PARI rank routines were unavailable, so the total Mordell–Weil ranks of the random lifted curves were not measured; every random row carries `rank_status = total rank unavailable` rather than a guessed rank. The span rank of the lifted points was likewise not measured.

PROVED The random data measure one explicit coefficient-minimizing procedure — balanced representatives, least-Euclidean-norm correction, two bounded nullspace offsets — not the minimum over all lifts and not a rank-conditioned attack distribution. The least-norm rule is a construction bias, stated as such.

EMPIRICAL: six primes, 31 to 32719 The fitted slopes and the log-versus-power comparison live on a short range; the conjecture of §7.3 carries its own quantitative falsifier at larger bits precisely because the range is short. The audit of §8 is exact but bounded (coefficients through eight) and at toy primes (bits 5–11). The $p = 17$ reproduction failed and contributes nothing quantitative (§9).

Against the problem's falsifier list: the refutation here is of the **literal uniform** statement, via constructions outside the attack constraints; it identifies the missing hypotheses rather than refuting an attack-constrained obstruction, which no experiment in this repository reaches.

12 Conclusion

P1.6 asked whether the canonical height of rational lifts is the obstruction that makes lifting attacks on ECDLP fail, in the strong literal form $\hat{h}(\tilde{P}) = \Omega(p^c)$ uniformly over attack-relevant lifts. It is not: single points always lift at height $O(\log p)$ (Theorem 2), tuples of up to four points lift simultaneously at height $O_k(\log p)$ under an explicit and generically satisfied row-rank condition (Theorem 4), fixed curves give p-independent counterexamples (§5), and the primary

failure analysis of xedni calculus never used such a bound in the first place — its mechanism is conditional rarity of bounded-coefficient dependence (§6). The measured growth of an explicit lift procedure is logarithmic on the observed range with slopes 4.5–10.3 (§7), and an exact audit of 144 lifted variants found bounded rational dependence only in two torsion artifacts (§8). What remains open is exactly the dependence-conditioned question: a structural theorem combining a no-small-relation input distribution with rational dependence or rank below k, whose required inputs we have stated (§10), or a genuinely faithful reproduction of the historical dependence rate, whose blocking gap we have recorded (§9). Height alone was the wrong gatekeeper; the gate is dependence.

Reproducibility

All measurements are seeded and re-runnable from the repository: `measure_height_growth.py` (six primes, three trials, three lift variants, five exact doublings, seed 16062026; about 41 s) wrote the 678 point rows, 282 curve/group rows, 102 summary rows, 30 fit rows, and 393 row-level residuals used in §7; `analyze_lift_relations.py` (bound eight, about 76 s) wrote the 144-row audit of §8; `reproduce_xedni_probability.py` wrote the exact $1/65$ reproduction of §6; `reproduce_xedni_p17.py` is retained as the failed-attempt prototype of §9 and must not be used for rate comparisons. The height routine and its LMFDB validation live in `lib/heights.py` and `lib/tests/test_heights.py`. Figures load the named CSVs directly. Every mathematical claim above carries one of the epistemic tags PROVED, CITED, EMPIRICAL: range, CONDITIONAL: hypothesis, CONJECTURE as used in the research log; untagged sentences are exposition, not claims.

References

[1] J. H. Silverman, "The xedni calculus and the elliptic curve discrete logarithm problem," *Designs, Codes and Cryptography*, vol. 20, no. 1, pp. 5–40, 2000.

[2] M. J. Jacobson, N. Koblitz, J. H. Silverman, A. Stein, and E. Teske, "Analysis of the xedni calculus attack," *Designs, Codes and Cryptography*, vol. 20, no. 1, pp. 41–64, 2000.

[3] The LMFDB Collaboration, "The L-functions and modular forms database: canonical height of a rational point (knowledge entry ec.canonical_height) and curve data for 37.a1, 389.a1, and 11.a2." 2026.

[4] The Sage Developers, "SageMath reference manual: canonical heights for elliptic curves over number fields." 2026.

[5] J. H. Silverman, *The Arithmetic of Elliptic Curves*, 2nd ed. in Graduate Texts in Mathematics, vol. 106. Springer, 2009.

[6] S. Lang, *Elliptic Curves: Diophantine Analysis*. in Grundlehren der mathematischen Wissenschaften, vol. 231. Springer, 1978.