One Query on One Side: An Elliptic-Curve-Backed Generic Oracle Separating Fixed-Argument Pairing Inversion from ECDLP

Stage separation, a Satoh distortion transfer, and a typed random-representation lower bound with its exact scope

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

with GPT 5.6 Sol · Claude Fable 5

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

ABSTRACT

Fixed-argument pairing inversion (FAPI-1) asks, given a bilinear pairing $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$, a fixed $P \in \mathbb{G}_1$, and a target $z \in \mathbb{G}_T$, for the point $Q \in \mathbb{G}_2$ with $e(P, Q) = z$. The task posed as P2.4 was to either prove a polynomial-time equivalence between FAPI-1 and ECDLP in $\mathbb{G}_2$ or construct an oracle separation. We construct the separation, and we make it elliptic-curve-backed: a typed random-representation (RR/Shoup) bilinear oracle whose three groups are random encodings of actual supersingular curve groups $E : y^2 = x^3 + x$ over $\mathbb{F}_p$ with $p \equiv -1 \pmod{4r}$, equipped with a free discrete-logarithm gate on $\mathbb{G}_2$ and no target-to-source operation. Relative to this oracle, ECDLP in $\mathbb{G}_2$ takes one query, while every algorithm creating at most t target handles inverts the fixed-argument pairing with probability at most $\big(\binom{t}{2} + 1\big)/r + O(q/2^L)$; constant success needs $t = \Omega(\sqrt{r})$, and a Markov–Borel–Cantelli argument extracts one fixed oracle defeating every probabilistic polynomial-time machine. On the constructive side, we separate the Miller and final-exponentiation stages at toy scale: exhaustive fibre enumeration on six curve realizations ($43 \le p \le 163$) confirms fibre size $d = (p^2 - 1)/r$ with exactly one Miller-compatible representative per nonidentity target; a proved distortion-map identity $f_{r,P}(\psi(R)) = i^{-r} f_{r,\psi^{-1}P}(R)$ transfers Satoh's polynomial-time Miller inversion to the FAPI-1 orientation on this family; the transferred inverse succeeds on all 82 raw targets, while the canonical final-exponentiation root is Miller-compatible in 0 of 82 cases — isolating representative selection at the final-exponentiation interface as the residual obstruction. The collision core of the oracle bound is verified over 541,966 exhaustive affine-form sets with zero violations and tight rows, and the theorem survived an independent audit and a second adversarial self-verification. The result is a statement about typed generic encodings: we claim no lower bound for algorithms given concrete $\mathbb{F}_{p^2}$ target-field arithmetic, coordinates, or a coordinate-to-label map.

Keywords. pairing inversion · FAPI-1 · Tate pairing · Miller functions · final exponentiation · generic group model · oracle separation

1 Introduction

Every pairing-based cryptosystem rests on the presumed one-wayness of the pairing itself. For a non-degenerate bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ on an elliptic curve with embedding degree k, the **fixed-argument pairing inversion** problems ask:

- **FAPI-1**: given e, $P \in \mathbb{G}_1$, and $z \in \mathbb{G}_T$, find $Q \in \mathbb{G}_2$ with $e(P, Q) = z$;
- **FAPI-2**: the same with the roles of $\mathbb{G}_1$ and $\mathbb{G}_2$ exchanged.

CITED Galbraith, Hess, and Vercauteren showed that these problems are dangerous well beyond re-evaluating the pairing: one-sided FAPI computes every nontrivial homomorphism between the source groups and solves the corresponding bilinear Diffie–Hellman problem, and both directions together solve CDH in all three groups [1]. What their paper does **not** contain is a reduction from ECDLP in $\mathbb{G}_2$ to FAPI-1. The research task for P2.4 was to settle this relationship in one of two ways: prove a polynomial-time equivalence between FAPI-1 and $\mathbb{G}_2$-ECDLP, or construct an oracle separation. A stated partial target was to expose the Miller loop and the final exponentiation as separate stages, measure their fibres and inversion costs at toy scale, characterize the fixed-argument Miller function, and state precisely where an attempted ECDLP-to-FAPI reduction stops.

This paper reports the completed resolution: the separation exists, and the staged analysis explains **why** the natural reduction cannot get off the ground.

Main result. PROVED There exists an oracle relative to which ECDLP in $\mathbb{G}_2$ is solvable with one query, while every probabilistic polynomial-time algorithm for FAPI-1 has negligible success. Concretely, at prime order r, an algorithm creating at most t target-group handles and making q total queries succeeds with probability at most $\binom{t}{2} + 1)/r + O(q/2^L)$, where $L \geq 3\lceil \log_2 r \rceil$ is the typed-label length; constant success requires $t = \Omega(\sqrt{r})$. The oracle's three groups are random encodings of actual supersingular elliptic-curve groups with embedding degree two, so the separation is elliptic-curve-backed, not purely abstract.

PROVED **Scope.** This is a generic bilinear oracle separation in the RR/Shoup typed-encoding model. It is **not** a lower bound for algorithms given ordinary curve coordinates, a coordinate-to-label map, or full $\mathbb{F}_{p^2}$ arithmetic on target elements.

1.1 Contributions and honest scope

We contribute (i) exact structural lemmas for the two stages of reduced Tate pairing inversion — power-map fibre sizes, the unique Miller-compatible representative, and the $(r-2, r-3)$ factor-degree law for the binary fixed-argument Miller expression (§3); (ii) an exhaustive toy-scale measurement of both stages on six curve realizations, with the published $\mathbb{F}_{43}$ test vector reproduced exactly (§4, Figure 2); (iii) a proved distortion-map identity transferring Satoh's polynomial-time Miller inversion [2] to the FAPI-1 orientation on the supersingular $j = 1728$ family, validated on all 82 raw targets (§5, Figure 3); (iv) the precise reduction obstruction (§6); (v) the elliptic-curve-backed generic oracle separation with a complete proof (§7) and an exhaustively verified collision core (§8, Figure 4); and (vi) the audit and adversarial self-verification record (§9).

We state the scope plainly. The separation theorem lives in a typed generic oracle model; the concrete-representation question — whether coordinate access to $\mathbb{F}_{p^2}$ target elements collapses or preserves the hardness gap — is explicitly outside the proved claim, and we mark it as a different,

stronger problem rather than unfinished work (§10). Toy-scale timings are implementation observations, never asymptotic claims.

2 Setting and notation

Fix a prime $p \equiv 3 \pmod 4$ and the supersingular curve

$$E : y^2 = x^3 + x \quad \text{over } \mathbb{F}_p, \quad \#E(\mathbb{F}_p) = p + 1, \tag{1}$$

with a prime $r \mid p+1, r > 2$. Then $p \equiv -1 \pmod r$, so the embedding degree is $k = 2$: $r \mid p^2 - 1$. Write $\mathbb{F}_{p^2} = \mathbb{F}_p(i)$ with $i^2 = -1$. We take $\mathbb{G}_1 = \langle P \rangle \subset E(\mathbb{F}_p)[r]$, $\mathbb{G}_2 = \langle Q \rangle$ the opposite Frobenius eigenspace of $E[r]$, and $\mathbb{G}_T = \mu_r \subset \mathbb{F}_{p^2}^\times$ the r-th roots of unity. The distortion map

$$\psi(x, y) = (-x, iy) \tag{2}$$

carries $\mathbb{G}_1$ to $\mathbb{G}_2$; we write $Q = \psi(P)$.

The **reduced Tate pairing** is computed in two stages: the Miller loop evaluates the function $f_{r,P}$ with divisor $r(P) - r(\mathcal{O})$ at Q, giving a **raw** value $v \in \mathbb{F}_{p^2}^\times$, and the **final exponentiation** raises it to $d = (p^2 - 1)/r$:

$$e(P, Q) = f_{r,P}(Q)^{(p^2-1)/r} \in \mu_r. \tag{3}$$

FAPI-1 splits accordingly into **final-exponentiation inversion** (FEI) — find some v with $v^d = z$ — followed by **Miller inversion** (MI) — find Q with $f_{r,P}(Q) = v$ [1]. The interplay of the two stages, and in particular **which** v FEI must produce for MI to have any chance, is the subject of §§3–5.

All experiments run at toy scale, $p \le 163$, under the repository's parameter ceiling; every dataset is seeded and deterministic.

3 The two stages and their exact structure

3.1 Final exponentiation is a regular d-to-one map

Lemma 1 (power-map fibres). PROVED Let $N = q^k - 1$, let $r \mid N$ be prime, and put $d = N/r$. On the cyclic group $\mathbb{F}_{q^k}^\times$, the map $x \mapsto x^d$ has kernel of size d, image of size r (namely μ_r), and exactly d preimages over every point of its image.

Proof. On a cyclic group of order N, the power map $x \mapsto x^d$ has kernel of size $\gcd(N, d) = d$ (since $d \mid N$), and the first isomorphism theorem gives image size $N/d = r$ and constant fibre size d over the image. □

FEI in isolation is therefore easy in a strong sense: the fibre is a coset of the kernel, and any method that produces one d-th root — for instance a discrete logarithm in the target subgroup, cheap at toy scale — produces a valid FEI answer. The trap is that FAPI-1 does not ask for **a** root.

Lemma 2 (unique compatible representative). PROVED For fixed nonzero P in a non-degenerate pairing of prime order r, the map $Q \mapsto e(P, Q)$ is a bijection from $\mathbb{G}_2$ to μ_r. Consequently, for every nonidentity target $z \in \mu_r$ whose direct Miller evaluation is defined, exactly one of the d final-exponentiation preimages of z is the raw Miller value of the unique $Q \in \mathbb{G}_2$ with $e(P, Q) = z$: the Miller-compatible fraction of each fibre is exactly $1/d$.

Proof. $Q \mapsto e(P,Q)$ is a group homomorphism between cyclic groups of prime order r; non-degeneracy makes it nontrivial, hence injective, hence bijective. For nonidentity z, let Q be its unique preimage and $v = f_{r,P}(Q)$; then $v^d = z$, so v lies in the FEI fibre of z. If v' in that fibre is also Miller-compatible, i.e. $v' = f_{r,P}(Q')$ for some Q' in the cyclic domain, then $e(P,Q') = (v')^d = z$ forces $Q' = Q$ and $v' = v$. □

This is the information-selection obstruction in its exact form: a fast arbitrary-root FEI algorithm hands MI a correct input only with probability $1/d$ on a uniform fibre element. CITED Galbraith–Hess–Vercauteren observe the same phenomenon in the opposite direction: on their larger Tate–Lichtenbaum divisor domain a **random** final root is useful with non-negligible probability — but that domain is not the cyclic $\mathbb{G}_2$ domain fixed by the formal P2.4 statement, and the distinction matters ([1], Example 18).

3.2 The fixed-argument Miller function

Proposition 3 (factor-degree law). PROVED For an odd order r, the left-to-right binary Miller expression for $f_{r,P}$, after its final vertical-line cancellation, has $r-2$ affine line factors in its numerator and $r-3$ vertical factors in its denominator.

Proof. Induction along the binary expansion: the intermediate expression for scalar n has $n-1$ factors on each side (doubling maps $(n-1, n-1)$ to $(2n-1, 2n-1)$ after counting the new line and vertical; addition of P adds one to each). The last bit of r doubles the accumulator at scalar $(r-1)/2$ to $-P$ and adds P to reach $\mathcal{O}$: the final addition introduces no denominator (the vertical through $\mathcal{O}$ is trivial) and its numerator line is the vertical through P, which cancels the vertical denominator introduced by the immediately preceding doubling. Net count: $r-2$ lines over $r-3$ verticals. □

EMPIRICAL: orders 3–41, seeded The binary-loop factor counter reproduces $(r-2, r-3)$ for every odd order $r \in \{3, 5, \ldots, 41\}$ (Figure 1 rows in `analyze_miller_function_p43_r11_20260624.csv`).

At the published test vector the function itself is small enough to write down. EMPIRICAL: p=43, r=11, all 10 nonidentity G2 points On $E/\mathbb{F}_{43} : y^2 = x^3 + x$ with $P = (23, 8)$, the validated raw fixed-argument Miller function is, in $\mathbb{F}_{43}[x,y]/(y^2 - x^3 - x)$,

$$f_{11,P}(x,y) = \frac{(y + 19x + 28)^4 (y + 39x + 20)^2 (y + 14x + 14)^2 (y + 36x + 24)}{(x + 29)^4 (x + 12)^2 (x + 30)^2}, \tag{4}$$

with factor degrees $9/8$ as the proposition predicts. After reducing powers of y, the numerator is $A(x) + B(x)y$ with $\deg A = 13$ and $\deg B = 12$, while the denominator has x-degree 8; the complete coefficient arrays are recorded in the dataset, and the symbolic function agrees with the numeric Miller loop at all ten nonidentity $\mathbb{G}_2$ points.

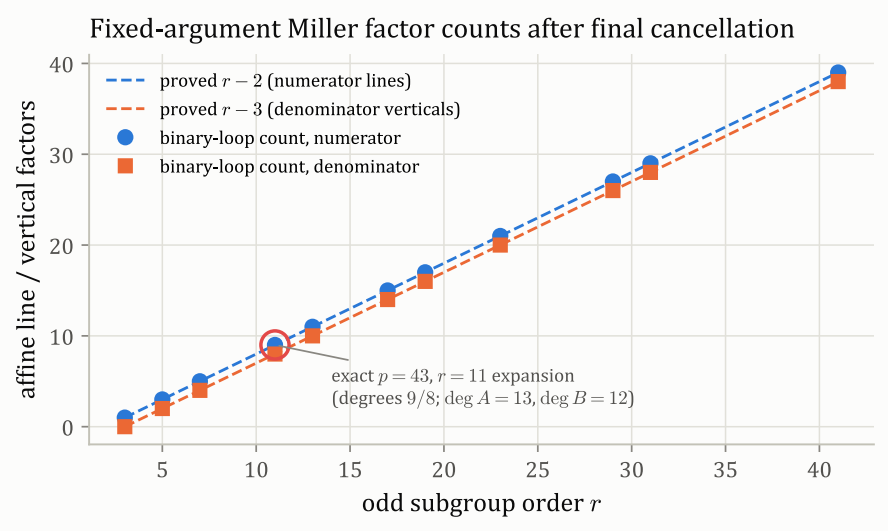

Figure 1: EMPIRICAL: orders 3–41 Fixed-argument Miller factor counts after the final vertical cancellation. Markers are the binary-loop counts from `analyze_miller_function_p43_r11_20260624.csv`; dashed lines are the proved $(r-2, r-3)$ law. The circled point is the exact $p = 43$, $r = 11$ expansion displayed in the text.

4 Toy-scale stage measurements

4.1 Validation and realizations

EMPIRICAL: p=43, r=11, k=2 The staged implementation reproduces Stögbauer's published vector exactly: with $P = (23, 8)$, $Q = (20, 8t)$, $t^2 = -1$, the composed pairing gives $e(P, Q) = 11 + 3t$ and $e(2P, Q) = 26 + 23t$, and the raw loop follows the published `double,double,add,double,add` trace [3]. Raw representatives may differ from the source's by a nonzero $\mathbb{F}_{43}$ factor because the source omits vertical denominators; the final exponent 168 kills every such factor since $42 \mid 168$, so the reduced values are the invariant comparison.

Six realizations of the family were then enumerated exhaustively (Table 1). EMPIRICAL: six curves, exhaustive enumeration In every case the final fibre had its predicted size d and every nonidentity pairing target had exactly one compatible raw Miller value on the cyclic FAPI-1 domain — the exact statements of Lemmas 1 and 2.

Table 1: The six toy realizations of $E : y^2 = x^3 + x$, $p \equiv -1 \pmod{4r}$. Exhaustive enumeration confirmed fibre size d for every final-exponentiation fibre and exactly one Miller-compatible raw value per nonidentity target; all 300 deterministic bilinearity trials passed and every base pairing was nonidentity of order r. Data: `measure_pairing_stages_p43-59-83-103-131-163_20260624.csv`.

p	r	P	p^2-1	$d=(p^2-1)/r$	targets	compatible raw per target
43	11	(23, 8)	1848	168	10	1
59	5	(35, 31)	3480	696	4	1
83	7	(69, 8)	6888	984	6	1
103	13	(49, 22)	10608	816	12	1
131	11	(121, 13)	17160	1560	10	1
163	41	(88, 63)	26568	648	40	1

4.2 Measured inversion costs

Four direct inversion strategies were timed on 50 seeded targets per curve: an arbitrary FEI root via a toy target-subgroup DLP, a full-field FEI scan, naive raw Miller inversion by scanning the cyclic domain, and composed pairing inversion (Table 2, Figure 2).

Table 2: EMPIRICAL: 50 seeded timings per curve Mean per-target inversion times with 95% normal-approximation confidence half-widths. Naive Miller inversion is 147–545 times slower than finding an arbitrary final root in this unoptimized implementation; the composed cost tracks the Miller stage almost exactly.

p	r	FEI root via target DLP	naive MI	composed inversion	ratio MI / FEI
43	11	$24.7 \pm 2.4\mu s$	8.15 ± 1.08 ms	8.21 ms	330
59	5	$13.6 \pm 1.2\mu s$	1.99 ± 0.25 ms	2.07 ms	147
83	7	$17.8 \pm 1.8\mu s$	4.22 ± 0.60 ms	4.34 ms	238
103	13	$40.9 \pm 6.0\mu s$	15.25 ± 2.03 ms	15.94 ms	373
131	11	$32.0 \pm 4.2\mu s$	13.04 ± 1.71 ms	13.32 ms	407
163	41	$102.0 \pm 13.1\mu s$	55.59 ± 8.67 ms	57.63 ms	545

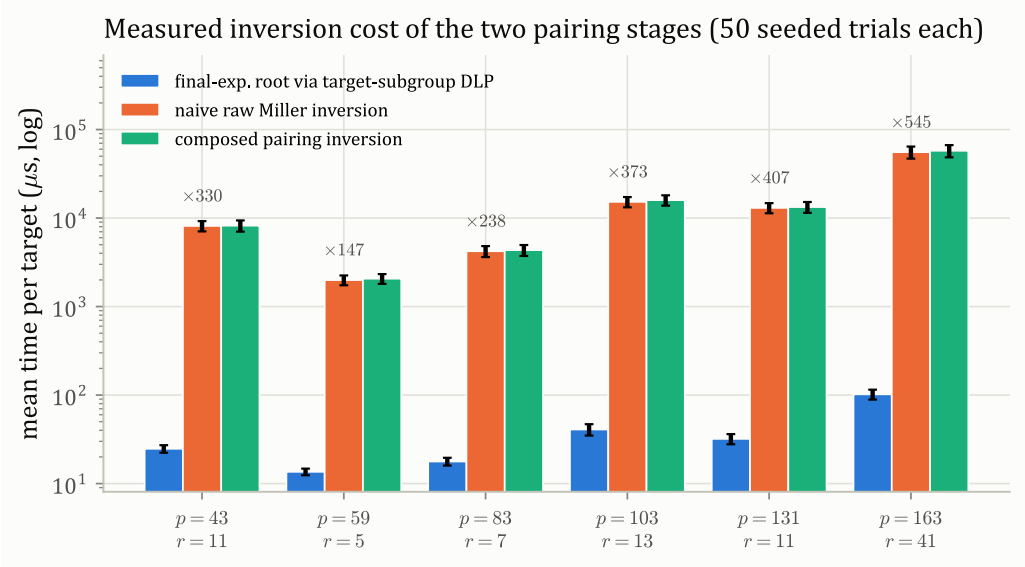

Figure 2: EMPIRICAL: six curves, 50 trials each Measured inversion cost of the two stages (log scale; error bars are 95% CI half-widths). Annotations give the naive-MI-to-FEI-root ratio. Data: `measure_pairing_stages_p43-59-83-103-131-163_20260624.csv`.

4.3 The bottleneck, stated carefully

Two distinct findings must not be conflated. EMPIRICAL: six curves If a **correct raw target** is already supplied, naive Miller inversion is the measured computational bottleneck; final exponentiation adds little to direct exhaustive FAPI-1 inversion. PROVED But if only the **pairing target** is supplied, arbitrary final-stage inversion is insufficient on the cyclic domain, because exactly a $1/d$ fraction of each fibre is Miller-compatible (Lemma 2). The information-selection obstruction therefore lies at the interface created by final exponentiation — even when a chosen Miller-inversion algorithm is fast, as the next section makes concrete.

EMPIRICAL: 82 targets, one deterministic selector Across all 82 nonidentity targets of the six curves, the canonical root returned by the experiment's target-subgroup-DLP method was Miller-compatible in 0 of 82 cases. This tests one deterministic root selector, not all root-extraction algorithms.

5 Satoh Miller inversion transfers to the FAPI-1 orientation

CITED Satoh proved that Miller inversion with a fixed q-eigenspace argument and a variable base-field argument is easy: deterministic $O((k \log q)^3)$ bit operations for even embedding degree, and probabilistic average $O(k^6 (\log q)^3)$ for odd k, assuming a stated square-root precomputation ([2], Algorithm 4.1, Theorems 4.2 and 5.3). Under P2.4′s convention $\mathbb{G}_1 = E(\mathbb{F}_q)[r]$, this is the FAPI-2 Miller orientation. Whether it says anything about FAPI-1 on our family is settled by a normalization identity.

Lemma 4 (distortion transfer). PROVED On $E : y^2 = x^3 + x$ with $\psi(x, y) = (-x, iy)$, $i^2 = -1$, normalize Miller functions to leading coefficient one in the local parameter $\tau = x/y$ at infinity, so that $\tau \circ \psi = i\tau$ and $f_{r,P}$ has leading term τ^{-r}. Then $f_{r,P} \circ \psi$ and $f_{r,\psi^{-1}P}$ have the same divisor, and comparing leading terms gives

$$f_{r,P}(\psi(R)) = i^{-r} f_{r,\psi^{-1}P}(R). \tag{5}$$

A raw FAPI-1 Miller target v therefore transfers to Satoh's orientation as $i^r v$.

Corollary 5. PROVED Combining the transfer identity with Satoh's theorem yields polynomial-time raw Miller inversion for the fixed-base FAPI-1 orientation on this supersingular degree-two family: compute $u = \left(v'^{(q+1)/d}\right)^{(q+1)/2}$ for the transferred target v', reject unless $u \in \mathbb{F}_q$, and test the at most four base-field points above $x(A) \pm u$.

Remark 6. PROVED This does **not** solve reduced FAPI-1. From $z = v^{(q^2-1)/r}$ the algorithm still needs the unique raw representative v compatible with the cyclic Miller image (Lemma 2); Satoh's own Example 4.4 makes the same distinction in the opposite orientation [2].

EMPIRICAL: Satoh Example 4.4 The implementation reproduces the published example exactly: $u = 131$, x-candidates $\{59, 75\}$, solution $(59, -54)$, and raw target $25\theta + 109$ for $p = 139$, $\ell = 35$, $d = 140$. EMPIRICAL: six curves, all 82 raw targets The pullback identity and the Satoh-based FAPI-1 raw inverse both passed exhaustively, testing at most four candidate points per target (Figure 3); and on the same 82 **reduced** targets the canonical final root in μ_r was Miller-compatible 0 times — fast raw MI leaves the final-fibre representative-selection obstruction fully intact.

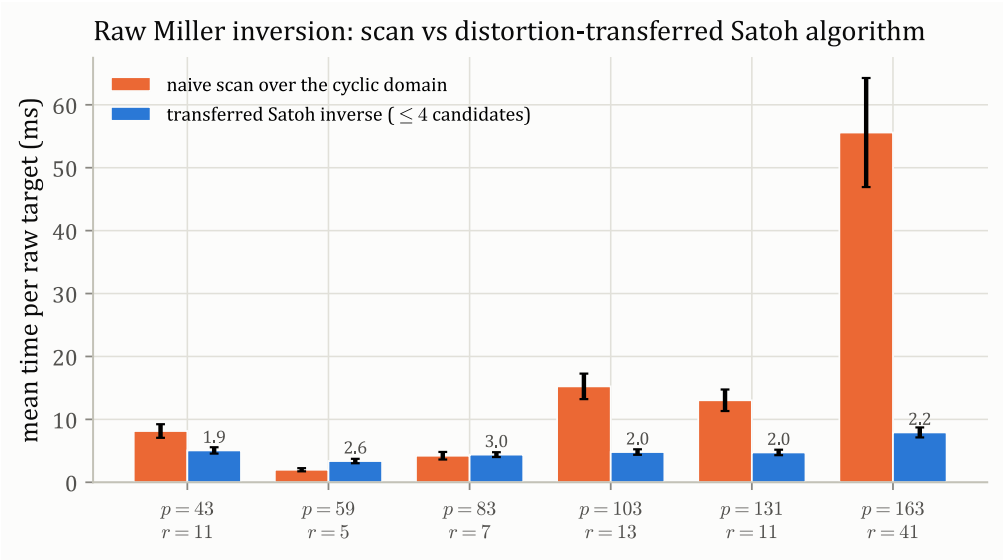

Figure 3: EMPIRICAL: six curves, 50 seeded timings each Raw Miller inversion: the naive cyclic-domain scan against the distortion-transferred Satoh inverse (error bars 95% CI; the number over each blue bar is the mean count of candidate points tested, always at most four). The Satoh route is flat in r where the scan grows. Data: `reproduce_satoh_mi_p43-59-83-103-131-163_20260627.csv` and the stage-measurement CSV.

6 Where the natural reduction stops

The obvious route from $\mathbb{G}_2$-ECDLP hardness to FAPI-1 hardness would be a reduction that turns a FAPI-1 solver into a discrete-log solver, or dually a proof that an ECDLP oracle solves FAPI-1. The second direction fails at a precise point, and the failure is structural.

Proposition 7 (the missing capability). PROVED A $\mathbb{G}_2$-ECDLP oracle alone does not implement the natural FAPI-1 route: the target $z \in \mathbb{G}_T$ must first be mapped back to some $\mathbb{G}_2$ element before that oracle has a valid input, and constructing precisely that cross-group inverse **is** FAPI-1.

Proposition 8 (exact equivalence relative to source DLP). PROVED With fixed generators and a $\mathbb{G}_2$-DLP oracle, FAPI-1 and DLP in $\mathbb{G}_T$ are polynomial-time equivalent: given $g_T = e(P, Q)$, target DLP computes $c = \log_{g_T} z$ and returns $[c]Q$, while FAPI-1 followed by source DLP returns the target scalar. Source DLP alone lacks exactly the target-to-source step.

The reduction attempt thus stops at a typed interface: nothing in the pairing API maps $\mathbb{G}_T$ back to a source group. The right formal home for that observation is an oracle model in which the interface is all there is — and in that model the separation can be proved outright.

7 The typed generic oracle and the separation theorem

7.1 Model

CITED Shoup's random-representation (RR) model exposes group elements through a random injective bit-string encoding and charges oracle group operations [4]. Maurer's hidden-state model expresses generic DLP values as affine forms and obtains information only through collisions [5]. Zhandry distinguishes the RR and type-safe (TS) models and proves their security notions coincide for single-stage games that exist in both models ([6], Theorems 1.5–1.6). PROVED FAPI-1 is a single-

stage search game, and the typed extension used here — three element sorts, a bilinear gate $\mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$, and a DLP gate only from $\mathbb{G}_2$ to scalar bits — exists in both models; the proof below is written for RR and remains valid in the corresponding type-safe game.

Definition 9 (typed bilinear oracle with one-sided DLP). For prime r and label length $L \geq 3\lceil \log_2 r \rceil$, sample independently random injections $\sigma_1, \sigma_2, \sigma_T$ from three copies of $\mathbb{Z}_r$ into disjoint typed L-bit label spaces. The oracle exposes typed group laws, equality, generators, the pairing

$$e(\sigma_1(a), \sigma_2(b)) = \sigma_T(ab), \tag{6}$$

and $\mathsf{DLOG}_2(\sigma_2(b)) = b$. Invalid or cross-typed strings return $\perp$. There is **no** target-to-source operation. Public data may include p, r, and the curve equation, but no coordinate-to-label map.

7.2 Elliptic-curve realization

The labels are not encodings of abstract cyclic groups: they encode actual curve groups, so the separation cannot be dismissed as concerning pairings that do not exist.

Lemma 10 (curve realization). PROVED For each security parameter $\lambda \geq 2$ there exist a prime $2^\lambda < r_\lambda < 2^{\lambda+1}$, a prime $p_\lambda \equiv -1 \pmod{4r_\lambda}$ of $O(\lambda)$ bits, and the curve $E_\lambda/\mathbb{F}_{p_\lambda} : y^2 = x^3 + x$ with $\#E_\lambda\big(\mathbb{F}_{p_\lambda}\big) = p_\lambda + 1$, an order-$r_\lambda$ point P, embedding degree two, and — via $\psi(x, y) = (-x, iy)$ in $\mathbb{F}_{p_\lambda^2}$ — a basis (P, Q) of $E[r_\lambda]$ with $Q = \psi(P)$ in the opposite Frobenius eigenspace whose Weil pairing value generates an order-r_λ target group.

Proof. Bertrand's postulate supplies r_λ. Dirichlet's theorem supplies primes in the class $-1 \bmod 4r_\lambda$, and a Linnik-type bound — with Xylouris's explicit fixed power [7] — bounds the least one by a fixed power of r_λ, hence $O(\lambda)$ bits. CITED The congruence gives $p_\lambda \equiv 3 \pmod 4$; pairing the quadratic-character contributions at x and $-x$ proves $\#E_\lambda\big(\mathbb{F}_{p_\lambda}\big) = p_\lambda + 1$, so an order-$r_\lambda$ point exists, and $p_\lambda \equiv -1 \pmod{r_\lambda}$ gives embedding degree two. The distortion-map and eigenspace facts are the standard $j = 1728$ supersingular structure, verified computationally on all six toy realizations. Oracle setup may store the selected prime and torsion data non-uniformly; the infinite oracle is the disjoint union of independently sampled components indexed by λ. □

7.3 Easy side and exact problem equivalence

PROVED ECDLP in $\mathbb{G}_2$ takes one DLOG_2 query. Moreover, relative to DLOG_2, FAPI-1 and target-group DLP are polynomial-time equivalent (Proposition 8 carries over verbatim to the oracle setting). The lower bound below therefore simultaneously separates FAPI-1 from $\mathbb{G}_2$-ECDLP and shows that, in this model, FAPI-1 retains the full generic $\Omega(\sqrt{r})$ hardness of target-group DLP.

7.4 The lower bound

Theorem 11 (elliptic-curve-backed generic oracle separation). PROVED There exists an oracle relative to which ECDLP in $\mathbb{G}_2$ is solvable with one query, while every probabilistic polynomial-time algorithm for FAPI-1 has negligible success. Quantitatively, at prime order r: give the algorithm $P = \sigma_1(1)$, $Q = \sigma_2(1)$, and $Z = \sigma_T(c)$ for uniform $c \in \mathbb{Z}_r$; it must output $\sigma_2(c)$. If it creates at most t target handles and makes q total queries, then

$$\Pr[\text{FAPI-1 success}] \leq \frac{\binom{t}{2} + 1}{r} + O(q/2^L). \tag{7}$$

Constant success requires $t = \Omega(\sqrt{r})$.

Proof. **Coupling.** Use Shoup's coupling in the form repaired by the audit (§9): first run a collision-free lazy simulator with fixed algorithm coins and fixed lazily sampled distinct labels, then choose c, and finally complete a uniformly random encoding consistent with the simulated labels whenever its formal values are distinct at c. For every fixed c this completion has exactly the distribution of a fresh uniform random encoding, so the coupled experiment has the same marginal distribution as independent uniform $(c, \sigma_1, \sigma_2, \sigma_T)$, except on the collision event charged below.

Affine invariant. Every target handle in the simulator has a formal value $\alpha X + \beta$ in the unknown X (the challenge exponent). Target operations preserve affinity. Source exponents may depend on earlier opaque label bits, but along the fixed simulated transcript they are constants, and no typed operation substitutes X into a source exponent; pairing queries therefore add only **constant** target forms.

Bad challenges. Let F be the final set of at most t distinct formal target forms on the simulated path. Each pair of distinct affine forms agrees at most once in $\mathbb{Z}_r$, so the bad set

$$R_F = \{c : \text{two forms in } F \text{ evaluate equally at } c\} \tag{8}$$

has size at most $B = \binom{t}{2}$. For every $c \notin R_F$, the completed encoding gives a real execution using exactly the simulator's path and labels, even when the algorithm branches arbitrarily on their bit patterns. (The encodings completed for different c need not coincide; the argument bounds the joint random-encoding experiment, not success over challenges for one already-fixed encoding.)

Registered outputs. On the collision-free path, every registered $\mathbb{G}_2$ output handle has one transcript-fixed exponent b; among $c \notin R_F$ it succeeds only when $c = b$, accounting for at most one further challenge. For every fixed simulator randomness, at most $B + 1$ of the r challenges succeed through registered handles; averaging preserves $(B + 1)/r$.

Blind labels. Any unregistered typed string supplied to a group, pairing, equality, or DLOG_2 interface is independent of the hidden encoding for its type. Hitting a particular challenge-dependent value such as $\sigma_1(c)$, $\sigma_2(c)$, or an affine target value has probability at most 2^{-L}; hitting some other valid label reveals only a transcript-fixed exponent and preserves the affine simulation. At most q blind typed inputs plus one unregistered output add $O(q/2^L)$.

Worst case. A bounded-error worst-case FAPI-1 solver would succeed with probability at least $2/3$ on every valid target, hence also on a uniform target, so the uniform-challenge bound excludes it; no average-case definition is substituted for the search problem. Cross-parameter queries cannot create a current-parameter handle and are independent of the current challenge; conditioning on their complete transcript leaves the per-parameter proof unchanged. □

Corollary 12 (one fixed oracle). PROVED There is a single fixed infinite oracle relative to which $\mathbb{G}_2$-ECDLP is polynomial-time and FAPI-1 is not solvable by any probabilistic polynomial-time machine.

Proof. Fix a machine M and let $S_{M,O}(\lambda)$ be its success for a fixed infinite oracle O, averaged over coins and the uniform challenge. The per-parameter bound gives $\mathbb{E}_O\big[S_{M,O}(\lambda)\big] \le \varepsilon_M(\lambda)$ with $\varepsilon_M(\lambda) = \mathrm{poly}(\lambda)/2^\lambda$. Put $\delta_M = \sqrt{\varepsilon_M}$; Markov's inequality gives $\Pr_O\big[S_{M,O}(\lambda) > \delta_M(\lambda)\big] \le \delta_M(\lambda)$, and $\sum_\lambda \delta_M(\lambda)$ converges, so by the first Borel–Cantelli lemma, with probability one over O only finitely many parameters violate $S_{M,O}(\lambda) \le \delta_M(\lambda)$. Since δ_M is negligible, M's success

is negligible for almost every fixed oracle. Probabilistic polynomial-time oracle machines form a countable set; the intersection of their probability-one good-oracle sets still has probability one, and any oracle in it works. □

Remark 13 (two randomness steps). PROVED The proof has two distinct randomness steps: first couple each challenge to a uniform encoding to obtain the query bound, then use Markov–Borel–Cantelli to select one fixed infinite oracle. Conflating these steps — holding one concrete encoding fixed while varying c, which does not preserve the initial challenge label — was the only genuine proof-level defect found during the audit, and its repair is exactly the coupling paragraph above.

8 Exhaustive verification of the collision core

The single load-bearing combinatorial fact in the theorem is the bad-set bound $|R_F| \le \min\left(p, \binom{t}{2}\right)$ for a set of t distinct affine forms over $\mathbb{Z}_p$. This is also the easiest place for a subtle error to hide, so it was checked exhaustively at small prime orders.

EMPIRICAL: p ∈ {5,7,11}, 2 ≤ t ≤ 4 The affine verifier checked 541,966 form sets exhaustively, plus 10,000 seeded sets at the one scale ($p = 11, t = 4$) whose 8,495,410 subsets exceeded the exhaustion limit, with zero violations. Every exhaustive row **attained** the applicable upper bound, so the check exercised tight collision patterns rather than only sparse cases (Table 3, Figure 4).

Table 3: EMPIRICAL: seed 2404 The affine-collision audit. Every exhaustive row attains its applicable bound and none exceeds it. Data: `verify_generic_oracle_bound_p5-7-11_t4_20260702.csv`.

p	t	mode	sets checked	bound $\min\left(p, \binom{t}{2}\right)$	max $\lvert R_F \rvert$	mean $\lvert R_F \rvert$	violations
5	2	exhaustive	300	1	1	0.833	0
5	3	exhaustive	2,300	3	3	2.283	0
5	4	exhaustive	12,650	5	5	3.765	0
7	2	exhaustive	1,176	1	1	0.875	0
7	3	exhaustive	18,424	3	3	2.439	0
7	4	exhaustive	211,876	6	6	4.224	0
11	2	exhaustive	7,260	1	1	0.917	0
11	3	exhaustive	287,980	3	3	2.611	0
11	4	seeded sample	10,000	6	6	4.758	0

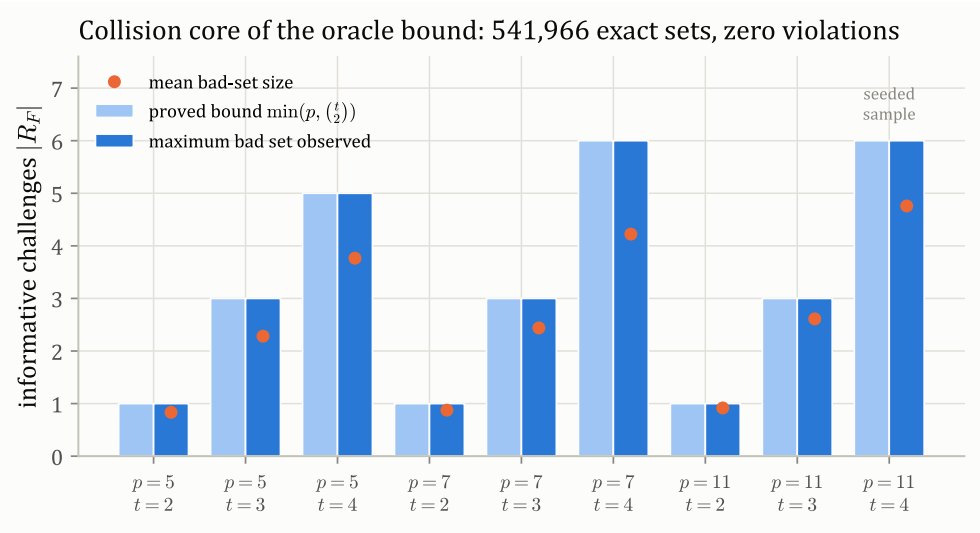

Figure 4: The proved union bound $\min\left(p, \binom{t}{2}\right)$ against the maximum and mean bad-set sizes over all checked affine-form sets. Observed maxima meet the bound exactly in every exhaustive row — the audit is tight, not vacuous.

9 Audit and adversarial self-verification

Because the separation is the paper's central claim, it was subjected to two independent hostile passes after being proved.

9.1 The A004 audit

The audit targeted four questions: **model** (RR, TS, or an undocumented hybrid?), **adaptivity** (are algorithms branching on opaque label bits covered?), **quantifiers** (does the ensemble bound really yield one fixed oracle?), and **realization** (is there an actual curve behind the labels?). CITED Primary sources were re-read for each: Shoup's random-encoding method [4], Maurer's affine collision argument ([5], §1.2), and Zhandry's RR/TS correspondence [6]. PROVED The audit's one substantive finding was the encoding/challenge conflation described in Remark 13; the repaired proof couples each challenge to a fresh uniform consistent encoding and only afterwards extracts a fixed oracle. All four questions closed, and the exact registered-handle bound $\left(\binom{t}{2} + 1\right)/r$ was confirmed.

9.2 The A005 self-verification

A second pass treated the audit itself as untrusted and attacked four failure surfaces: arbitrary unregistered typed labels, adaptive target-label branching, the ensemble-to-fixed-oracle quantifiers, and the curve realization.

- PROVED **One wording gap found and repaired:** the blind-label paragraph originally covered only DLOG_2 probes. Arbitrary unregistered strings can be fed to group, equality, pairing, or DLP interfaces; independent typed encodings keep the total contribution $O(q/2^L)$, so the theorem statement is unchanged.
- EMPIRICAL: deterministic rerun A separately generated affine CSV reproduced the recorded dataset byte-for-byte, SHA-256 `8A69F4ACE2D21F1AA34105603A295121A4DE8A6F7746FD54779BD3682EE2A042`.

- EMPIRICAL: six curves Independent recomputation verified $p \equiv -1 \pmod{4r}$, $\#E(\mathbb{F}_p) = p + 1$, $rP = \mathcal{O}$, pairing nonidentity, and target r-torsion for every measured realization.
- EMPIRICAL: negative control The exhaustive affine rows attain $\min\left(p, \binom{t}{2}\right)$, so the verifier would not pass merely because its collision sets were atypically sparse.
- EMPIRICAL: Python 3.13.4 After the correction, all 66 shared-library tests, all 7 P2.4 tests, and compile checks passed. One initial ad-hoc curve-check invocation failed before any arithmetic on a nonexistent placeholder import and wrote nothing; the corrected invocation passed all rows.

The resolved status survived both passes. No open P2.4 claim remains.

10 Limitations and open directions

PROVED **The scope boundary, stated exactly.** The theorem is an elliptic-curve-backed **generic bilinear oracle separation**. It is not a lower bound for algorithms given ordinary curve coordinates, a coordinate-to-label map, or full $\mathbb{F}_{p^2}$ arithmetic on target elements. Concrete target-field addition, coordinates, or a coordinate-to-label map can destroy the affine invariant on which the entire counting argument rests; no lower bound is claimed in that stronger model. A coordinate-exposing finite-field separation would be a strictly stronger new problem — with its own claim and sub-goals — not unfinished work in the stated theorem.

Three further honest limits. First, the toy timings of §4 are properties of one unoptimized Python implementation; the 147–545 ratio is an observation, not an asymptotic statement, and Satoh's algorithm shows how orientation alone can move the raw-MI cost class [2]. Second, the 0-of-82 canonical-root incompatibility tests a single deterministic root selector; it rules out that selector, not every FEI strategy, and on Galbraith–Hess–Vercauteren's larger Tate–Lichtenbaum domain a random root is provably useful with non-negligible probability [1]. Third, the Satoh transfer is proved for the supersingular $j = 1728$ degree-two family the repository measures; other families would need their own normalization identities.

The natural open direction is the model question the scope boundary raises: formulate the weakest **coordinate-exposing** target-field model in which either the affine invariant survives in some form, or an explicit algorithm exploits $\mathbb{F}_{p^2}$ structure to beat the generic bound for FAPI-1. Either outcome would be informative; the present theorem deliberately claims neither.

11 Conclusion

P2.4 asked for a polynomial-time equivalence between FAPI-1 and $\mathbb{G}_2$-ECDLP or an oracle separation. We delivered the separation, in a strong form: relative to an explicitly elliptic-curve-backed typed generic oracle with a free $\mathbb{G}_2$-DLP gate, source ECDLP is one query while FAPI-1 retains full generic square-root hardness, $\left(\binom{t}{2} + 1\right)/r + O(q/2^L)$, with one fixed oracle defeating all probabilistic polynomial-time machines (Theorem 11, Corollary 12). The staged analysis explains the mechanism: final exponentiation is a regular d-to-one map whose fibres contain exactly one Miller-compatible representative (Lemmas 1–2), the fixed-argument Miller function has exact factor structure $(r - 2, r - 3)$ (Proposition 3), Satoh's algorithm — carried to the FAPI-1 orientation by the distortion identity (Lemma 4) — makes raw Miller inversion cheap on the measured family, and yet the representative-selection obstruction at the final-exponentiation interface remains untouched: 0 of 82 canonical roots were Miller-compatible. The missing capability in every reduction attempt is the same typed gap the oracle formalizes — nothing maps $\mathbb{G}_T$ back to

a source group. The collision core of the bound was verified exhaustively with tight rows, and the theorem survived an audit and a second adversarial self-verification. The claim stops exactly where the model stops: coordinate-exposing target-field lower bounds are outside it, by design and by statement.

Reproducibility

All experiments run at toy scale ($p \leq 163$) with seed 2404 on Python 3.13.4, using only the standard library and shared repository modules. Four scripts produce the four datasets used here: `measure_pairing_stages.py` (staged validation, exhaustive fibre enumeration, and inversion timings; 16.6 s for the six curves; `measure_pairing_stages_p43-59-83-103-131-163_20260624.csv`), `analyze_miller_function.py` (exact fixed-argument Miller expansion and the factor-degree sequence; `analyze_miller_function_p43_r11_20260624.csv`), `reproduce_satoh_mi.py` (Satoh Example 4.4 and the exhaustive FAPI-1 transfer; 3.0 s; `reproduce_satoh_mi_p43-59-83-103-131-163_20260627.csv`), and `verify_generic_oracle_bound.py` (the affine-collision audit; 0.9 s; `verify_generic_oracle_bound_p5-7-11_t4_20260702.csv`). The staged pairing is validated against Stögbauer's published $\mathbb{F}_{43}$ vector [3] and Satoh's published Example 4.4 [2]; 66 shared and 7 problem-specific tests pass. Every mathematical claim above carries one of the epistemic tags PROVED, CITED, EMPIRICAL: range exactly as in the research log; untagged sentences are exposition, not claims.

References

[1] S. D. Galbraith, F. Hess, and F. Vercauteren, "Aspects of pairing inversion," *IEEE Transactions on Information Theory*, vol. 54, no. 12, pp. 5719–5728, 2008, doi: 10.1109/TIT.2008.2006431.

[2] T. Satoh, "Miller inversion is easy for the reduced Tate pairing of embedding degree greater than one." IACR ePrint 2019/385, revised January 2025, 2025.

[3] M. Stögbauer, "Efficient Algorithms for Pairing-Based Cryptosystems," Diploma thesis, Technische Universität Darmstadt, 2004.

[4] V. Shoup, "Lower bounds for discrete logarithms and related problems," in *EUROCRYPT 1997*, in LNCS, vol. 1233. 1997, pp. 256–266.

[5] U. Maurer, "Abstract models of computation in cryptography," in *IMA Cryptography and Coding 2005*, in LNCS, vol. 3796. 2005, pp. 1–12.

[6] M. Zhandry, "To label, or not to label (in generic groups)," in *CRYPTO 2022*, in LNCS, vol. 13509. 2022, pp. 66–96. doi: 10.1007/978-3-031-15982-4_3.

[7] T. Xylouris, "On the least prime in an arithmetic progression and estimates for the zeros of Dirichlet L-functions," *Acta Arithmetica*, vol. 150, no. 1, pp. 65–91, 2011.