Barriers for Transparent Algebraic Generation of Hard Supersingular Curves

Anis Bkakria

IRT SystemX, France

Abstract. We study transparent public generation of hard supersingular curves: a public, seeded, rerunnable algorithm outputs a supersingular curve while exposing the seed, verification transcript, and all algebraic information reconstructible from the implementation. This setting is distinct from trusted or distributed ceremonies, where a witness may be hidden, erased, or zero-knowledge protected. We define a transcript-security model for this setting and develop barriers for several modeled algebraic generation routes. For modular-conjugacy samplers, we analyze loci of supersingular invariants admitting an isogeny to their Frobenius conjugate; small public parameters give efficiently recognizable sparse support, while extractable high-parameter witnesses yield non-scalar endomorphisms by Frobenius composition. We extend this leakage principle to bounded correspondence witnesses, including kernel, rational-map, Hecke–Brandt, and elliptic-return detour traces. For direct samplers, we separate predicate-only search, local-neighbor exploration, bounded-relation witnesses, low-formal-degree ambient algebraic maps, and degree-one indexed maps from $\mathbb{F}_p$ to $\mathbb{F}_{p^2}$. The resulting theorems are deliberately model-bounded: they do not rule out hidden-witness ceremonies, higher-degree indexed maps, compact high-degree root indexing, endpoint-only mechanisms, or nonlocal aggregate sampling. The formal proofs are collected in the appendices.

Table of Contents

1 Introduction . . . 4
2 Technical Overview . . . 5
2.1 Transparent traces versus hidden witnesses . . . 5
2.2 Small modular-conjugacy loci . . . 5
2.3 High-degree witnesses leak OneEnd . . . 5
2.4 Direct sampler regimes . . . 6
2.5 Residual primitives . . . 6
3 Preliminaries . . . 6
3.1 Fields, curves, and Frobenius twists . . . 7
3.2 Isogenies, paths, and endomorphism witnesses . . . 7
3.3 Admissible represented witnesses . . . 7
3.4 Modular polynomials and conjugacy loci . . . 7
3.5 Projective level covers and finite-field Chebotarev . . . 7
4 Transparent Generation Model . . . 8
4.1 Public seeded generators and traces . . . 8
4.2 Leakage experiments . . . 8
4.3 Extractability interfaces . . . 9
4.4 Scope of the model . . . 10
5 Modular-Conjugacy Barriers . . . 10
5.1 Modular-conjugacy loci . . . 10
5.2 The counting statement . . . 10
5.3 Distributional distinguisher . . . 11
5.4 Interpretation for transparent generation . . . 11
6 Extractable High-Degree and Bounded-Correspondence Barriers . . . 11
6.1 Extractable conjugacy witnesses . . . 11
6.2 Bounded correspondence witnesses . . . 12
6.3 Classification of high-parameter routes . . . 13
7 Direct Classical Sampler Barriers . . . 13
7.1 Predicate-only search . . . 13
7.2 Local-neighbor exploration . . . 13
7.3 Bounded algebraic relations . . . 14
7.4 Low-degree ambient algebraic maps . . . 14
7.5 Degree-one indexed maps . . . 14
7.6 What remains . . . 15
8 Residual Routes and Non-Covered Mechanisms . . . 15
8.1 Root-only high-degree indexing . . . 15
8.2 Nonlocal aggregate sampling . . . 15
8.3 Endpoint-only higher-dimensional detours . . . 16
8.4 Ceremonies and hidden witnesses . . . 16
8.5 Residual-route summary . . . 16
9 Related Work . . . 16
9.1 Hard supersingular-curve generation . . . 16
9.2 Conjugate isogenies, level structure, and small structured loci . . . 17
9.3 Endomorphism rings, paths, and one endomorphism . . . 17
9.4 Modular polynomials and explicit algebraic computation . . . 17
9.5 Post-SIDH context and other isogeny directions . . . 17
9.6 Arithmetic tools used in the proofs . . . 17
10 Limitations and Open Problems . . . 17
11 Conclusion . . . 18
References . . . 18
A Proof-Coverage Map . . . 21
B Proofs for the Transparent-Trace Leakage Interface . . . 21
B.1 Frobenius composition and OneEnd leakage . . . 21
B.2 Bounded correspondences and route classification . . . 22
C Proofs for Modular-Conjugacy Barriers . . . 22
C.1 Imported arithmetic inputs . . . 22
C.2 Normal and non-normal witnesses . . . 23

C.3 Small-level support and distinguisher . . . 24
D Proofs for Direct Sampler Barriers . . . 25
D.1 Predicate, local-neighbor, and bounded-relation proofs . . . 25
D.2 Low-degree ambient sparse-hit bound . . . 25
E Affine-Line Monodromy and Chebotarev Proof . . . 26
E.1 Imported modular-cover and Chebotarev inputs . . . 26
E.2 Typed branch loci and no accidental diagonal quotients . . . 26
E.3 Frobenius-stable descent and coset density . . . 27
E.4 Chebotarev optimization and indexed-sampler corollary . . . 28

1 Introduction

Transparent hard-curve generation. Many isogeny-based protocols rely, explicitly or implicitly, on supersingular elliptic curves for which no party knows an endomorphism ring or an isogeny path from a curve with known endomorphism ring. Producing such curves without a trusted party is a recurring open problem. Booher et al. formulate this problem and document several failed attempts to hash into supersingular isogeny graphs without revealing an endomorphism ring or a path to a known-ring curve, including iterative root finding for the supersingular polynomial, specialized modular-polynomial gcds, division-polynomial systems, abelian-surface walks, and quantum random walks [Boo+22]. The present paper studies a sharper, deliberately adversarial variant: the generation algorithm is public, seeded, and rerunnable, and the adversary may inspect the transcript and algebraic trace generated by the implementation.

This transparent setting differs from trusted or distributed ceremonies. Basso et al. design a distributed trusted-setup ceremony for generating supersingular curves whose endomorphism rings are unknown to all parties [Bas+23]; Mokrani and Jao propose a multiparty protocol for generating supersingular curves over $\mathbb{F}_p$ with unknown endomorphism ring [MJ23]. Such protocols are outside our model because their security goal is precisely to hide, erase, distribute, or zero-knowledge-protect the witness. Our barriers instead target transparent algebraic procedures whose success witness remains extractable from a public rerunnable trace.

Why traces matter. The distinction between an output curve and its generation trace is essential. A random walk from a known curve may output a curve that appears hard when viewed alone, but the walk itself is a path witness. A modular-polynomial computation may output only a j-invariant, but the public computation may reveal a kernel, a rational map, an isogeny to a Frobenius conjugate, or a small structured algebraic locus. In the transparent model, such information is not an implementation detail; it is adversarial input.

The leakage criteria used in this paper are standard in the modern supersingular-isogeny landscape. Finding one non-scalar endomorphism is equivalent, under probabilistic polynomial-time reductions, to computing the full supersingular endomorphism ring [PW24]. Path-finding in supersingular isogeny graphs is likewise tightly related to endomorphism-ring computation [Wes22], and recent work further clarifies unconditional reductions between isogeny, endomorphism-ring, maximal-order, and related problems [MW25]. Consequently, a transparent sampler that reveals a path, orientation, represented correspondence, or non-scalar endomorphism witness should not be considered transcript-secure.

Our contributions. We give a theorem-level barrier map for modeled transparent algebraic generation routes. The contributions are as follows.

1. **Transcript-security model.** We formalize public seeded supersingular samplers with rerunnable traces, and we define leakage experiments for OneEnd witnesses, known-root paths, extractable correspondences, and support distinguishability.
2. **Small modular-conjugacy barrier.** For
$$\mathcal{C}_N(p) = \{j \in \mathrm{SS}_p : \Phi_N(j, j^p) = 0\},$$
we show that small public parameters give efficiently recognizable sparse support. The proof uses a normal/non-normal dichotomy and an Eichler–Deuring optimal-embedding count; the full arithmetic derivation is placed in the appendices.
3. **Extractable high-parameter leakage.** If a public trace yields an admissibly represented prime-to-p isogeny $E \to E^{(p)}$, then Frobenius composition yields a non-scalar endomorphism of E. This turns extractable high-degree modular-conjugacy witnesses into OneEnd leakage.
4. **Bounded-correspondence barriers.** Kernel certificates, rational maps, Hecke–Brandt traces, and elliptic-return higher-dimensional detours are unsafe when their success witnesses are extractable as elliptic paths, Frobenius-conjugacy correspondences, or represented endomorphisms.
5. **Direct-sampler barriers.** We classify modeled direct samplers: predicate-only access requires linear-in-p work in the unstructured model; local-neighbor exploration leaks a path; bounded algebraic-relation witnesses reduce to the correspondence barrier; low-formal-degree ambient algebraic maps hit SS_p with probability only $O(\Delta/p)$; and degree-one indexed maps $t \mapsto a + bt$ are blocked by an affine-line nonconcentration theorem.

Scope. The paper is not a universal impossibility theorem for classical sampling. The residual routes are explicit: hidden-witness ceremonies, higher-degree indexed maps $J : \mathbb{F}_p \to \mathbb{F}_{p^2}$, compact high-degree root indexing, endpoint-only higher-dimensional detours, and nonlocal aggregate sampling remain outside the proved barriers. This scope discipline is intentional. The goal is to identify which transparent algebraic routes collapse to leakage or small support, and which routes require genuinely new primitives.

Paper organization. Section 3 fixes notation and the algebraic conventions used throughout the paper. The later sections define the transcript-security model, prove the modular-conjugacy and extractability barriers, classify direct samplers, and isolate residual routes. All formal proofs are placed in the appendices.

2 Technical Overview

This section gives the proof architecture without proofs. The paper studies public generation algorithms whose seed, transcript, and reconstructible algebraic trace are available to the adversary. The objective is not to rule out every conceivable classical sampler, but to identify the modeled transparent algebraic routes that either reveal a hard witness or land in a distinguishable support.

2.1 Transparent traces versus hidden witnesses

A hard supersingular curve is useful only if the generation process does not reveal enough information to solve a nearby hard problem. In a trusted or distributed ceremony, the witness may be erased, hidden across parties, or protected by zero-knowledge arguments. In a transparent public generator, by contrast, the adversary can rerun the algorithm on the public seed and inspect the algebraic objects generated along the way. The same output curve can therefore be secure in one model and insecure in another.

The central modeling choice is to treat the transcript as part of the public output. A transcript may contain an explicit path from a known curve, a kernel polynomial, an isogeny representation, a modular-polynomial gcd trace, a Brandt or Hecke edge trace, or an endpoint returned by a higher-dimensional detour. If such information is extractable in polynomial time, then the generator has not merely sampled a curve; it has also published a witness about that curve. Our results classify when that witness is already enough to win a path, OneEnd, correspondence, or support-distinguishing experiment.

2.2 Small modular-conjugacy loci

The first route asks for supersingular invariants satisfying

$$\Phi_N(j, j^p) = 0.$$

Equivalently, the output curve admits a cyclic degree-N isogeny to its Frobenius conjugate. This relation is attractive because it is algebraic and publicly testable. It is also dangerous: when N is small, the support is far smaller than the full supersingular set.

The counting argument separates two cases. Non-normal witnesses compose with Frobenius to produce small non-scalar endomorphisms, so they are confined to a small-endomorphism locus. Normal witnesses correspond to trace-zero orientation data satisfying an equation of the form

$$\mu^2 = -pN.$$

Those are counted through Eichler–Deuring optimal embeddings. The main body only states the resulting support bound and its distributional consequence; the arithmetic proof, including local embedding factors and the weighted-to-unweighted conversion, is placed in the appendix.

The cryptographic conclusion is distributional. A transparent sampler supported on

$$\bigcup_{N \le B} \mathcal{C}_N(p), \qquad B = \mathrm{poly}(\log p),$$

can be distinguished from a uniform-like supersingular sampler by testing membership in this union. Thus small public modular-conjugacy parameters cannot provide transcript-secure generation.

2.3 High-degree witnesses leak OneEnd

Large N changes the computational regime. The small-support argument no longer rules out a high-degree root-only solver. The relevant question becomes whether the public trace contains an extractable witness.

Suppose the trace yields a represented prime-to-p isogeny

$$\psi : E \to E^{(p)}.$$

Then composing ψ with the canonical Frobenius map back to E yields an endomorphism

$$\alpha : E \to E.$$

Its degree has odd p-adic valuation, hence cannot equal the square degree of scalar multiplication. Therefore α is a non-scalar endomorphism. This turns extractable high-degree conjugacy witnesses into OneEnd leakage.

The same principle applies beyond modular polynomials. If a division-polynomial system, kernel certificate, rational map, Hecke–Brandt trace, or elliptic-return higher-dimensional detour exposes an efficiently represented elliptic correspondence, then the transcript is not harmless auxiliary data. It is a witness that can be converted into path leakage, OneEnd leakage, or an explicitly named residual assumption. The paper phrases this as a modeled bounded-correspondence barrier rather than a universal impossibility statement.

2.4 Direct sampler regimes

The direct-sampler part asks what remains if the generator does not visibly produce a modular-conjugacy witness. We separate five modeled regimes.

First, a sampler with only predicate access to supersingularity is just rejection sampling from a hidden subset of size $\Theta(p)$ inside an ambient universe of size p^2. It needs linear-in-p trials in the unstructured model.

Second, local-neighbor exploration from a known-root curve records parent pointers and edge witnesses. If the transcript is public and rerunnable, the output comes with a path from the known root.

Third, bounded algebraic-relation witnesses reduce to the bounded-correspondence barrier whenever the relation is convertible into an elliptic path, kernel, rational map, or Frobenius-conjugacy correspondence.

Fourth, low-formal-degree ambient maps

$$J : \mathbb{F}_{p^2}^m \dashrightarrow \mathbb{F}_{p^2}$$

hit SS_p with probability only $O(\Delta/p)$, up to denominator exceptions. This is a formal-degree density statement, not a runtime lower bound: compact high-degree operations remain outside the claim.

Fifth, degree-one indexed maps

$$t \longmapsto a + bt, \qquad t \in \mathbb{F}_p,$$

are ruled out by affine-line nonconcentration. The proof uses projective level-ℓ modular covers, a Frobenius-stable descent for non-$\mathbb{F}_p$-rational lines, a no-accidental-diagonal quotient lemma, coset-level density, and finite-field Chebotarev. Because this is the most delicate arithmetic-monodromy part of the paper, the dedicated appendix gives the complete argument.

2.5 Residual primitives

The results leave explicit residual routes. A successful transparent sampler may still require a genuinely root-only primitive

$$\mathsf{ImplicitSSRootIndex},$$

which indexes supersingular roots without revealing a path, kernel, correspondence, or small structured support. Another possible residual is a nonlocal aggregate primitive

$$\mathsf{NonlocalBrandtAggregateSample},$$

which samples using global Brandt or Hecke information without materializing a path or a $\Theta(p)$-size object. Higher-degree indexed maps $J : \mathbb{F}_p \to \mathbb{F}_{p^2}$, compact high-degree circuits, endpoint-only higher-dimensional detours, and hidden-witness ceremonies lie outside the proved barriers.

This residual list is part of the contribution. It prevents the barrier results from being read as a global impossibility theorem and identifies the mechanisms that a positive construction must instantiate.

3 Preliminaries

This section fixes notation and conventions. It intentionally avoids long textbook background; only the objects used in the main statements are introduced here. The algebraic lemmas used later are proved in the appendices.

3.1 Fields, curves, and Frobenius twists

Let $p > 3$ be a prime and let $\mathbb{F}_{p^2}$ be the quadratic extension of $\mathbb{F}_p$. For an element $x \in \mathbb{F}_{p^2}$, write x^p for its p-power Frobenius conjugate. If $E/\mathbb{F}_{p^2}$ is an elliptic curve, then $E^{(p)}$ denotes its Frobenius twist obtained by applying p-power Frobenius to the defining coefficients. Thus

$$j(E^{(p)}) = j(E)^p.$$

Let $\mathrm{SS}_p \subseteq \mathbb{F}_{p^2}$ denote the set of supersingular j-invariants. We use the following convention throughout: Deuring correspondence is a geometric correspondence over $\overline{\mathbb{F}}_p$, while algorithmic representatives of supersingular j-invariants are taken in $\mathbb{F}_{p^2}$. This distinction matters because the paper counts geometric isomorphism classes but tests support membership using explicit $\mathbb{F}_{p^2}$-representatives.

3.2 Isogenies, paths, and endomorphism witnesses

An isogeny $\varphi : E \to E'$ is always a nonzero morphism of elliptic curves preserving the identity. Its degree is written $\deg \varphi$. The endomorphism ring of E over $\overline{\mathbb{F}}_p$ is denoted $\mathrm{End}(E)$. A *non-scalar endomorphism* is an element of $\mathrm{End}(E) \setminus \mathbb{Z}$. The OneEnd problem asks for one such element. The equivalence between OneEnd and EndRing, and the relation between path-finding and endomorphism-ring computation, motivate the leakage experiments used later [PW24; Wes22; MW25].

A path in an ℓ-isogeny graph is a sequence of small-degree isogenies between supersingular curves. A transparent generator leaks a path if the public seed and trace determine such a sequence from a known-root curve to the output curve. In this paper, path leakage is treated as a security failure, not as an efficiency artifact.

3.3 Admissible represented witnesses

The paper works with extractable algebraic witnesses. A represented isogeny or endomorphism is *admissible* if it supports polynomial-time evaluation, domain/codomain verification, composition with the canonical Frobenius map, degree checking, and non-scalar checking when the target is an endomorphism. The precise interface is defined in the model section and used by the leakage theorems.

The most important example is a represented prime-to-p isogeny

$$\psi : E \to E^{(p)}.$$

Composing ψ with the p-power Frobenius map back to E gives an endomorphism of degree $p \cdot \deg \psi$. When $p \nmid \deg \psi$, this degree is not a square, so the endomorphism cannot be scalar multiplication. The formal statement and proof are in the appendices.

3.4 Modular polynomials and conjugacy loci

Let $\Phi_N(X, Y) \in \mathbb{Z}[X, Y]$ be the classical modular polynomial. For $N \geq 2$ with $p \nmid N$, define the modular-conjugacy locus

$$\mathcal{C}_N(p) = \{j \in \mathrm{SS}_p : \Phi_N(j, j^p) = 0\}.$$

Equivalently, $j \in \mathcal{C}_N(p)$ when some elliptic curve with invariant j admits a cyclic degree-N isogeny to a curve with invariant j^p. These loci are the first family of transparent algebraic samplers studied in the paper.

The small-parameter analysis separates witnesses into non-normal and normal cases. Non-normal witnesses lead to small non-scalar endomorphisms. Normal witnesses lead to trace-zero orientation data and are counted using Eichler–Deuring optimal-embedding formulas. The arithmetic proof is given in the appendices.

3.5 Projective level covers and finite-field Chebotarev

The affine-line indexed barrier uses projective full-level covers of the punctured j-line. For $\ell \geq 5$, the projective level-ℓ cover over $\mathbb{A}^1_j \setminus \{0, 1728\}$ has geometric monodromy between

$$\mathrm{PSL}_2(\mathbb{F}_\ell) \quad \text{and} \quad \mathrm{PGL}_2(\mathbb{F}_\ell),$$

with modular-curve complexity polynomial in ℓ. We use Katz–Mazur for modular curves with level structure [KM85], Cojocaru–Hall for uniform large-image results over function fields [CH05], and Kowalski's large-sieve framework for finite-field Frobenius/Chebotarev estimates [Kow08].

The main body states the affine-line nonconcentration theorem. Its arithmetic-monodromy package–typed branch loci, Frobenius-stable descent, no accidental nonabelian diagonal quotient, coset-density, and Chebotarev optimization–is proved in the appendix.

4 Transparent Generation Model

This section defines the public-generation setting studied in the paper. The definitions are intentionally distributional: they specify the public seed, the generated curve, the transcript, and the algebraic trace available to an adversary. The experiments are written as main-body objects because they define the claims; all reductions and proofs appear in the appendices.

4.1 Public seeded generators and traces

A transparent generator is a public algorithm. Anyone can rerun it on the same seed, instrument its implementation, and inspect any algebraic object that is reconstructible from the run. This is the point at which our model differs from trusted or distributed ceremonies: in those protocols, a witness may be generated by hidden parties, zero-knowledge protected, erased, or unavailable to the public verifier.

Definition 1 (Public seeded supersingular generator). *A public seeded supersingular generator is a pair of polynomial-time algorithms*

$$\mathsf{Gen}(pp, \rho) \to (j, \tau), \qquad \mathsf{Verify}(pp, \rho, j, \tau) \to \{0, 1\},$$

where pp contains the prime p, field encodings for $\mathbb{F}_{p^2}$, curve encodings, validation rules, and any fixed public anchors; ρ is a public seed; $j \in \mathbb{F}_{p^2}$ is the output invariant; and τ is the public verification transcript. Correctness requires that, except with negligible probability over the seed distribution, $j \in \mathrm{SS}_p$ and $\mathsf{Verify}(pp, \rho, j, \tau) = 1$.

Definition 2 (Reconstructible trace). *Fix an implementation of* Gen. *The reconstructible trace*

$$\mathsf{Trace}_{\mathsf{Gen}}(pp, \rho)$$

is the deterministic algebraic execution record available by rerunning or instrumenting the public algorithm on input (pp, ρ). Depending on the sampler, the trace may include intermediate curves, modular-polynomial roots, kernel polynomials, rational maps, edge certificates, Brandt or Hecke data, CM orientations, or higher-dimensional correspondences. If an implementation computes an object by a public deterministic subroutine, that object belongs to $\mathsf{Trace}_{\mathsf{Gen}}$, even if it is not returned as part of τ.

The adversarial view in all experiments below is therefore

$$(pp, \rho, j, \tau, \mathsf{Trace}_{\mathsf{Gen}}(pp, \rho)).$$

This convention prevents a public sampler from treating intermediate algebraic data as private merely because it is not included in the official output.

4.2 Leakage experiments

The paper uses three failure modes. The first captures endomorphism leakage. It is motivated by the equivalence between OneEnd and EndRing and by known relations between path and endomorphism-ring problems [PW24; Wes22; MW25].

Experiment $\mathsf{Exp}^{\mathsf{one}}_{\mathsf{Gen},\mathcal{A}}(\lambda)$: OneEnd leakage.

1. The challenger samples public parameters pp containing a prime $p \approx 2^\lambda$, then samples a public seed ρ and computes $(j, \tau) \leftarrow \mathsf{Gen}(pp, \rho)$.
2. The adversary receives $(pp, \rho, j, \tau, \mathsf{Trace}_{\mathsf{Gen}}(pp, \rho))$.
3. The adversary outputs a represented endomorphism $\alpha : E_j \to E_j$.
4. The experiment outputs 1 iff $j \in \mathrm{SS}_p$, the representation verifies as an endomorphism of E_j, and $\alpha \in \mathrm{End}(E_j) \setminus \mathbb{Z}$.

The second failure mode captures explicit path leakage from a known-ring anchor. This is the transparent analogue of a public walk-and-forget attempt: if the walk can be rerun from the seed, then the path was not forgotten in the public-trace model.

Experiment $\mathsf{Exp}^{\mathsf{path}}_{\mathsf{Gen},\mathcal{A}}(\lambda)$: known-root path leakage.

1. Public parameters include a known-ring supersingular curve E_0, its endomorphism ring description, and a prescribed family of small-degree isogeny graphs.
2. The challenger samples ρ, computes $(j,\tau) \leftarrow \mathsf{Gen}(pp,\rho)$, and gives the adversarial view together with E_0 and $\mathrm{End}(E_0)$ to $\mathcal{A}$.
3. The adversary outputs a sequence of represented small-degree isogenies.
4. The experiment outputs 1 iff the sequence verifies as a path from E_0 to an elliptic curve with invariant j.

The third failure mode is distributional. A sampler can be insecure even without exposing an endomorphism or path if its outputs lie in an efficiently recognizable sparse subset of SS_p.

Experiment $\mathsf{Exp}^{\mathsf{dist}}_{\mathsf{Gen},\mathcal{D}}(\lambda)$: trace-aware distributional distinguishing.

1. The challenger samples a hidden bit $b \leftarrow \{0,1\}$.
2. If $b=0$, it samples ρ, computes $(j,\tau) \leftarrow \mathsf{Gen}(pp,\rho)$, and gives $(pp, j, \rho, \tau, \mathsf{Trace}_{\mathsf{Gen}}(pp,\rho))$ to $\mathcal{D}$.
3. If $b=1$, it samples $j \leftarrow U(\mathrm{SS}_p)$, sets $\rho, \tau, \mathsf{Trace}$ to the prescribed dummy symbols, and gives $(pp, j, \rho, \tau, \mathsf{Trace})$ to $\mathcal{D}$.
4. The distinguisher outputs b', and the experiment outputs 1 iff $b'=b$.

Definition 3 (Uniform TS-cSRS distributional security). *A public generator* Gen *satisfies uniform transcript-secure cSRS distributional security if, for every polynomial-time distinguisher* $\mathcal{D}$*,*

$$\left|\Pr[\mathsf{Exp}^{\mathsf{dist}}_{\mathsf{Gen},\mathcal{D}}(\lambda)=1] - \frac{1}{2}\right|$$

is negligible in λ*.*

Definition 4 (Leakage security). *A public generator is OneEnd-leakage secure if every polynomial-time adversary wins* $\mathsf{Exp}^{\mathsf{one}}$ *with negligible probability. It is known-root-path secure if every polynomial-time adversary wins* $\mathsf{Exp}^{\mathsf{path}}$ *with negligible probability. A transcript-secure generator in this paper is expected to satisfy both leakage notions and the distributional notion of Definition 3.*

4.3 Extractability interfaces

The barriers are not triggered by the mere mathematical existence of a witness. They are triggered when the public trace contains enough information to extract and verify a witness efficiently. We therefore separate existence from extractability.

Definition 5 (Admissible represented isogeny). *An admissible represented isogeny from* E *to* E' *is a polynomial-size description* $\mathsf{rep}(\psi)$ *equipped with deterministic polynomial-time procedures for:*

1. *verifying the domain and codomain;*
2. *evaluating* ψ *on curve points or on the algebraic representation required by the subsequent reduction;*
3. *checking the degree* $D = \deg \psi$*;*
4. *composing* ψ *with represented isomorphisms and with the canonical Frobenius map.*

The representation may be a rational map, a kernel polynomial together with quotient data, a decomposed path of small isogenies, or another explicitly verified format.

Definition 6 (Admissible represented endomorphism). *An admissible represented endomorphism of* E *is a polynomial-size description* $\mathsf{rep}(\alpha)$ *equipped with deterministic polynomial-time procedures for verifying that the domain and codomain are both* E*, evaluating* α*, checking* $\deg \alpha$*, and certifying non-scalarity whenever the degree certificate is sufficient to rule out scalar multiplication.*

Definition 7 (PPT-extractable witness). *Let* $\mathcal{W}$ *be a witness relation, such as a path relation, a conjugacy-isogeny relation, or a correspondence relation. A public trace is PPT-extractable for* $\mathcal{W}$ *if there exists a polynomial-time extractor*

$$\mathsf{Ext}(pp, \rho, j, \tau, \mathsf{Trace}_{\mathsf{Gen}}(pp,\rho)) \to w$$

such that, with the stated probability over public seeds, w *verifies as an admissible witness for* $\mathcal{W}$*. The witness is* not *considered extractable merely because a witness exists existentially.*

4.4 Scope of the model

The model covers transparent algebraic samplers whose public implementation and seed determine the success witness, support test, or local exploration trace. It does not cover hidden-witness ceremonies, distributed generation with erased secrets, or zero-knowledge proofs that assert the existence of a witness without revealing an extractable representative. Those mechanisms move the problem into a different setup model and are discussed in the related-work and limitations sections.

The later theorems instantiate the model as follows. Small modular-conjugacy samplers fail Definition 3 because their support is efficiently recognizable and sparse. Extractable high-degree conjugacy witnesses fail $\mathsf{Exp}^{\mathsf{one}}$. Local-neighbor exploration from a known anchor fails $\mathsf{Exp}^{\mathsf{path}}$. Low-degree ambient and affine-line indexed samplers fail by distributional density or nonconcentration. The residual mechanisms are exactly those that avoid these modeled extractability and support conditions.

5 Modular-Conjugacy Barriers

This section states the first barrier family. A natural transparent route is to output a supersingular invariant that is linked to its Galois conjugate by a public cyclic isogeny relation. Such curves arise in higher-degree supersingular group actions, where one studies supersingular curves over $\mathbb{F}_{p^2}$ equipped with an isogeny to their Galois conjugate [CS22]. The level-structure and quaternionic viewpoint used to count these objects is part of the broader Deuring–Eichler dictionary for supersingular curves with additional structure [Arp24; Voi21]. We use the route in the opposite direction: as a public generation mechanism, small-level conjugacy support is efficiently recognizable and too sparse to look uniform.

Throughout this section, the counting and leakage proofs are given in the appendices. The main body records the exact objects and theorem statements needed by the cryptographic model.

5.1 Modular-conjugacy loci

Let $N \geq 1$ be prime to p. Write $\Phi_N(X, Y) \in \mathbb{Z}[X, Y]$ for the classical modular polynomial, so that away from the usual automorphism and characteristic exceptions, $\Phi_N(j_1, j_2) = 0$ records the existence of a cyclic degree-N isogeny between elliptic curves with invariants j_1 and j_2. For the small levels used in the distinguishers below, the needed modular-polynomial evaluations are efficient using standard algorithms for computing or evaluating modular polynomials [BLS12; Sut13].

Definition 8 (Modular-conjugacy locus). *For $N \geq 1$ with $p \nmid N$, define*

$$\mathcal{C}_N(p) = \{j \in \mathrm{SS}_p : \Phi_N(j, j^p) = 0\}.$$

For a bound B, define the small-level union

$$\mathcal{C}_{\leq B}(p) = \bigcup_{1 \leq N \leq B,\, p \nmid N} \mathcal{C}_N(p).$$

The set $\mathcal{C}_N(p)$ is public and efficiently testable when N is small. Therefore a generator supported on $\mathcal{C}_{\leq B}(p)$ can be distinguished from uniform supersingular sampling whenever the union has negligible mass inside SS_p.

5.2 The counting statement

The counting argument separates two cases. Non-normal witnesses induce non-scalar endomorphisms of degree exactly N^2 by the scalar-degeneracy argument, and are bounded by small-endomorphism rarity. Normal witnesses yield trace-zero orientations satisfying a quadratic relation of the form $\mu^2 = -pN$, and are counted by an Eichler–Deuring optimal-embedding formula. The local embedding factors, the weighted-to-unweighted conversion, and the multiplicity from $\pm\mu$ are handled in the appendix.

Theorem 1 (Small-level modular-conjugacy support). *There are absolute constants $c_1, c_2 > 0$ and an arithmetic factor $\kappa(N) = N^{o(1)}$ such that, for every prime $p > 3$ and every $N \geq 1$ with $p \nmid N$,*

$$|\mathcal{C}_N(p)| \leq c_1 N^3 \kappa(N) + c_2 N \kappa(N) p^{1/2}.$$

Consequently, for every $B = \mathrm{poly}(\lambda)$ and every $p \asymp 2^\lambda$,

$$|\mathcal{C}_{\leq B}(p)| \leq p^{1/2}\,\mathrm{poly}(\lambda),$$

whereas $|\mathrm{SS}_p| = \Theta(p)$. Hence

$$\frac{|\mathcal{C}_{\leq B}(p)|}{|\mathrm{SS}_p|} \leq \mathrm{negl}(\lambda).$$

Proof. See Section C.3.

Remark 1 (What the bound does and does not say). Theorem 1 is a support bound for small public levels. It does not rule out a high-level root-only solver that reveals no isogeny, kernel, path, or correspondence witness. Such a solver is treated later as a residual primitive rather than as a consequence of the small-level theorem.

5.3 Distributional distinguisher

The distinguisher is conceptually simple: test whether the output lies in the efficiently recognizable union $\mathcal{C}_{\leq B}(p)$. The rejection probability under the sampler may be high or low depending on the generator, but under a uniform supersingular sample the event has negligible probability by Theorem 1.

Theorem 2 (Small-parameter modular-conjugacy distinguisher). *Let $B(\lambda) = \mathrm{poly}(\lambda)$. Let* Gen *be a public seeded generator whose output distribution is supported on $\mathcal{C}_{\leq B}(p)$, except with negligible probability. Then* Gen *does not satisfy uniform TS-cSRS distributional security in the sense of Definition 3. More precisely, there exists a polynomial-time trace-aware distinguisher $\mathcal{D}_B$ with advantage*

$$\frac{1}{2} - \mathrm{negl}(\lambda)$$

against uniform supersingular sampling.

Proof. See Section C.3.

5.4 Interpretation for transparent generation

The result blocks a broad but specific route: choose a small public level N, solve or search for roots of

$$\Phi_N(j, j^p) = 0,$$

and publish the resulting supersingular invariant as if it were uniform or hard. The output support is then public, efficiently testable, and negligible inside SS_p. Thus the sampler fails distributionally before one asks whether it leaks an explicit path or endomorphism.

The high-level case is different. If the public trace contains an admissibly represented degree-N isogeny $E \to E^{(p)}$, then the next section shows that Frobenius composition yields OneEnd leakage. If the high-level method reveals only an endpoint root and no extractable witness, then it is not covered by Theorem 2; it becomes an explicit residual route.

6 Extractable High-Degree and Bounded-Correspondence Barriers

The small-level result of Section 5 is a distributional barrier. It says that transparent samplers supported on small modular-conjugacy loci are distinguishable from uniform supersingular sampling. Large algebraic parameters require a different argument. A high-degree equation may have many solutions, and root-only access to one solution is not by itself a leakage certificate. The relevant question is whether the public trace contains an efficiently represented witness.

This section states the witness-leakage barriers. The full reductions appear in the proof appendices. The main mechanism is simple: an extractable prime-to-p correspondence from a curve to its Frobenius conjugate can be composed with Frobenius to obtain a non-scalar endomorphism. Thus a transparent generator that publishes such a witness fails OneEnd leakage security.

6.1 Extractable conjugacy witnesses

Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve with invariant j. Write $E^{(p)}$ for its Frobenius twist. A public trace may produce a represented isogeny to $E^{(p)}$ directly, or it may produce a represented isogeny to a curve C with $j(C) = j^p$ together with a represented isomorphism $C \simeq E^{(p)}$. Both cases are covered by the following theorem.

Theorem 3 (Extractable Frobenius-conjugacy witnesses leak OneEnd). *Let* Gen *be a public seeded supersingular generator. Suppose that, with non-negligible probability over the public seed, the reconstructible trace yields:*

1. *a supersingular curve $E/\mathbb{F}_{p^2}$ with output invariant j;*

2. *an admissibly represented isogeny* $\psi : E \to C$ *of degree* D *with* $p \nmid D$*;*
3. *a represented isomorphism* $u : C \to E^{(p)}$*.*

Then Gen *is not OneEnd-leakage secure in the sense of Definition 4. More precisely, there is a polynomial-time adversary that outputs an admissibly represented non-scalar endomorphism of* E *with non-negligible probability.*

Proof. See Section B.1.

The endomorphism produced in the proof is the Frobenius-composition witness

$$\alpha = F_{E^{(p)}} \circ u \circ \psi : E \to E,$$

where $F_{E^{(p)}} : E^{(p)} \to E$ is the canonical p-power Frobenius map over $\mathbb{F}_{p^2}$. Its degree is pD, which cannot be the square degree of scalar multiplication because $p \nmid D$. The appendix spells out the representation interface, target-isomorphism ambiguity, and exceptional automorphism cases.

Corollary 1 (High-degree modular-conjugacy traces). *Let* N *be arbitrary with* $p \nmid N$*. A transparent sampler that solves*

$$\Phi_N(j, j^p) = 0$$

and whose public trace contains an admissibly represented cyclic degree-N *isogeny from* E_j *to a curve isomorphic to* $E_j^{(p)}$ *fails OneEnd-leakage security.*

Proof. See Section B.1.

6.2 Bounded correspondence witnesses

The preceding theorem is stated for a direct Frobenius-conjugacy isogeny. Many transparent algebraic routes expose a slightly different object: a kernel polynomial, a rational map, a decomposed path, a Hecke or Brandt trace, a CM orientation, or a higher-dimensional detour that returns to an elliptic curve. The following proposition is the modeled bridge from such traces to the leakage experiments of Section 4.

Definition 9 (Extractable elliptic correspondence trace). *An extractable elliptic correspondence trace for output* E *is a public trace from which a polynomial-time extractor obtains at least one of the following admissible objects:*

1. *a represented non-scalar endomorphism of* E*;*
2. *a represented path from a known-ring public anchor to* E*;*
3. *a represented prime-to-*p *isogeny from* E *to a curve isomorphic to* $E^{(p)}$*;*
4. *a represented elliptic correspondence that is polynomial-time convertible into one of the preceding objects.*

The definition is intentionally extractability-based. It does not apply to an existential witness that remains hidden, erased, or zero-knowledge protected.

Proposition 1 (Modeled bounded-correspondence barrier). *Let* Gen *be a public generator whose successful outputs have, with non-negligible probability, a PPT-extractable elliptic correspondence trace in the sense of Definition 9. Then at least one of the following holds:*

1. Gen *fails OneEnd-leakage security;*
2. Gen *fails known-root-path security;*
3. *the trace is not an extractable elliptic correspondence and the sampler relies on an explicit residual primitive.*

Proof. See Section B.2.

The proposition covers the transparent versions of the failed-route families identified in hard-curve generation work: division-polynomial systems, kernel certificates, modular-polynomial gcd traces, explicit Hecke or Brandt edge data, and abelian-surface detours that return an extractable elliptic witness [Boo+22]. It does not say that these mathematical ideas are useless in hidden-witness ceremonies. It says that, in the transparent-trace model, an extractable elliptic witness is already a leakage object.

6.3 Classification of high-parameter routes

The high-parameter regime should therefore be read as a classification rather than as a universal impossibility theorem.

Proposition 2 (High-parameter transparent-route classification). *Consider a transparent algebraic sampler whose success condition uses large modular-polynomial, division-polynomial, Hecke, Brandt, or higher-dimensional data. Within the extractability model of Definitions 7 and 9, every successful route falls into one of the following classes:*

1. *an extractable Frobenius-conjugacy witness, which leaks OneEnd by Theorem 3;*
2. *an extractable path from a known public anchor, which leaks a known-root path;*
3. *an extractable orientation, endomorphism, or correspondence that is convertible into one of the preceding witnesses;*
4. *a root-only, endpoint-only, compact high-degree, or nonlocal aggregate mechanism outside the modeled extractability interface.*

Proof. See Section B.2.

This classification is deliberately conservative. The last case is not treated as impossible; it is promoted to a residual design requirement. A positive transparent construction must instantiate such a residual mechanism without producing an extractable path, endomorphism, or Frobenius-conjugacy correspondence.

7 Direct Classical Sampler Barriers

The preceding sections address samplers whose traces contain modular-conjugacy or correspondence witnesses. A direct sampler may try to avoid those witnesses entirely: it may query a predicate, explore the graph locally, output roots through an implicit indexer, or compute an algebraic map into $\mathbb{F}_{p^2}$. This section states model-bounded barriers for the direct regimes that can be analyzed without assuming a hidden witness.

7.1 Predicate-only search

The most basic direct route samples candidates in $\mathbb{F}_{p^2}$ and tests supersingularity. Since $|\mathrm{SS}_p| = \Theta(p)$ inside an ambient universe of size p^2, unstructured predicate-only search is linear in p.

Proposition 3 (Predicate-only hidden-subset lower bound). *Let U be a universe of size p^2, and let $S \subset U$ be a hidden subset of size $\Theta(p)$. Any randomized algorithm that obtains information only by membership queries $x \mapsto 1_S(x)$ and makes $q = o(p)$ queries hits S with probability $o(1)$. In particular, predicate-only supersingularity search over $\mathbb{F}_{p^2}$ needs $\Omega(p)$ queries for constant success probability.*

Proof. See Section D.1.

The proof in the appendix uses the all-negative path for deterministic adaptive algorithms and then averages over the algorithm's coins. This avoids treating adaptive queries as independent events.

7.2 Local-neighbor exploration

Another direct route starts from a public known-ring curve and walks through small-degree neighbors until it reaches an output curve. This does not fit the predicate-only model, but it is unsafe in a public-trace setting because the rerunnable computation exposes the path.

Proposition 4 (Local-neighbor exploration leaks a path). *Let Gen be a public generator that starts from a public known-ring anchor E_0 and reaches its output by an explicit sequence of represented small-degree isogeny steps. If the sequence of intermediate curves, parent pointers, edge kernels, or equivalent edge certificates is reconstructible from $\mathsf{Trace}_{\mathsf{Gen}}(pp, \rho)$, then Gen fails known-root-path security.*

Proof. See Section D.1.

This statement is not a claim about private random walks. It is a claim about transparent walks whose public seed and implementation reconstruct the edge trace.

7.3 Bounded algebraic relations

A sampler may instead output a curve together with a bounded algebraic relation that is not syntactically a path. If that relation can be converted into an elliptic correspondence witness, then it is already covered by Proposition 1.

Proposition 5 (Bounded-relation reduction). *Let* Gen *be a public generator whose successful trace contains a bounded algebraic relation R involving the output curve. Suppose there is a polynomial-time converter that maps R to an extractable elliptic correspondence trace in the sense of Definition 9. Then* Gen *satisfies one of the leakage alternatives in Proposition 1.*

Proof. See Section D.1.

The proposition is intentionally conditional on the converter. Relations that are root-only, endpoint-only, or aggregate and do not yield a represented elliptic witness remain outside this barrier.

7.4 Low-degree ambient algebraic maps

A different route is to compute a low-degree rational map from a large seed space into $\mathbb{F}_{p^2}$ and hope that its image hits SS_p often. This is ruled out by a formal-degree density argument using Schwartz–Zippel [Sch80; Zip79].

Theorem 4 (Low-degree ambient sparse-hit bound). *Let*

$$J = N/D : \mathbb{F}_{p^2}^m \dashrightarrow \mathbb{F}_{p^2}$$

be a rational map of formal degree at most Δ, not constant with supersingular value. Then

$$\Pr_{X \leftarrow \mathbb{F}_{p^2}^m}[J(X) \in \mathrm{SS}_p] \le O\left(\frac{\Delta}{p}\right) + \Pr[D(X) = 0].$$

Consequently, if denominator exceptions occur with negligible probability and the map hits SS_p with constant probability, then $\Delta = \Omega(p)$.

Proof. See Section D.2.

The theorem is a formal-degree statement, not a runtime lower bound. A compact circuit can compute a high-degree function, for example through repeated squaring or Frobenius powering. Such compact high-degree behavior is therefore an explicit residual route.

Corollary 2 (Bounded-branch ambient samplers). *A public sampler represented by $B = \mathrm{poly}(\log p)$ ambient algebraic branches of formal degree $\Delta = \mathrm{poly}(\log p)$ cannot hit SS_p with non-negligible probability, except by denominator exceptions outside the model or by concentrating on a $\mathrm{poly}(\log p)$-size set of constant supersingular outputs.*

Proof. See Section D.2.

7.5 Degree-one indexed maps

The previous theorem does not cover true indexed maps whose seed space has size $\Theta(p)$. The first such case is a degree-one map

$$t \longmapsto a + bt, \qquad t \in \mathbb{F}_p, \qquad a, b \in \mathbb{F}_{p^2},\ b \neq 0.$$

Its image is an affine $\mathbb{F}_p$-line in $\mathbb{F}_{p^2}$. The relevant arithmetic input is affine-line nonconcentration for supersingular invariants.

Theorem 5 (Affine-line nonconcentration). *There exist constants $C > 0$ and $\delta > 0$ such that, for all sufficiently large primes $p > 3$ and all affine $\mathbb{F}_p$-lines*

$$L = a + b\mathbb{F}_p \subseteq \mathbb{F}_{p^2}, \qquad b \neq 0,$$

one has

$$|\mathrm{SS}_p \cap L| \le Cp^{1-\delta}.$$

Proof. See Section E.4.

The appendix proves this theorem through projective level-ℓ modular covers, Frobenius-stable descent for non-$\mathbb{F}_p$-rational lines, a typed-branch no-accidental-diagonal quotient lemma, coset-level bad-density bounds, and finite-field Chebotarev. The theorem is the only place in the paper where affine-line monodromy is used.

Corollary 3 (Degree-one indexed samplers fail). *Every degree-one indexed sampler*

$$t \leftarrow \mathbb{F}_p, \qquad j = a + bt, \qquad b \neq 0,$$

hits SS_p *with probability at most* $Cp^{-\delta}$*. In particular, no such sampler has constant success probability.*

Proof. See Section E.4.

7.6 What remains

The direct barriers leave several explicit residual mechanisms. They do not rule out a true root indexer

$$\mathsf{ImplicitSSRootIndex}(p, k) \to j \in \mathrm{SS}_p$$

that reveals no path, kernel, correspondence, or small structured support. They also do not rule out higher-degree indexed maps $J : \mathbb{F}_p \to \mathbb{F}_{p^2}$, compact high-degree circuits, endpoint-only higher-dimensional detours, or a nonlocal aggregate sampler

$$\mathsf{NonlocalBrandtAggregateSample}.$$

These residual mechanisms are the remaining positive design space for transparent generation.

8 Residual Routes and Non-Covered Mechanisms

The preceding sections rule out several transparent algebraic regimes, but they do not constitute a general impossibility theorem for public generation of hard supersingular curves. This section records the remaining routes explicitly. We state them as residual primitives because a proof that one of them is impossible would require new arithmetic or complexity lower bounds beyond the methods used here.

8.1 Root-only high-degree indexing

The small-level modular-conjugacy theorem covers levels $N \leq \mathrm{poly}(\lambda)$, while extractable high-degree witnesses leak endomorphism information. A residual route is a root-only high-degree solver that outputs

$$j \in \mathrm{SS}_p$$

from a high-degree relation, such as a modular-conjugacy equation, without revealing a kernel, correspondence, path, rational map, or certificate that can be converted into a leakage witness.

Definition 10 (Implicit supersingular-root indexing). *An* implicit supersingular-root indexer *is a public algorithm*

$$\mathsf{ImplicitSSRootIndex}(p, k) \to j \in \mathrm{SS}_p$$

that maps an index k to a supersingular invariant while revealing no extractable path, endomorphism, modular-conjugacy witness, bounded relation, or small-support membership certificate.

The paper gives barriers for low-degree ambient maps and degree-one subfield-indexed maps, but Definition 10 remains outside the theorem package. Addressing it requires either an arithmetic theorem for higher-degree indexed images

$$J : \mathbb{F}_p \dashrightarrow \mathbb{F}_{p^2}, \qquad \deg J > 1,$$

or a complexity-theoretic lower bound against compact high-degree algebraic circuits. Neither is claimed here.

8.2 Nonlocal aggregate sampling

Another residual route is a sampler that uses global spectral, Brandt, Hecke, or other aggregate information without materializing a path or local edge trace.

Definition 11 (Nonlocal Brandt aggregate sampling). *A* nonlocal Brandt aggregate sampler *is a public algorithm that samples a supersingular invariant from global Brandt, Hecke, or spectral information while revealing neither an explicit path from a known root nor enough edge data to reconstruct such a path.*

The bounded-correspondence barrier applies when the aggregate computation exposes edge witnesses, decomposed paths, kernel data, or an elliptic return certificate. It does not rule out a genuinely nonlocal primitive that samples endpoints without exposing any such witness. The distinction is important: explicit graph traces leak paths, while a compact aggregate sampler must avoid all local reconstruction.

8.3 Endpoint-only higher-dimensional detours

Booher et al. analyze several attempted approaches to hard-curve generation, including abelian-surface detours and higher-dimensional walks [Boo+22]. The bounded-correspondence barrier covers transparent detours that return an elliptic endpoint together with an extractable elliptic correspondence, decomposition, or certificate. It does not cover endpoint-only higher-dimensional procedures whose internal witness is hidden or erased.

Such procedures fall outside the transparent-trace model unless their public execution leaves a reconstructible algebraic trace. This is a genuine residual route rather than a contradiction: our claims are about public traces, not about hidden-witness ceremonies or zero-knowledge-protected computations.

8.4 Ceremonies and hidden witnesses

Trusted or distributed ceremonies deliberately hide the generation witness. Basso et al. give a distributed ceremony for producing supersingular curves whose endomorphism rings are unknown to all parties [Bas+23]; Mokrani and Jao propose a multiparty generation protocol for curves over $\mathbb{F}_p$ with unknown endomorphism ring [MJ23]. These works are complementary to the present paper. They change the setup model by using private randomness, interaction, erasure, proofs, or distributed trust; our barriers instead target public, rerunnable, transparent algorithms.

8.5 Residual-route summary

Route	Covered here?	Status
Extractable $E \to E^{(p)}$ witness	Yes	OneEnd leakage
Known-root local walk	Yes	path leakage
Low-degree ambient algebraic map	Yes	sparse-hit bound
Degree-one indexed map $a + bt$	Yes	affine-line nonconcentration
Higher-degree indexed map	No	residual
Root-only high-degree solver	No	residual
Nonlocal aggregate sampler	No	residual
Hidden-witness ceremony	No	outside model

This scope is part of the statement of the contribution. The contribution is a set of barriers for modeled transparent routes, together with a precise list of routes that remain open.

9 Related Work

We organize related work by the role it plays in the barrier. The comparison is deliberately model-sensitive: some prior works construct hidden-witness ceremonies or cryptographic protocols, whereas this paper studies public transparent traces.

9.1 Hard supersingular-curve generation

Closest to the motivation of this paper is the line of work on generating hard supersingular curves without revealing an endomorphism ring or a path from a known curve. Booher et al. document failed attempts to hash into supersingular isogeny graphs, including supersingular-polynomial root finding, modular-polynomial gcds, division-polynomial systems, abelian-surface walks, and quantum walks [Boo+22]. Their work studies why natural generation procedures tend to leak or fail; our contribution is a transcript-security model and a theorem package for transparent public algebraic traces.

Basso et al. take a different direction: they design a distributed trusted setup for producing supersingular curves whose endomorphism rings remain unknown [Bas+23]. Mokrani and Jao target multiparty generation of supersingular curves over $\mathbb{F}_p$ with unknown endomorphism ring [MJ23]. These ceremony-based works are not contradicted by our barriers because their witnesses are hidden, distributed, erased, or protected, while our adversary observes the public seed and reconstructible algebraic trace.

9.2 Conjugate isogenies, level structure, and small structured loci

The modular-conjugacy loci studied here are related to supersingular curves equipped with isogenies to their Galois conjugates. Chenu and Smith study higher-degree supersingular group actions arising from such conjugate-isogeny structures [CS22]. Arpin studies supersingular isogeny graphs with level structure and their quaternionic interpretation [Arp24]. Our use of modular-conjugacy loci is different: we treat them as public support sets for transparent generators and show that small-level support is distinguishable.

Small structured supersingular loci also appear in the study of curves with small non-integer endomorphisms. Love and Boneh identify and analyze such special subsets of supersingular curves over $\mathbb{F}_{p^2}$ [LB20]. This provides useful context for our non-normal branch: a transparent construction supported on curves with small endomorphism witnesses is not distributionally generic and may reveal hard information.

9.3 Endomorphism rings, paths, and one endomorphism

The leakage criteria in this paper are motivated by equivalences between central isogeny problems. Page and Wesolowski prove that computing one non-scalar endomorphism of a supersingular elliptic curve is equivalent, under probabilistic polynomial-time reductions, to computing the full endomorphism ring [PW24]. Wesolowski proves an equivalence between supersingular isogeny path finding and endomorphism-ring computation [Wes22]. Le Merdy and Wesolowski further develop unconditional foundations relating isogeny, endomorphism-ring, maximal-order, and related problems [MW25]. These results justify treating a transcript-recoverable path or non-scalar endomorphism as a serious leakage event.

9.4 Modular polynomials and explicit algebraic computation

Transparent modular-conjugacy samplers rely on modular-polynomial computation or evaluation. Bröker, Lauter, and Sutherland compute modular polynomials via isogeny volcanoes [BLS12]; Sutherland studies direct evaluation of modular polynomials [Sut13]. These works provide efficient algebraic tools for moderate levels. Our small-level theorem uses this algorithmic availability adversarially: if the support test is efficient and the support is sparse, the generator is distinguishable.

9.5 Post-SIDH context and other isogeny directions

The present paper is not about SIDH/SIKE-style key exchange, but the post-SIDH landscape matters for positioning. Castryck and Decru gave an efficient key-recovery attack on SIDH [CD22]; Maino and Martindale extended the attack surface to arbitrary starting curves [MM22]. These attacks exploit auxiliary structure specific to SIDH-style protocols. Our work instead studies transparent generation traces for supersingular curves.

Other isogeny directions remain active and structurally different. CSIDH instantiates a commutative group action on supersingular curves over $\mathbb{F}_p$ [Cas+18]. SQISign constructs compact signatures from quaternion and isogeny techniques [Feo+20]. These works use different assumptions and expose different auxiliary information. They are included here to clarify scope, not as targets of the barrier.

9.6 Arithmetic tools used in the proofs

The proof appendices rely on standard arithmetic inputs. Voight's text is used for quaternion algebras, Eichler orders, and optimal-embedding background [Voi21]. Katz and Mazur provide the foundation for modular curves and level structures [KM85]. Cojocaru and Hall provide uniform large-image results for elliptic curves over function fields [CH05]. Kowalski's large-sieve framework supplies the finite-field Frobenius/Chebotarev perspective used in the affine-line appendix [Kow08]. The low-degree ambient sampler barrier uses the Schwartz–Zippel polynomial identity method [Sch80; Zip79].

10 Limitations and Open Problems

The results in this paper are barriers for modeled transparent algebraic generation. They should not be read as a general impossibility theorem for all classical or quantum procedures that output supersingular curves. We isolate the adversarial information that is present in transparent, rerunnable generation: public seeds, public or reconstructible transcripts, and algebraic witnesses that can be extracted from those transcripts. The negative results apply to routes that expose such information or that fall into the stated low-degree and affine-indexed models.

Hidden-witness generation is outside the model. Trusted ceremonies, multiparty generation procedures, and zero-knowledge-protected setup protocols are not contradicted by our barriers. Those mechanisms may use private randomness, secret intermediate paths, or hidden witnesses that are erased or proved only in zero knowledge. Our model deliberately studies the complementary setting in which the generation procedure is public and rerunnable, so that the adversary can inspect or reconstruct the algebraic trace. Consequently, a protocol may avoid our leakage results by ensuring that all path, correspondence, orientation, and endomorphism witnesses remain hidden.

Root-only and aggregate primitives remain residual. Several residual mechanisms are intentionally left open. A sampler that realizes

$$\mathsf{ImplicitSSRootIndex}(p, k) \to j \in \mathrm{SS}_p$$

without revealing a path, an endomorphism, a modular-conjugacy witness, or a bounded algebraic relation is not ruled out. Similarly, a nonlocal Brandt or Hecke aggregate sampler that samples from a global distribution without materializing a path or a large explicit matrix lies outside this proof framework. Our contribution is to show that many natural transparent routes collapse to leakage or distinguishability; it does not prove that these named residual primitives cannot exist.

Higher-degree indexed maps are not closed. The affine-line theorem rules out degree-one subfield-indexed maps

$$t \mapsto a + bt, \qquad t \in \mathbb{F}_p.$$

It does not rule out rational maps

$$J(t) = A(t)/B(t) : \mathbb{F}_p \dashrightarrow \mathbb{F}_{p^2}$$

of degree larger than one. Extending the affine-line argument to bounded-degree rational images would require uniform control of projective level-cover monodromy after arbitrary low-degree pullback, including exceptional low-monodromy families and compact high-degree root-indexing escapes. This identifies the main mathematical boundary of the present barrier.

Formal degree is not running time. The low-degree ambient barrier is a formal-degree statement. It shows that a rational map from a large ambient seed space cannot hit SS_p with constant probability unless its formal degree, branch count, or exceptional behavior is large. It does not imply that every compact high-degree circuit is slow. Operations such as Frobenius powering may have high formal degree and small arithmetic circuit size. Compact high-degree root-indexing mechanisms therefore remain outside the formal-degree lower bound.

No implementation or side-channel claim. This paper is a foundations and barrier paper. It does not propose an implementable generator, does not benchmark a system, and does not claim constant-time behavior or side-channel resistance. Any construction that uses one of the residual mechanisms requires a separate implementation and leakage analysis.

Open problems. The main open problems are: prove or refute the existence of an efficient $\mathsf{ImplicitSSRootIndex}$ primitive; extend affine-line nonconcentration to bounded-degree rational images; formalize nonlocal Brandt/Hecke aggregate sampling and determine whether it can be implemented without path leakage; and understand whether compact high-degree algebraic operations can sample hard supersingular curves without revealing exploitable structure.

11 Conclusion

We studied transparent algebraic generation of hard supersingular curves in a model where the seed, transcript, and reconstructible algebraic trace are visible to the adversary. In this setting, several natural routes fail for structural reasons. Small modular-conjugacy loci are distinguishable from the supersingular set. Extractable high-degree conjugacy witnesses give non-scalar endomorphisms by Frobenius composition. Local walks and explicit graph traces leak paths. Bounded algebraic relation witnesses reduce to correspondence leakage. Low-formal-degree ambient maps hit the supersingular set too rarely, and affine-line indexed maps are ruled out by affine-line nonconcentration.

The resulting picture is not a universal impossibility theorem. It is a barrier map for modeled transparent algebraic mechanisms. The surviving routes are precisely those that avoid public extractable witnesses: root-only indexing, nonlocal aggregate sampling, compact high-degree operations, higher-degree indexed maps, endpoint-only higher-dimensional detours, and hidden-witness ceremonies. Instantiating any of these routes as a secure and efficient public generator is an open problem.

References

[Arp24] Sarah Arpin. "Adding Level Structure to Supersingular Elliptic Curve Isogeny Graphs". In: *Journal de Théorie des Nombres de Bordeaux* 36.2 (2024), pp. 405–443. DOI: 10.5802/jtnb.1283. URL: https://www.numdam.org/articles/10.5802/jtnb.1283/.

[Bas+23] Andrea Basso, Giulio Codogni, Deirdre Connolly, Luca De Feo, Tako Boris Fouotsa, Guido Maria Lido, Travis Morrison, Lorenz Panny, Sikhar Patranabis, and Benjamin Wesolowski. "Supersingular Curves You Can Trust". In: *Advances in Cryptology – EUROCRYPT 2023*. Vol. 14005. Lecture Notes in Computer Science. Springer, 2023, pp. 405–437. DOI: 10.1007/978-3-031-30617-4_14. URL: https://eprint.iacr.org/2022/1469.

[BLS12] Reinier Bröker, Kristin Lauter, and Andrew V. Sutherland. "Modular Polynomials via Isogeny Volcanoes". In: *Mathematics of Computation* 81.278 (2012), pp. 1201–1231. DOI: 10.1090/S0025-5718-2011-02508-1. URL: https://arxiv.org/abs/1001.0402.

[Boo+22] Jeremy Booher, Ross Bowden, Javad Doliskani, Tako Boris Fouotsa, Steven D. Galbraith, Sabrina Kunzweiler, Simon-Philipp Merz, Christophe Petit, Benjamin Smith, Katherine E. Stange, Yan Bo Ti, Christelle Vincent, José Felipe Voloch, Charlotte Weitkämper, and Lukas Zobernig. *Failing to Hash into Supersingular Isogeny Graphs*. Cryptology ePrint Archive, Paper 2022/518. 2022. URL: https://eprint.iacr.org/2022/518.

[Cas+18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. *CSIDH: An Efficient Post-Quantum Commutative Group Action*. Cryptology ePrint Archive, Paper 2018/383. 2018. URL: https://eprint.iacr.org/2018/383.

[CD22] Wouter Castryck and Thomas Decru. *An Efficient Key Recovery Attack on SIDH*. Cryptology ePrint Archive, Paper 2022/975. 2022. URL: https://eprint.iacr.org/2022/975.

[CH05] Alina Carmen Cojocaru and Chris Hall. "Uniform Results for Serre's Theorem for Elliptic Curves". In: *International Mathematics Research Notices* 2005.50 (2005), pp. 3065–3080. DOI: 10.1155/IMRN.2005.3065. URL: https://academic.oup.com/imrn/article-abstract/2005/50/3065/694981.

[CS22] Mathilde Chenu and Benjamin Smith. "Higher-Degree Supersingular Group Actions". In: *Mathematical Cryptology* 1.2 (2022), pp. 85–101. URL: https://inria.hal.science/hal-03288075/.

[Feo+20] Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski. "SQISign: Compact Post-Quantum Signatures from Quaternions and Isogenies". In: *Advances in Cryptology – ASIACRYPT 2020*. Vol. 12491. Lecture Notes in Computer Science. Springer, 2020, pp. 64–93. DOI: 10.1007/978-3-030-64837-4_3. URL: https://eprint.iacr.org/2020/1240.

[KM85] Nicholas M. Katz and Barry Mazur. *Arithmetic Moduli of Elliptic Curves*. Vol. 108. Annals of Mathematics Studies. Princeton University Press, 1985.

[Kow08] Emmanuel Kowalski. *The Large Sieve and Its Applications: Arithmetic Geometry, Random Walks and Discrete Groups*. Vol. 175. Cambridge Tracts in Mathematics. Cambridge University Press, 2008. URL: https://www.cambridge.org/core/books/large-sieve-and-its-applications/69A33E8682B2B989AAB7D56C8F6A209C.

[LB20] Jonathan Love and Dan Boneh. "Supersingular Curves with Small Noninteger Endomorphisms". In: *ANTS XIV—Proceedings of the Fourteenth Algorithmic Number Theory Symposium*. Vol. 4. The Open Book Series. Mathematical Sciences Publishers, 2020, pp. 7–22. DOI: 10.2140/obs.2020.4.7. URL: https://msp.org/obs/2020/4-1/p02.xhtml.

[MJ23] Youcef Mokrani and David Jao. *Generating Supersingular Elliptic Curves over $\mathbb{F}_p$ with Unknown Endomorphism Ring*. Cryptology ePrint Archive, Paper 2023/984. 2023. URL: https://eprint.iacr.org/2023/984.

[MM22] Luciano Maino and Chloe Martindale. *An Attack on SIDH with Arbitrary Starting Curve*. Cryptology ePrint Archive, Paper 2022/1026. 2022. URL: https://eprint.iacr.org/2022/1026.

[MW25] Arthur Herlédan Le Merdy and Benjamin Wesolowski. *Unconditional Foundations for Supersingular Isogeny-Based Cryptography*. arXiv:2502.17010. 2025. URL: https://arxiv.org/abs/2502.17010.

[PW24] Aurel Page and Benjamin Wesolowski. "The Supersingular Endomorphism Ring and One Endomorphism Problems Are Equivalent". In: *Advances in Cryptology – EUROCRYPT 2024*. Vol. 14656. Lecture Notes in Computer Science. Springer, 2024, pp. 388–417. DOI: 10.1007/978-3-031-58751-1_14. URL: https://eprint.iacr.org/2023/1399.

[Sch80] Jacob T. Schwartz. "Fast Probabilistic Algorithms for Verification of Polynomial Identities". In: *Journal of the ACM* 27.4 (1980), pp. 701–717. DOI: 10.1145/322217.322225. URL: https://dl.acm.org/doi/10.1145/322217.322225.

[Sut13] Andrew V. Sutherland. "On the Evaluation of Modular Polynomials". In: *The Open Book Series* 1.1 (2013), pp. 531–555. DOI: 10.2140/obs.2013.1.531. URL: https://msp.org/obs/2013/1-1/p26.xhtml.

[Voi21] John Voight. *Quaternion Algebras*. Vol. 288. Graduate Texts in Mathematics. Springer, 2021. DOI: 10.1007/978-3-030-56694-4. URL: https://link.springer.com/book/10.1007/978-3-030-56694-4.

[Wes22] Benjamin Wesolowski. "The Supersingular Isogeny Path and Endomorphism Ring Problems Are Equivalent". In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2022, pp. 1100–1111. DOI: 10.1109/FOCS52979.2021.00109. URL: https://arxiv.org/abs/2111.01481.

[Zip79] Richard E. Zippel. "Probabilistic Algorithms for Sparse Polynomials". PhD thesis. Massachusetts Institute of Technology, 1979. URL: https://search.worldcat.org/title/642421899.

A Proof-Coverage Map

This appendix records where each main-body theorem, proposition, and corollary is proved. The main body contains definitions, experiments, theorem statements, and concise proof pointers; all formal proofs appear after the bibliography.

Main-body statement	Label	Proof location
Small-level modular-conjugacy support	Theorem 1	Section C.3
Small-parameter distinguisher	Theorem 2	Section C.3
Extractable conjugacy leakage	Theorem 3	Section B.1
High-degree modular traces	Corollary 1	Section B.1
Bounded-correspondence barrier	Proposition 1	Section B.2
High-parameter classification	Proposition 2	Section B.2
Predicate-only lower bound	Proposition 3	Section D.1
Local-neighbor leakage	Proposition 4	Section D.1
Bounded-relation reduction	Proposition 5	Section D.1
Low-degree ambient sparse-hit bound	Theorem 4	Section D.2
Bounded-branch ambient samplers	Corollary 2	Section D.2
Affine-line nonconcentration	Theorem 5	Section E.4
Degree-one indexed samplers fail	Corollary 3	Section E.4

B Proofs for the Transparent-Trace Leakage Interface

This appendix proves the leakage statements used in Sections 6 and 7. The experiments and witness interfaces are defined in Section 4.

B.1 Frobenius composition and OneEnd leakage

Lemma 1 (Canonical represented Frobenius). *Let $E/\mathbb{F}_{p^2}$ be given by a validated Weierstrass model. The p-power Frobenius map*

$$F_E : E \to E^{(p)}, \qquad (x, y) \mapsto (x^p, y^p),$$

and its twist counterpart $F_{E^{(p)}} : E^{(p)} \to E$ have deterministic polynomial-size admissible representations. Their degree is p.

Proof. The field representation of $\mathbb{F}_{p^2}$ is part of the public parameters. Raising a coordinate to the p-th power is the nontrivial Galois automorphism of $\mathbb{F}_{p^2}/\mathbb{F}_p$, hence is computable by a fixed linear map over $\mathbb{F}_p$. Applying this map to the curve equation sends the coefficients of E to those of $E^{(p)}$, and sends points of E to points satisfying the equation of $E^{(p)}$. Since the absolute inseparable degree of the relative p-Frobenius on an elliptic curve is p, its degree is p. The representation consists of the coordinate-wise Frobenius rule together with the source and target models, and verification consists of checking the target equation and the public degree certificate.

Lemma 2 (Degree certifies non-scalarity). *Let E be an elliptic curve and let $\alpha \in \mathrm{End}(E)$ have degree pD, where p is prime and $p \nmid D$. Then $\alpha \notin \mathbb{Z}$.*

Proof. If $\alpha = [m]$ were scalar multiplication by an integer m, then $\deg(\alpha) = m^2$. The p-adic valuation of m^2 is even. The p-adic valuation of pD is exactly one because $p \nmid D$. Hence pD is not a square in $\mathbb{Z}$, so α cannot be scalar multiplication.

Proof (Proof of Theorem 3). Let $\eta(\lambda)$ be the non-negligible probability with which the public trace yields the three objects in the theorem statement. The OneEnd adversary reruns or instruments Gen on the public seed, applies the assumed extractor to obtain (E, ψ, C, u), and forms the represented composition

$$\alpha = F_{E^{(p)}} \circ u \circ \psi : E \to E.$$

Each factor has an admissible representation: ψ by assumption, u as a represented isomorphism of degree one, and $F_{E^{(p)}}$ by Lemma 1. The admissible-representation interface is closed under composition by Definitions 5 and 6, so the adversary obtains an admissibly represented endomorphism of E.

The degree is multiplicative under composition:

$$\deg(\alpha) = \deg(F_{E^{(p)}}) \cdot \deg(u) \cdot \deg(\psi) = p \cdot 1 \cdot D = pD.$$

Since $p \nmid D$, Lemma 2 implies that α is non-scalar. Thus the adversary wins $\mathsf{Exp}^{\mathsf{one}}$ whenever the extractor succeeds, namely with probability at least $\eta(\lambda)$, which is non-negligible. Hence Gen is not OneEnd-leakage secure.

Proof (Proof of Corollary 1). A successful trace for $\Phi_N(j, j^p) = 0$ that contains an admissibly represented cyclic degree-N isogeny from E_j to a curve isomorphic to $E_j^{(p)}$ satisfies the hypotheses of Theorem 3 with $D = N$. The condition $p \nmid N$ is the prime-to-p condition required in the theorem. Applying the theorem gives a polynomial-time OneEnd adversary with the same non-negligible success probability.

B.2 Bounded correspondences and route classification

Proof (Proof of Proposition 1). Let $\eta(\lambda)$ be the non-negligible probability that the successful trace contains a PPT-extractable elliptic correspondence trace. On input the adversarial view, run the extractor promised by Definition 9. If the extracted object is a represented non-scalar endomorphism of E, output it and win the OneEnd experiment. If it is a represented path from a known-ring public anchor to E, output the path and win the known-root path experiment. If it is a represented prime-to-p isogeny from E to a curve isomorphic to $E^{(p)}$, invoke the adversary from Theorem 3 and win OneEnd. If it is a represented elliptic correspondence convertible into one of these objects, run the converter and then use the corresponding branch. Therefore every extractable case yields OneEnd or path leakage with non-negligible probability. The only remaining case is that the trace is not extractable as an elliptic correspondence in the sense of the definition, in which case the sampler is outside the modeled barrier and must be treated as a residual mechanism.

Proof (Proof of Proposition 2). The proposition is a case distinction inside the extractability model. If the trace contains a prime-to-p Frobenius-conjugacy witness, Theorem 3 applies. If it contains a represented path from a known public anchor, the adversary wins $\mathsf{Exp}^{\mathsf{path}}$ by outputting that path. If it contains an orientation, endomorphism, or correspondence convertible into either of these objects, the converter reduces the case to the first or second alternative. If none of these extractors or converters exists, then by Definitions 7 and 9 the route is outside the modeled extractability interface. Such a route is exactly one of the residual root-only, endpoint-only, compact high-degree, or nonlocal aggregate mechanisms listed in the proposition.

C Proofs for Modular-Conjugacy Barriers

This appendix proves the statements of Section 5. We use standard Deuring and Eichler–Deuring optimal-embedding inputs in imported form. The role of this appendix is to make the reduction from modular-conjugacy witnesses to the stated support bound explicit. In particular, the non-normal branch uses the scalar-degeneracy argument $\beta = \psi^{(p)} \circ \psi$, whose degree is exactly N^2; no Frobenius factor p enters that branch.

C.1 Imported arithmetic inputs

Theorem 6 (Imported optimal-embedding bound). *Let $B_{p,\infty}$ be the quaternion algebra over $\mathbb{Q}$ ramified exactly at p and ∞. Let $\mathcal{R}$ range over maximal orders corresponding, through geometric Deuring correspondence, to supersingular elliptic curves over $\overline{\mathbb{F}}_p$. Let $\mathcal{O}$ be an imaginary quadratic order of discriminant $\Delta < 0$ with conductor prime-to-p except for the specified ramified local factor. The number of geometric supersingular j-classes admitting an optimal embedding of $\mathcal{O}$ into the corresponding maximal order is bounded by*

$$O\big(h(\mathcal{O}) \prod_{q \mid \mathrm{cond}(\mathcal{O})} (v_q(\mathrm{cond}(\mathcal{O})) + 1)\big),$$

up to an absolute automorphism factor. In the parameter regimes used below this is

$$O\big(|\Delta|^{1/2} \kappa(|\Delta|)\big)$$

for a factor $\kappa(x) = x^{o(1)}$.

The statement is the standard consequence of the Eichler–Deuring optimal-embedding formula: a weighted count of optimal embeddings equals a class-number term times local embedding factors; nonoptimal embeddings are counted by summing over overorders; the passage from weighted to unweighted geometric j-classes costs at most $|\mathrm{Aut}(E)| \le 24$. The local factor at a split or ramified finite prime contributes at most a divisor-type factor, absorbed in κ. The paper uses only this coarse upper bound, not exact constants.

Lemma 3 (Imported small-endomorphism count). *There is an absolute constant $c > 0$ such that, for every $M \ge 1$, the number of supersingular j-invariants admitting a non-scalar endomorphism of degree at most M is at most*

$$cM^{3/2}\kappa(M),$$

where $\kappa(M) = M^{o(1)}$ absorbs divisor-type losses. In the range $M < p$, this is the small-endomorphism rarity input; for larger M, the same upper bound follows after increasing c from the trivial bound $|\mathrm{SS}_p| = O(p)$ whenever $M^{3/2} \ge p$, and otherwise the imported small-endomorphism estimate applies.

Proof. A non-scalar endomorphism of degree at most M places the curve in a small-endomorphism locus. We use the standard small-endomorphism rarity estimate of Love–Boneh [LB20] in the range needed for small public levels. Equivalently, one may obtain the same shape by summing the Eichler–Deuring embedding bound over the imaginary quadratic orders generated by non-scalar endomorphisms of norm at most M, with trace and local-factor losses absorbed by $\kappa(M)$. If M is so large that the imported estimate is not the sharp input, the trivial bound $|\mathrm{SS}_p| = O(p)$ is weaker only when $M^{3/2} < p$, which is precisely the small-endomorphism range. Thus the displayed uniform bound is sufficient for the uses below.

C.2 Normal and non-normal witnesses

Let $\psi : E \to E^{(p)}$ be a cyclic separable degree-N isogeny with $p \nmid N$. Its Frobenius twist is an isogeny

$$\psi^{(p)} : E^{(p)} \to E^{(p^2)} = E,$$

and its dual is

$$\widehat{\psi} : E^{(p)} \to E.$$

We call the witness normal if $\psi^{(p)} = \varepsilon\widehat{\psi}$ for some $\varepsilon \in \{\pm 1\}$, up to the explicit automorphism ambiguity at $j = 0, 1728$. Otherwise it is non-normal.

Lemma 4 (Scalar-degeneracy lemma). *Let $\psi : E \to E^{(p)}$ be a cyclic separable degree-N isogeny with $p \nmid N$, and define*

$$\beta = \psi^{(p)} \circ \psi \in \mathrm{End}(E).$$

Then $\deg(\beta) = N^2$. If $\beta \in \mathbb{Z}$, then $\psi^{(p)} = \varepsilon\widehat{\psi}$ for some $\varepsilon \in \{\pm 1\}$. Consequently, if ψ is non-normal and $j(E) \notin \{0, 1728\}$, then β is a non-scalar endomorphism of degree exactly N^2.

Proof. Degree multiplicativity gives

$$\deg(\beta) = \deg(\psi^{(p)}) \deg(\psi) = N^2.$$

Assume that β is scalar, say $\beta = [m]$. Since $\deg([m]) = m^2$, we have $m^2 = N^2$, and hence $m = \pm N$. The dual isogeny satisfies

$$\widehat{\psi} \circ \psi = [N]_E.$$

If $m = N$, then

$$\psi^{(p)} \circ \psi = [N]_E = \widehat{\psi} \circ \psi.$$

Because ψ is a finite surjective morphism of curves, right cancellation holds in the category of morphisms of curves: two morphisms out of $E^{(p)}$ that agree after precomposition with the surjective map ψ are equal. Thus $\psi^{(p)} = \widehat{\psi}$. If $m = -N$, the same argument applied to

$$\psi^{(p)} \circ \psi = [-N]_E = (-\widehat{\psi}) \circ \psi$$

gives $\psi^{(p)} = -\widehat{\psi}$. Therefore scalarity of β forces normality. The contrapositive proves that a non-normal witness gives a non-scalar β of degree N^2 away from the extra automorphism cases.

Lemma 5 (Non-normal witnesses give degree-N^2 small endomorphisms). *Let $\mathcal{C}_N^{\mathrm{nonnormal}}(p)$ be the subset of $\mathcal{C}_N(p)$ admitting a non-normal cyclic degree-N conjugacy isogeny. Then*

$$|\mathcal{C}_N^{\mathrm{nonnormal}}(p)| \le O(N^3 \kappa(N)).$$

More precisely,

$$\mathcal{C}_N^{\mathrm{nonnormal}}(p) \subseteq \mathsf{Small}_{N^2}(p) \cup \{0, 1728\},$$

where $\mathsf{Small}_M(p)$ is the supersingular locus admitting a non-scalar endomorphism of degree at most M.

Proof. Let $j \in \mathcal{C}_N^{\mathrm{nonnormal}}(p)$. Then there is a supersingular curve $E/\mathbb{F}_{p^2}$ with $j(E) = j$ and a cyclic separable degree-N isogeny

$$\psi : E \to E^{(p)}$$

which is non-normal. If $j \notin \{0, 1728\}$, Lemma 4 shows that

$$\beta = \psi^{(p)} \circ \psi \in \mathrm{End}(E) \setminus \mathbb{Z}$$

has degree exactly N^2. Therefore $j \in \mathsf{Small}_{N^2}(p)$.

At $j = 0, 1728$, the automorphism group is larger than $\{\pm 1\}$. These give at most two j-invariants. If the only failure of the strict $\pm$-normality convention is $\psi^{(p)} = u\widehat{\psi}$ for a nontrivial automorphism u, then $u \circ [N]$ is already a non-scalar endomorphism of degree N^2. Hence these exceptional cases are either counted by $\mathsf{Small}_{N^2}(p)$ or contribute $O(1)$ points.

By Lemma 3, applied with $M = N^2$,

$$|\mathsf{Small}_{N^2}(p)| \le O((N^2)^{3/2}\kappa(N^2)) = O(N^3\kappa(N)).$$

The displayed inclusion gives the result.

Lemma 6 (Normal witnesses give trace-zero orientations). *Let $\psi : E \to E^{(p)}$ be a normal cyclic degree-N conjugacy isogeny. Then the associated Frobenius-composition endomorphism yields an element $\mu \in \mathrm{End}(E)$ satisfying*

$$\mu^2 = -pN$$

up to the sign convention determined by normality. Conversely, for fixed E, the number of distinguished elements μ over a fixed embedded quadratic order is at most two, namely $\pm\mu$, outside the already bounded automorphism ambiguity at $j = 0, 1728$.

Proof. Let $F_{E^{(p)}} : E^{(p)} \to E^{(p^2)} = E$ be the relative p-power Frobenius back to E, and put $\mu = F_{E^{(p)}} \circ \psi$. Taking Frobenius twists and using compatibility of duals with Frobenius twist gives

$$\widehat{\psi}^{(p)} = \widehat{\psi^{(p)}}.$$

Normality says $\psi^{(p)} = \varepsilon\widehat{\psi}$. Therefore the square of $F_{E^{(p)}} \circ \psi$, after identifying the double Frobenius twist with multiplication by $-p$ in the supersingular trace-zero convention, is $-pN$ up to the harmless sign convention. Thus $\mathbb{Z}[\mu]$ is an imaginary quadratic order in $\mathrm{End}(E)$. For a fixed embedded order, the elements satisfying the same trace and norm equation are exactly the two roots of the corresponding quadratic polynomial, hence at most $\pm\mu$. Extra automorphisms at $j = 0, 1728$ multiply counts only by an absolute factor already dominated by the unweighted automorphism bound.

C.3 Small-level support and distinguisher

Proof (Proof of Theorem 1). Fix N with $p \nmid N$. Let $\mathcal{C}_N^{\mathrm{norm}}(p)$ be the subset of $\mathcal{C}_N(p)$ admitting at least one normal cyclic degree-N conjugacy isogeny. Every element of $\mathcal{C}_N(p)$ either lies in $\mathcal{C}_N^{\mathrm{norm}}(p)$, or else all its cyclic degree-N conjugacy witnesses are non-normal; in the latter case Lemma 5 places it in $\mathsf{Small}_{N^2}(p) \cup \{0, 1728\}$. Hence

$$|\mathcal{C}_N(p)| \le |\mathcal{C}_N^{\mathrm{norm}}(p)| + O(N^3\kappa(N)).$$

For the normal part, Lemma 6 injects the data, up to an absolute automorphism and $\pm\mu$ multiplicity factor, into embeddings of quadratic orders generated by an element μ with $\mu^2 = -pN$. The relevant discriminants and conductors have size polynomial in N times the ramified factor at p. Applying Theorem 6, summing over possible overorders, and converting from weighted embeddings to unweighted geometric j-classes gives

$$|\mathcal{C}_N^{\mathrm{norm}}(p)| \le O(N^3\kappa(N) + N\kappa(N)p^{1/2}).$$

The first term accounts for conductor and local-factor summations independent of the ambient supersingular mass; the second is the ramified class-number scale. Combining the normal and non-normal bounds gives

$$|\mathcal{C}_N(p)| \le O(N^3\kappa(N) + N\kappa(N)p^{1/2}),$$

which is the displayed bound after naming the absolute constants.

For $B = \mathrm{poly}(\lambda)$, a union bound gives

$$|\mathcal{C}_{\le B}(p)| \le \sum_{2 \le N \le B} \left(c_1 N^3\kappa(N) + c_2 N\kappa(N)p^{1/2}\right) \le p^{1/2}\,\mathrm{poly}(\lambda)$$

after increasing the polynomial factor. Since $|\mathrm{SS}_p| = \Theta(p)$, the ratio is at most $p^{-1/2}\,\mathrm{poly}(\lambda) = \mathrm{negl}(\lambda)$ for $p \asymp 2^\lambda$.

Proof (Proof of Theorem 2). The distinguisher $\mathcal{D}_B$ validates supersingularity and then tests all equations $\Phi_N(j, j^p) = 0$ for $2 \le N \le B$. For $B = \mathrm{poly}(\lambda)$, classical modular-polynomial evaluation at these levels is polynomial time in λ and B. If j is sampled from Gen, the support condition implies $j \in \mathcal{C}_{\le B}(p)$ except with negligible probability, so $\mathcal{D}_B$ outputs the generator side with probability $1 - \mathrm{negl}(\lambda)$. If $j \leftarrow U(\mathrm{SS}_p)$, Theorem 1 implies

$$\Pr[j \in \mathcal{C}_{\le B}(p)] \le \mathrm{negl}(\lambda).$$

Thus the statistical gap between the two cases is $1 - \mathrm{negl}(\lambda)$, giving distinguishing advantage $1/2 - \mathrm{negl}(\lambda)$ in the experiment normalized around success probability $1/2$.

D Proofs for Direct Sampler Barriers

D.1 Predicate, local-neighbor, and bounded-relation proofs

Proof (Proof of Proposition 3). Fix the randomness of the algorithm. It becomes a deterministic adaptive algorithm. Condition on the all-negative path, i.e., the event that every previous membership query returns zero. Along this path the algorithm's sequence of at most q queried points is a fixed set $Q \subseteq U$ with $|Q| \le q$, independent of the hidden subset once the all-negative answers are fixed. If the algorithm hits S within q queries, then $S \cap Q \neq \emptyset$. For a uniformly hidden subset of size $s = \Theta(p)$ in a universe of size $|U| = p^2$,

$$\Pr[S \cap Q \neq \emptyset] \le \frac{|Q|s}{|U|} \le O\left(\frac{q}{p}\right).$$

Thus any deterministic adaptive algorithm making $q = o(p)$ queries succeeds with probability $o(1)$. Averaging over the algorithm's random coins gives the same bound for randomized algorithms. Taking $U = \mathbb{F}_{p^2}$ and $S = \mathrm{SS}_p$, where $|\mathrm{SS}_p| = \Theta(p)$, gives the supersingularity-search statement.

Proof (Proof of Proposition 4). The adversary receives the public seed and the reconstructible trace. By assumption, the trace contains the intermediate curves, parent pointers, edge kernels, or equivalent certificates for each small-degree step from the known-ring anchor E_0 to the output curve. The adversary reads these represented edges from the trace and outputs them in order. Verification of the known-root path experiment checks exactly that the sequence composes to a path from E_0 to a curve with invariant j. Hence the adversary wins whenever the generator succeeds with a reconstructible local-neighbor trace.

Proof (Proof of Proposition 5). Run the promised polynomial-time converter on the bounded relation R extracted from the trace. The output is an extractable elliptic correspondence trace in the sense of Definition 9. Applying Proposition 1 gives OneEnd leakage, known-root path leakage, or the residual non-extractable alternative. This is exactly the conclusion of the proposition.

D.2 Low-degree ambient sparse-hit bound

Lemma 7 (Rational root-hit bound). *Let $S \subseteq \mathbb{F}_q$ have size s. Let*

$$J = N/D : \mathbb{F}_q^m \dashrightarrow \mathbb{F}_q$$

be a rational map of formal degree at most Δ, and suppose J is not constant with value in S. Then

$$\Pr_{X \leftarrow \mathbb{F}_q^m}[D(X) \neq 0 \ \wedge \ J(X) \in S] \le \frac{s\Delta}{q}.$$

Proof. Define

$$P_S(X) = \prod_{\alpha \in S} (N(X) - \alpha D(X)).$$

If $D(X) \neq 0$ and $J(X) = N(X)/D(X) \in S$, then one factor vanishes, hence $P_S(X) = 0$. The degree of P_S is at most $s\Delta$. The polynomial is not identically zero: if it were, since $\mathbb{F}_q[X_1, \ldots, X_m]$ is an integral domain, some factor $N - \alpha D$ would be the zero polynomial, making J the constant $\alpha \in S$, contrary to the hypothesis. Schwartz–Zippel therefore gives

$$\Pr[P_S(X) = 0] \le \frac{\deg P_S}{q} \le \frac{s\Delta}{q}.$$

The desired event is contained in $\{P_S = 0\}$, proving the bound.

Proof (Proof of Theorem 4). Apply Lemma 7 with $q = p^2$, $S = \mathrm{SS}_p$, and $|\mathrm{SS}_p| = \Theta(p)$. The event $J(X) \in \mathrm{SS}_p$ is contained in the union of the denominator-exception event $D(X) = 0$ and the event bounded by the lemma. Hence

$$\Pr[J(X) \in \mathrm{SS}_p] \le \Pr[D(X) = 0] + \frac{|\mathrm{SS}_p|\Delta}{p^2} \le \Pr[D(X) = 0] + O\left(\frac{\Delta}{p}\right).$$

If the denominator exception probability is negligible and the left-hand side is bounded below by a positive constant, then necessarily $\Delta = \Omega(p)$.

Proof (Proof of Corollary 2). For every nonconstant branch, Theorem 4 bounds its supersingular-hit probability by $O(\Delta/p)$, plus its denominator-exception probability. The branch condition can only shrink the event, and a union bound over B branches gives total nonconstant-branch success probability $O(B\Delta/p)$ plus the total denominator-exception probability. If $B, \Delta = \mathrm{poly}(\log p)$, this is negligible. Constant branches with supersingular values can output at most $B = \mathrm{poly}(\log p)$ distinct invariants, which is a negligible-support distribution inside $|\mathrm{SS}_p| = \Theta(p)$ and is therefore excluded except as the small fixed-support alternative in the statement.

E Affine-Line Monodromy and Chebotarev Proof

This appendix proves Theorem 5. The proof imports standard modular-cover monodromy and finite-field Chebotarev theorems, and proves the affine-line descent, no-accidental-diagonal, and coset-density steps needed for arbitrary affine $\mathbb{F}_p$-lines in $\mathbb{F}_{p^2}$.

E.1 Imported modular-cover and Chebotarev inputs

Theorem 7 (Imported projective level-cover monodromy). *For every prime $\ell \geq 5$, the projective full-level-ℓ modular cover over the coarse punctured j-line*

$$U_j = \mathbb{P}^1_j \setminus \{0, 1728, \infty\}$$

is finite etale, has geometric projective monodromy G_ℓ satisfying

$$\mathrm{PSL}_2(\mathbb{F}_\ell) \subseteq G_\ell \subseteq \mathrm{PGL}_2(\mathbb{F}_\ell),$$

and has degree, compactified genus, and conductor complexity bounded by $O(\ell^A)$ for an absolute constant A. The compactified cover has typed branch over $0, 1728, \infty$ with inertia of orders $3, 2, \ell$, respectively.

This is the standard projective form of full-level modular-curve monodromy. Projectivization is essential: over the coarse j-line, quadratic twists differ by scalar action on $E[\ell]$, and scalars vanish in $\mathrm{PGL}_2(\mathbb{F}_\ell)$.

Theorem 8 (Imported finite-field Chebotarev for curve covers). *Let $\pi : W \to U$ be a finite etale cover of a geometrically irreducible affine curve over $\mathbb{F}_p$, with arithmetic monodromy group Γ, geometric monodromy group Γ^0, and cover complexity H. Let $C \subseteq \gamma\Gamma^0$ be a conjugacy-stable subset of the Frobenius coset relevant to $U(\mathbb{F}_p)$. Then*

$$\#\{t \in U(\mathbb{F}_p) : \mathrm{Frob}_t \in C\} = \frac{|C|}{|\gamma\Gamma^0|}\,|U(\mathbb{F}_p)| + O(Hp^{1/2}).$$

The same bound holds componentwise for the Galois closure when the cover is not geometrically connected, with constants depending polynomially on the cover complexity.

E.2 Typed branch loci and no accidental diagonal quotients

For an affine line $L = a + b\mathbb{F}_p$, set

$$j_1(t) = a + bt, \qquad j_2(t) = j_1(t)^p = a^p + b^p t.$$

The pullbacks of the compactified projective level cover along j_1 and j_2 have typed branch loci on the t-line:

$$B_1 = \left\{-\frac{a}{b}, \frac{1728 - a}{b}, \infty\right\}, \qquad B_2 = \left\{-\frac{a^p}{b^p}, \frac{1728 - a^p}{b^p}, \infty\right\},$$

where the three points have types $3, 2, \ell$, respectively.

Lemma 8 (Typed branch equality). *For $p > 3$, the typed branch loci of j_1 and j_2 coincide over the same t-line if and only if $a, b \in \mathbb{F}_p$.*

Proof. If $a, b \in \mathbb{F}_p$, then $j_1 = j_2$, so the typed branch loci coincide. Conversely, assume typed branch equality. Since the inertia types $3, 2, \ell$ are pairwise distinct for $\ell \geq 5$, equality preserves each typed point. Hence

$$-\frac{a}{b} = -\frac{a^p}{b^p} \qquad \text{and} \qquad \frac{1728 - a}{b} = \frac{1728 - a^p}{b^p}.$$

Put $\lambda = b^p/b$. The first equality gives $a^p = \lambda a$. The second gives

$$1728 - a^p = \lambda(1728 - a).$$

Substituting $a^p = \lambda a$ yields

$$1728 - \lambda a = 1728\lambda - \lambda a.$$

Since $p > 3$, the element 1728 is nonzero in $\mathbb{F}_p$, so $\lambda = 1$. Therefore $b^p = b$, hence $b \in \mathbb{F}_p$, and then $a^p = a$, hence $a \in \mathbb{F}_p$.

Lemma 9 (No accidental nonabelian diagonal quotient). *Assume $\ell \geq 5$ and not both a, b lie in $\mathbb{F}_p$. Let X_1 and X_2 be the two geometric pullbacks of the projective level-ℓ cover along j_1 and j_2. Then the geometric monodromy of the compositum has no common nonabelian simple projective quotient between the two factors. Equivalently, after passing to the simple derived projective quotient $\mathrm{PSL}_2(\mathbb{F}_\ell)$, the monodromy is product-like. Only bounded abelian correlations, such as the $\mathrm{PGL}_2 / \mathrm{PSL}_2$ quotient, may remain.*

Proof. Suppose a common nonabelian simple quotient exists. By Goursat's lemma, after quotienting the two factor monodromy groups by normal subgroups, one obtains an isomorphism between simple quotients of the two pulled-back covers. Since $\mathrm{PSL}_2(\mathbb{F}_\ell)$ is simple for $\ell \geq 5$, the relevant nonabelian quotient is a projective simple quotient of this type, up to automorphisms.

Consider a typed branch point P of the first pullback that is not a branch point of the second pullback with the same type. In the first simple quotient, the inertia image at P is nontrivial of order 2, 3, or ℓ, depending on the type. In the second quotient the cover is unramified at P, so the corresponding inertia image is trivial. The diagonal isomorphism would identify these two images, forcing the nontrivial inertia image to be trivial in the common simple quotient. The normal closure of any nontrivial element in a nonabelian simple quotient is the whole quotient, so this would kill the quotient, contradiction.

Therefore the typed branch loci must coincide. By Lemma 8, this forces $a, b \in \mathbb{F}_p$, contradicting the assumption. Hence no accidental common nonabelian simple quotient exists. Bounded abelian quotients are not excluded and are harmless for density estimates because their size is bounded independently of ℓ after passing to PSL_2.

E.3 Frobenius-stable descent and coset density

Lemma 10 (Frobenius-stable product descent). *For arbitrary $a, b \in \mathbb{F}_{p^2}$, $b \neq 0$, the product of the two pullbacks along j_1 and j_2 descends to a finite etale cover of*

$$U_L = \mathbb{A}^1_{\mathbb{F}_p} \setminus \{\text{the closed branch points above } 0, 1728, \infty\}.$$

Its arithmetic Frobenius coset acts by a factor swap, possibly composed with projective automorphisms of the level group.

Proof. The map j_2 is the coefficient-wise Frobenius conjugate of j_1. Therefore the p-power Frobenius sends the pullback cover along j_1 to the pullback cover along j_2, and sends the second back to the first. The fiber product of the two covers is stable under this Frobenius action. Galois descent gives a cover over the $\mathbb{F}_p$-parameter line after removing the full closed branch divisor, not merely its $\mathbb{F}_p$-rational points. On geometric fibers, arithmetic Frobenius exchanges the two level factors and may apply the natural projective automorphism induced by Frobenius on coefficients. This is the asserted swap action.

Lemma 11 (Supersingular projective square constraint). *Let $t \in U_L(\mathbb{F}_p)$, and suppose $j = a + bt \in \mathrm{SS}_p$. In the descended projective level cover, the arithmetic Frobenius class x_t satisfies that x_t^2 is projectively trivial on the relevant level factor.*

Proof. The first application of arithmetic Frobenius sends the fiber with invariant $j = a + bt$ to the conjugate fiber with invariant $j^p = a^p + b^p t$. The second application returns to the original j-fiber and realizes the p^2-power Frobenius action on the ℓ-torsion of an elliptic curve with invariant j. For $j \neq 0, 1728$, the projective level structure is independent of quadratic twist. If $E/\mathbb{F}_{p^2}$ is supersingular, then the p^2-Frobenius acts as a scalar on the two-dimensional ℓ-torsion representation; projectively, this scalar is the identity. Thus the square of the arithmetic Frobenius class is projectively trivial on the relevant factor. The removed points $0, 1728$ contribute only the excluded branch points.

Lemma 12 (Coset-level bad-density bound). *Let $\gamma\Gamma^0$ be a Frobenius coset arising from the descended cover in Lemma 10. Let $B_{L,\ell} \subseteq \gamma\Gamma^0$ be the subset whose square is projectively trivial on the relevant level factor. Then*

$$\frac{|B_{L,\ell}|}{|\gamma\Gamma^0|} \leq \frac{C}{\ell}$$

for an absolute constant C, uniformly in L, p, ℓ.

Proof. If $a, b \in \mathbb{F}_p$, the cover has one projective monodromy factor G with $\mathrm{PSL}_2(\mathbb{F}_\ell) \subseteq G \subseteq \mathrm{PGL}_2(\mathbb{F}_\ell)$. The bad set is contained in the set of projective involutions $\{g \in G : g^2 = 1\}$. In $\mathrm{PGL}_2(\mathbb{F}_\ell)$, nontrivial involutions have centralizer containing a split or nonsplit torus of size $\ell - 1$ or $\ell + 1$. Thus each involution class has size $O(\ell^2)$, and there are only boundedly many projective involution types, so the total number of involutions is $O(\ell^2)$. Since $|G| = \Theta(\ell^3)$, the bad density is $O(1/\ell)$.

If not both a, b lie in $\mathbb{F}_p$, Lemma 9 shows that the nonabelian projective part of the two-factor monodromy is product-like. After quotienting bounded abelian correlations, a Frobenius-coset element has the form $(g, h)\sigma$, where σ swaps the two factors and may apply a projective automorphism. Its square has first factor $g\sigma(h)$. For fixed g, the condition that $g\sigma(h)$ be projectively trivial determines h up to a bounded abelian ambiguity. If one uses the weaker condition that $g\sigma(h)$ lie in the projective involution set, then for fixed g there are at most $O(|G|/\ell)$ choices of h. Hence the bad set has size at most $O(|G|^2/\ell)$ inside a coset of size $\Theta(|G|^2)$, giving density $O(1/\ell)$. This bound is uniform over all connected components because the argument is applied to each component's Frobenius coset after Galois closure.

E.4 Chebotarev optimization and indexed-sampler corollary

Lemma 13 (Affine-line Chebotarev count). *There are absolute constants A, C_1, C_2 such that for every affine line $L = a + b\mathbb{F}_p$ and every auxiliary prime $\ell \geq 5$, $\ell \neq p$,*

$$|\mathrm{SS}_p \cap L| \leq C_1 \frac{p}{\ell} + C_2 \ell^A p^{1/2} + O(1).$$

Proof. Construct the descended cover of Lemma 10 and pass to its Galois closure componentwise. By Lemma 11, every $t \in U_L(\mathbb{F}_p)$ with $a + bt \in \mathrm{SS}_p$ has Frobenius class in the bad subset $B_{L,\ell}$. By Lemma 12, this subset has density at most C/ℓ in the relevant Frobenius coset. Applying Theorem 8 gives

$$\#\{t \in U_L(\mathbb{F}_p) : a + bt \in \mathrm{SS}_p\} \leq C_1 \frac{|U_L(\mathbb{F}_p)|}{\ell} + C_2 H(\ell) p^{1/2}.$$

The cover complexity $H(\ell)$ is polynomial in ℓ by Theorem 7; write $H(\ell) = O(\ell^A)$. The removed branch points contribute only $O(1)$ seeds, giving the stated bound.

Proof (Proof of Theorem 5). Let A be the exponent in Lemma 13. Choose a prime $\ell \asymp p^{1/(2A+2)}$, which exists for sufficiently large p by Bertrand's postulate and is different from p. Substituting into Lemma 13 gives

$$|\mathrm{SS}_p \cap L| \leq C_1 p \ell^{-1} + C_2 \ell^A p^{1/2} + O(1) \leq C p^{1-1/(2A+2)}.$$

Setting $\delta = 1/(2A + 2)$ proves the theorem.

Proof (Proof of Corollary 3). The map $t \mapsto a + bt$ is injective from $\mathbb{F}_p$ onto the affine line $L = a + b\mathbb{F}_p$ because $b \neq 0$. Therefore

$$\Pr_{t \leftarrow \mathbb{F}_p}[a + bt \in \mathrm{SS}_p] = \frac{|\mathrm{SS}_p \cap L|}{p} \leq C p^{-\delta}$$

by Theorem 5. This probability is subconstant, so no such sampler has constant success probability.