Spectral Theory of Isogeny Graphs and Quantum Sampling of Secure Supersingular Elliptic Curves

Maher Mamah[1], Jake Doliskani[2] and David Jao[1]

[1] University of Waterloo, Department of Combinatorics and Optimization, Waterloo, Canada
[2] McMaster University, Department of Computing and Software, Hamilton, Canada

Abstract. In this paper, we study the problem of sampling random supersingular elliptic curves with unknown endomorphism rings. This problem has recently gained considerable attention as many isogeny-based cryptographic protocols require such "secure" curves for instantiation, while existing methods achieve this only in a trusted-setup setting. We present the first *provable* quantum polynomial-time algorithms for sampling such curves with high probability, one of which builds on an algorithm of Booher *et al* (2024). One variant runs heuristically in $O(\log^{6.5} p)$ quantum gate complexity, and in $O(\log^{20} p)$ under the Generalized Riemann Hypothesis, and outputs a curve that is *provably* secure assuming quantum average-case hardness of the endomorphism ring problem and a plausible number theoretic conjecture. Another unconditional variant samples uniform $\mathcal{O}$-oriented curves with unknown endomorphism rings, for any imaginary quadratic order $\mathcal{O}$, with security based on the average quantum-hardness of VECTORIZATION problem. Our algorithms provide a secure instantiation of the CGL hash function and related primitives and show a quantum advantage over classical algorithms.

Our analysis relies on a spectral delocalization result for supersingular ℓ-isogeny graphs: we prove a strong version of the sup-norm conjecture of eigenvectors and provide numerical evidence for complete eigenvector delocalization. We also prove a stronger ε-separation property for eigenvalues of isogeny graphs than that predicted in the quantum money protocol of Kane, Sharif, and Silverberg, thereby removing a key heuristic assumption in their construction.

Keywords: Isogeny-based cryptography · Secure curves · Quantum algorithms · Endomorphism rings

1 Introduction

Isogeny-based cryptography is a promising approach to post-quantum cryptography, aiming to construct cryptographic schemes that remain secure against quantum adversaries while relying only on classical algorithms. Its central hardness assumption is the *isogeny path problem*: given two supersingular elliptic curves E and E', find an isogeny $\varphi : E \to E'$ connecting them. Despite its recent development, the area has attracted significant cryptanalytic attention.

One of the earliest primitives in this area is the CGL hash function of Charles, Goren, and Lauter [CGL09], defined via random walks on the supersingular ℓ-isogeny graph. Collision resistance reduces to the hardness of the *endomorphism ring problem* ENDRING: given a supersingular elliptic curve E, compute its ring of endomorphisms $\mathrm{End}(E)$. The reduction proceeds via the intermediate problem ONEEND of computing a nontrivial endomorphism [PW24]. Since the supersingular isogeny path problem is computationally

E-mail: mmamah@uwaterloo.ca (Maher Mamah), jake.doliskani@mcmaster.ca (Jake Doliskani), djao@uwaterloo.ca (David Jao)

equivalent to EndRing [EHL+18, Wes22, Mam24], initializing the hash function from a curve whose endomorphism ring is known renders the scheme insecure: the random walk reveals an explicit path from E_0 to E, which allows one to recover $\mathrm{End}(E)$ and hence find collisions.

A key requirement is therefore to sample the starting curve without revealing information about its endomorphism ring. This leads to a fundamental problem: how to generate a supersingular elliptic curve whose endomorphism ring is unknown, often called a *secure curve*. This assumption underlies many constructions, including the CGL hash function [CGL09], commitment schemes [Ste22], isogeny-based VDFs [DFMPS19], oblivious pseudorandom functions [Bas24], verifiable random functions [LP25], identity-based encryption [GS25], and related primitives.

Despite its importance, no efficient classical or quantum algorithm is known to sample guaranteed secure curves. Approaches based on generating supersingular j-invariants require exponential resources [MMP24]. Other classical and quantum proposals [BBD+24, Dol23], building on ideas of [KSS22], remain heuristic and provide no provable security.

Currently, the only method guaranteeing secure curves relies on a trusted setup. Basso *et al.* [BCC+23] proposed a distributed protocol in which multiple parties jointly generate a curve by sequentially applying isogenies and proving correctness in zero knowledge; the output curve is secure provided at least one participant behaves honestly. However, this approach only produces secure curves over $\mathbb{F}_{p^2}$. Many protocols such as CSIDH-based VDFs [DFMPS19], oblivious transfer [LGDdSG21], and delay encryption [BDF21]—require secure curves over $\mathbb{F}_p$ equipped with an orientation by a quadratic order ‰, typically induced by Frobenius. Mokrani and Jao [MJ23] proposed a related protocol for generating such curves over $\mathbb{F}_p$, but it likewise relies on a trusted setup.

These limitations motivate the central problem addressed in this work: efficiently sampling supersingular elliptic curves whose endomorphism rings are unknown, without relying on a trusted setup. We develop quantum algorithms for this task by exploiting spectral properties of supersingular isogeny graphs and group actions arising from oriented isogenies.

Contributions. We present two quantum polynomial-time algorithms for sampling *provably secure* supersingular elliptic curves, together with an explicit gate-complexity analysis. The first is based on continuous-time quantum walks, and admits both a heuristic variant optimized for efficiency and a rigorous conditional variant with provable guarantees. A heuristic sketch appeared in [BBD+24], but without any running-time or security analysis; here we provide provable guarantees for both. We give the first formal security proof for the sampled curves, relating their security to the quantum average-case hardness of EndRing. Under a plausible number-theoretic conjecture, any quantum polynomial-time adversary has only negligible probability of computing the endomorphism ring of a curve sampled by our algorithm.

The second contribution is an unconditional quantum algorithm for sampling secure $\mathfrak{O}$-oriented supersingular elliptic curves uniformly for any imaginary quadratic order $\mathfrak{O}$. To the best of our knowledge, this is the first instance of such a sampling algorithm for oriented curves. Its security relies on the group action inverse problem, commonly known as Vectorization.

A central technical contribution is to replace prior heuristic assumptions with rigorous conditional statements. Under the generalized Riemann hypothesis (GRH), we prove the ε-separation of eigenvalues required by Algorithm 1, obtaining a stronger bound than earlier heuristic results [Kan18, KSS22, BBD+24, Dol23].
Our main ingredient for proving security is a new eigenvector delocalization theorem for the ℓ-isogeny graphs (Theorem 3), showing that the relevant eigenvectors are sufficiently flattened; the proof crucially uses the theory of automorphic forms. Our techniques can

also be adapted to curves with N-level structure, but we leave this extension for future work.

Threat model. We note that the threat model is important in our setting. In the cryptographic literature it is common to distinguish between an *honest-but-curious* adversary, who follows the prescribed protocol but attempts to learn additional information from the execution, and an *active* adversary, who may deviate arbitrarily from the protocol.

One might ask why Alice, acting honestly, cannot simply perform a classical random walk in an isogeny graph, output the final curve, and discard the walk. The difficulty is that the walk itself may be observable: an adversary with access to Alice's internal computation, for instance through the cloud, could record the intermediate curves and reconstruct the entire path. Moreover, if Alice is honest-but-curious, she may attempt to extract additional information from the execution. By contrast, Algorithm 1 produces the output curve through Hamiltonian simulation, so that no explicit isogeny walk is ever generated. This makes the quantum approach more robust against honest-but-curious leakage.

The remaining concern is the active setting, where the generator may output a malicious trapdoor curve. One possible mitigation is to combine our algorithms with interactive quantum computation verification protocols [FK17, ABOEM17], in order to verify that the curve was sampled by an honest execution. We note, however, that there is a distinction between a trusted-setup setting and an authority-free setting. For instance, a standardization authority such as NIST could publicly generate a supersingular curve $E_\star$ using Algorithms 1 or 3, depending on application, together with such a quantum computation verification protocol. From such a publicly trusted secure base curve, one could bootstrap further candidate hard curves by taking random walks, thereby removing much of the practical parameter-generation bottleneck.

However, we note a subtlety in the problem of generating secure curves in previous proposals [MMP24, BBD+24]. As defined in [MMP24, Section 3.3], a curve E is secure if EndRing(E) is computationally hard. This notion is ambiguous, however: hard for whom? In particular, the definition should account for the generator itself, who must also be unable to compute the endomorphism ring of the output curve.

Ideally, one would like an algorithm that, given a random seed, outputs a secure curve, or j-invariant, without revealing any path information. Even if such an algorithm exists, one must still convince other parties that it was executed honestly, or indeed that it was executed at all. Algorithm 1 satisfies the former requirements in the quantum setting: it is inherently random, does not generate an explicit walk in the graph, and outputs only auxiliary spectral data. Our main contribution is showing that this auxiliary data is computationally useless: any curve honestly generated by Algorithm 1 remains secure under the quantum average-case hardness of EndRing, whose best known attack has complexity $O(p^{1/4})$ [BJS14].

Problem 1 (EndRing)**.** *Given a uniformly random supersingular elliptic curve $E/\mathbb{F}_{p^2}$, compute four endomorphisms that generate* $\mathrm{End}(E)$*.*

We note that the worst-case version of EndRing reduces to the average-case version by taking a sufficiently long random isogeny walk from the challenge curve. Next, we show that any oriented curve generated using algorithm 3 is secure assuming the average case hardness of Vectorization problem.

Problem 2 (Vectorization)**.** *Given an $\mathfrak{O}$-oriented supersingular elliptic curve (E_0, ι_0) and a uniformly random $\mathfrak{O}$-oriented curve (E_1, ι_1) defined over $\mathbb{F}_{p^2}$, compute an $\mathfrak{O}$-ideal $\mathfrak{a}$ such that $(E_1, \iota_1) = \mathfrak{a} * (E_0, \iota_0)$.*

Problem 2 is self-reducible, thus enjoying a worst-case to average-case reduction. Its hardness assumption underlies CSIDH [CLM+18], CSI-FiSh [BKV19], and related

primitives; it is believed classically hard, but admits a subexponential quantum attack via Kuperberg's algorithm [Kup05].

Outline. Section 2 reviews the material needed on supersingular elliptic curves, isogeny graphs, and the Deuring correspondence. Section 3 gives the analytic ingredients for our applications, including eigenvector delocalization and eigenvalue separation under GRH. Section 4 presents and analyzes the quantum sampling algorithms for the full isogeny graph and for oriented curves.

1.1 Technical overview

Our first algorithm exploits the spectral structure of the supersingular ℓ-isogeny graph $\mathcal{G}(p,\ell)$, whose vertices are supersingular elliptic curves over $\mathbb{F}_{p^2}$ and whose edges are ℓ-isogenies. We denote the set of such curves by $\mathcal{S}_p$. These graphs are $(\ell+1)$-regular Ramanujan graphs, and their adjacency matrices A_ℓ are Brandt matrices arising from Hecke operators. Since the operators A_ℓ commute for distinct primes ℓ, they admit a simultaneous orthonormal eigenbasis

$$\{\phi_i\}_i \subset \mathbb{C}^{\mathcal{S}_p}.$$

From the automorphic point of view, the ϕ_i correspond to Hecke eigenforms. A central analytic question is the extent to which these eigenvectors are delocalized, as predicted by the Quantum Unique Ergodicity conjecture. This is usually measured through the sup-norm

$$\|\phi_i\|_\infty = \max_E |\phi_i(E)|.$$

We prove the bound

$$\|\phi_i\|_\infty \ll p^{-1/4+\varepsilon}$$

in Theorem 3, together with an average form of complete delocalization in Theorem 4. These estimates, obtained from recent sup-norm bounds for quaternionic automorphic forms [KNS24], ensure that no eigenvector places too much mass on a small set of curves and form the main analytic ingredient in our security proof.

Sampling from isogeny graphs.

The sampling algorithm uses the simultaneous spectral decomposition of several commuting adjacency operators. Starting from a quantum state on the vertex space, we implement Hamiltonian simulation for $A_{\ell_1},\dots,A_{\ell_k}$ and apply quantum phase estimation. This produces a state of the form

$$\sum_i \alpha_i |\boldsymbol{\lambda}_i\rangle |\phi_i\rangle,$$

where $\boldsymbol{\lambda}_i$ is the vector of eigenvalues, or spectral tag, associated with ϕ_i. For the procedure to identify eigenvectors reliably, distinct eigenvectors must have distinguishable spectral tags. Previous works treated this separation heuristically; here we prove, under GRH, a stronger ε-separation statement than the heuristic bounds predicted in [Kan18, KSS22, Dol23].

After phase estimation, the eigenvector register is measured in the curve basis. The output probability of a curve E is governed by the values $|\phi_i(E)|^2$. The delocalization bounds above imply that these probabilities cannot be concentrated on a small subset of vertices, and hence the resulting distribution is close to uniform over the supersingular graph. We then show that curves sampled by this procedure are *quantumly secure* assuming the average-case hardness of EndRing and conjecture 1.

Sampling oriented supersingular curves.

Our second algorithm treats $\mathfrak{O}$-oriented supersingular elliptic curves. Here the relevant object is the regular class group action

$$(\mathrm{Cl}(\mathfrak{O}), \mathcal{X}, *),$$

where $\mathcal{X}$ is the set of $\mathfrak{O}$-oriented supersingular curves and $\mathrm{Cl}(\mathfrak{O})$ acts by oriented isogenies. For such actions, $\mathbb{C}^{\mathcal{X}}$ has a natural Fourier basis

$$|G^{(h)} * x\rangle = \frac{1}{\sqrt{|G|}} \sum_{g \in G} \chi(g,h)|g * x\rangle,$$

whose elements are simultaneous eigenvectors of the action operators $U_k : |y\rangle \mapsto |k * y\rangle$. The algorithm prepares a superposition of these Fourier states and recovers the index h using phase kickback together with the quantum Fourier transform over the class group. Measuring then yields a uniformly random $\mathfrak{O}$-oriented supersingular curve. Its security follows from the hardness of the VECTORIZATION problem for the class group action.

Notation. We write $f = O(g)$ and $f \ll g$ to mean that there exist constants $C > 0$ such that $|f(n)| \leq C\,|g(n)|$ for all sufficiently large n, in particular $f \ll 1 \iff f \leq c$. We also use the analytic number theory notation $f \asymp g$ to mean that there exist constants $c, c' > 0$ such that $c\,|g(n)| \leq |f(n)| \leq c'\,|g(n)|$ for all sufficiently large n (equivalently, $f \ll g$ and $g \ll f$). Unless stated otherwise, every theorem and proposition whose statement involves GRH is conditional on the Generalized Riemann Hypothesis.

2 Preliminaries

2.1 Background on Elliptic curves.

We briefly recall the elliptic-curve notions used below; see [Sil09] for background. Let K be a field of characteristic $p > 3$. An elliptic curve over K is given by $E : y^2 = x^3 + Ax + B$, for $A, B \in K$, $4A^3 + 27B^2 \neq 0$, with K-rational points together with the point at infinity ∞. An isogeny $\varphi : E_1 \to E_2$ defined over K is a nonconstant morphism that is also a group homomorphism and sends ∞_1 to ∞_2. Its degree is its degree as a rational map. Every isogeny of composite degree factors into prime-degree isogenies, and if $\deg(\varphi) = d$, then there is a unique dual isogeny $\hat{\varphi} : E_2 \to E_1$ such that $\hat{\varphi} \circ \varphi = [d]$. Isogenies of degree 1 are isomorphisms. The j-invariant classifies elliptic curves up to $\overline{\mathbb{F}}_p$-isomorphism, though over smaller fields several K-isomorphism classes may share the same j-invariant.

Since we require explicit representatives over $\mathbb{F}_{p^2}$, we use the following standard parametrization.

Proposition 1. *Let $p \neq 2, 3$, set $K = \mathbb{F}_{p^2}$, and fix a generator $\alpha \in K^\times$. Represent K-isomorphism classes of elliptic curves over K by pairs $(j, b) \in K \times \mathbb{Z}$, where*

$$b \in \{0,1\} \text{ if } j \neq 0, 1728, \quad b \in \{0,1,2,3\} \text{ if } j = 1728, \quad b \in \{0,1,2,3,4,5\} \text{ if } j = 0.$$

Define

$$\begin{aligned} \text{if } j \neq 0, 1728: &\quad E_{j,b} : y^2 = x^3 + \alpha^{2b}\frac{3j}{1728 - j}\,x + \alpha^{3b}\frac{2j}{1728 - j}, \\ \text{if } j = 1728: &\quad E_{1728,b} : y^2 = x^3 + \alpha^b x, \\ \text{if } j = 0: &\quad E_{0,b} : y^2 = x^3 + \alpha^b. \end{aligned}$$

Then $E_{j,b} \cong_K E_{j,b'}$ if and only if $b \equiv b' \pmod 2$ for $j \neq 0, 1728$, $b \equiv b' \pmod 4$ for $j = 1728$, and $b \equiv b' \pmod 6$ for $j = 0$.

Proof. See [Sil09, Corollary X.5.4.1]. □

If $\deg(\varphi)$ is coprime to p, then φ is uniquely determined by its kernel. Given the kernel, Vélu's formulas compute φ in polynomial time in $\deg(\varphi)$ and $\log p$ [Vél71]; asymptotically faster methods, with square-root complexity, are also known [BFLS20].

An endomorphism of E is an isogeny from E to itself, and these form the ring $\mathrm{End}(E)$. If $\mathrm{End}(E)$ is an order in an imaginary quadratic field, then E is called *ordinary*; if $\mathrm{End}(E)$ is a maximal order in the quaternion algebra $B_{p,\infty}$, then E is *supersingular*. Every supersingular elliptic curve in characteristic p is $\overline{\mathbb{F}}_p$-isomorphic to one defined over $\mathbb{F}_{p^2}$, since its j-invariant lies in $\mathbb{F}_{p^2}$. We denote by $\mathcal{S}_p$ the set of supersingular elliptic curves over $\mathbb{F}_{p^2}$ up to $\mathbb{F}_{p^2}$-isomorphism; it has size about $\lfloor p/12 \rfloor$.

For a prime $\ell \neq p$, the supersingular ℓ-isogeny graph $\mathcal{G}(p,\ell)$ has vertex set $\mathcal{S}_p$, with one edge for each ℓ-isogeny. Since every curve has exactly $\ell+1$ subgroups of order ℓ, the graph is $(\ell+1)$-regular. Moreover, it is Ramanujan, so random walks mix rapidly, and any two vertices are connected by an isogeny of degree ℓ^m with $m = O(\log p)$ [Koh96].

We also consider supersingular curves with level structure; see [Arp24] for background. Let N be square-free and coprime to p, and define $\mathcal{S}_p(N)$ to be the set of equivalence classes of pairs (E,G), where $E/\mathbb{F}_{p^2}$ is supersingular and $G \subseteq E[N]$ is cyclic of order N. Two pairs (E_1,G_1) and (E_2,G_2) are equivalent if there is an isomorphism $\iota : E_1 \to E_2$ with $\iota(G_1) = G_2$. We write $\mathcal{G}(p,\ell,N)$ for the corresponding ℓ-isogeny graph with level-N structure.

To such a pair $(E,G) \in \mathcal{S}_p(N)$, we associate the order

$$\mathcal{O}(E,G) = \{\alpha \in \mathrm{End}(E) : \alpha(G) \subseteq G\},$$

consisting of endomorphisms preserving the chosen subgroup.

2.2 Quaternion Algebras and the Deuring Correspondence

For a detailed account on the arithmetic of quaternion algebras we refer the reader to [Voi21]. For $a,b \in \mathbb{Q}^\times$, let $B(a,b)$ denote the quaternion algebra over $\mathbb{Q}$, with basis $1,i,j,ij$, i.e.

$$B(a,b) = \mathbb{Q} + \mathbb{Q}i + \mathbb{Q}j + \mathbb{Q}ij,$$

such that $i^2 = a$, $j^2 = b$ and $ij = -ji$.

An order $\mathcal{O}$ in B is a full rank lattice that is also a subring. It is maximal if it is not contained in any other order. For a lattice Λ we define the *left order* and *right order* of Λ to be

$$\mathcal{O}_L(\Lambda) = \{\alpha \in B |\ \alpha\Lambda \subseteq \Lambda\}, \quad \mathcal{O}_R(\Lambda) = \{\alpha \in B |\ \Lambda\alpha \subseteq \Lambda\}.$$

If $\mathcal{O}$ is a maximal order and I is a left $\mathcal{O}$-ideal, then $\mathcal{O}_R(I)$ is also a maximal order. Given two maximal orders $\mathcal{O}$ and $\mathcal{O}'$, then there exist a lattice I, called a connecting ideal, such that $\mathcal{O}_L(I) = \mathcal{O}$ and $\mathcal{O}_R(I) = \mathcal{O}'$. Additionally, we say left $\mathcal{O}$-ideals I and J are equivalent if there exists an $\alpha \in B^\times$ such that $I = \alpha J$. The set of classes of this equivalence relation is called the *left*-ideal class set of $\mathcal{O}$ and is denoted by $\mathrm{Cl}(\mathcal{O})$.
The endomorphism rings of supersingular elliptic curves in characteristic p are maximal orders in the unique quaternion algebra $B_{p,\infty}$ over $\mathbb{Q}$ ramified exactly at p and ∞. An Eichler order $\mathfrak{O}$ is an intersection $\mathfrak{O} = \mathcal{O}_1 \cap \mathcal{O}_2$ for two (not necessarily distinct) maximal orders $\mathcal{O}_1$ and $\mathcal{O}_2$. The level N of an Eichler order is given by its index in one of the maximal orders whose intersection defines the order, i.e. $N = [\mathcal{O}_1 : \mathfrak{O}] = [\mathcal{O}_2 : \mathfrak{O}]$ (this index will be the same for either order). Trivially, a maximal order is an Eichler order of level $N = 1$.

The Deuring Correspondence. For a detailed account of on the Deuring correspondence, see Chapter 42 of [Voi21]. Let E be a supersingular elliptic curve defined over $\mathbb{F}_{p^2}$, Deuring

showed that endomorphism rings of supersingular elliptic curves are in bijection with maximal order in $B_{p,\infty}$ up to conjugation with norms and traces on the endomorphisms side correspond to norms and traces in the quaternion algebra.

Deuring's correspondence establishes that, for a supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$, every isogeny $\varphi : E \to E'$ of degree n can be associated with a left ideal I in $\mathrm{End}(E)$ whose norm is n. According to Kohel's work [Koh96, Section 5.3], every nonzero left ideal of $\mathrm{End}(E)$ arises from such an isogeny.

We now look more closely at the analogous structure for Eichler Orders. The next result gives a characterization of Eichler orders in terms of supersingular elliptic curves with level-N structure.

Theorem 1. *For every $(E, G) \in \mathcal{S}_p(N)$, the order $\mathcal{O}(E, G)$ is isomorphic to an Eichler order of level N. Conversely, every Eichler order of level N is isomorphic to $\mathcal{O}(E, G)$ for some $(E, G) \in \mathcal{S}_p(N)$.*

Proof. See Theorem 3.7 and proposition 3.8 of [Arp24]. □

We also define ℓ-isogeny graphs with level-N structure. As we'll see, these graphs will also satisfy our delocalization results. Let $(E, G) \in \mathcal{S}_p(N)$ with $G \subseteq E[N]$ and $|G| = N$, and let $\varphi : E \to E'$ be an isogeny of prime degree ℓ with $\gcd(\ell, pN) = 1$. Then $\varphi(G) \subseteq E'[N]$, so φ induces a morphism in $\mathcal{S}_p(N)$,

$$\varphi : (E, G) \longrightarrow (E', \varphi(G)).$$

Just as isogenies between supersingular curves correspond to connecting ideals between their endomorphism rings, isogenies between curves equipped with level-N structure correspond to connecting ideals between the associated Eichler orders. In particular, an isogeny $\varphi : (E, G) \to (E', \varphi(G))$ corresponds to an ideal I_φ connecting $\mathcal{O}(E, G)$ and $\mathcal{O}(E', \varphi(G))$. For details, see Propositions 3.9 and 3.10 of [Arp24].

Group Actions. We briefly recall the notion of oriented supersingular elliptic curves and the associated class group action. Let K be a quadratic number field and $\mathfrak{O}$ an order in K. A K-orientation on a supersingular elliptic curve E is an embedding

$$\iota : K \hookrightarrow \mathrm{End}(E) \otimes_{\mathbb{Z}} \mathbb{Q}.$$

An $\mathfrak{O}$-orientation is a K-orientation such that $\iota(\mathfrak{O}) \subseteq \mathrm{End}(E)$, and it is called *primitive* if

$$\iota(\mathfrak{O}) = \iota(K) \cap \mathrm{End}(E).$$

We write (E, ι) for an elliptic curve equipped with such an orientation and assume from now on that all orientations discussed are primitive.

Any isogeny $\phi : E \to F$ induces an orientation on F. If $\hat{\phi} : F \to E$ denotes the dual isogeny, then the pushforward orientation is defined by

$$\phi_*(\iota)(\alpha) = \left(\phi \circ \iota(\alpha) \circ \hat{\phi}\right) \otimes \frac{1}{\deg \phi}.$$

Hence, for two K-oriented elliptic curves (E_1, ι_1) and (E_2, ι_2), an isogeny $\phi : E_1 \to E_2$ is called K-oriented if $\phi_*(\iota_1) = \iota_2$.

Let $\mathrm{SS}_{\mathfrak{O}}(p)$ denote the set of isomorphism classes of supersingular elliptic curves over $\overline{\mathbb{F}}_p$ equipped with a $\mathfrak{O}$-orientation. Under the standard assumptions that p does not divide the conductor of $\mathfrak{O}$ and does not split in K, this set is nonempty [Onu21, Prop. 3.2]. We define ideal-induced isogenies by the following: for $(E, \iota) \in \mathrm{SS}_{\mathfrak{O}}(p)$ and an ideal $\mathfrak{a} \subseteq \mathfrak{O}$, set

$$E[\mathfrak{a}] = \bigcap_{\alpha \in \mathfrak{a}} \ker(\iota(\alpha)).$$

This subgroup determines an isogeny $\phi_{\mathfrak{a}} : E \to E_{\mathfrak{a}}$, and thus a new oriented curve

$$\mathfrak{a} * (E, \iota) := \big(E_{\mathfrak{a}}, (\phi_{\mathfrak{a}})_*(\iota)\big).$$

In this way one obtains an action

$$\mathrm{Cl}(\mathfrak{O}) \times \mathrm{SS}_{\mathfrak{O}}(p) \longrightarrow \mathrm{SS}_{\mathfrak{O}}(p), \qquad (\mathfrak{a}, (E, \iota)) \longmapsto \mathfrak{a} * (E, \iota).$$

When p does not split in K and does not divide the conductor of $\mathfrak{O}$, the class group action of $\mathrm{Cl}(\mathfrak{O})$ on $SS_{\mathfrak{O}}(p)$ is free and transitive [Onu21, Theorem 3.4]. Moreover, recent advances show that this action on $\mathfrak{O}$-oriented curves can be evaluated efficiently, notably via the CLAPOTIS algorithm [PR23] and PEGASIS [DEF+25].

2.3 Brandt Operators of Isogeny Graphs

For background on this section, see Chapter 41 of [Voi21]. Let $\mathcal{S}_p(N)$ be the set of supersingular elliptic curves over $\mathbb{F}_{p^2}$ with level-N structure, and fix a prime $\ell \neq p$. The adjacency operator A_ℓ of the supersingular ℓ-isogeny graph $\mathcal{G}(p, \ell, N)$ acts on $\mathbb{C}^{\mathcal{S}_p(N)}$, with basis $\{\delta_{(E,G)}\}$, by

$$(A_\ell F)(E, G) = \sum_{(E,G)\to(E',G')} F(E', G'),$$

where the sum runs over all ℓ-isogenies $\phi : E \to E'$ such that $G' = \phi(G)$, counted with multiplicity. Since each vertex has $\ell + 1$ outgoing ℓ-isogenies, the graph is $(\ell + 1)$-regular.

Via Deuring's correspondence, this graph admits an algebraic description in terms of Brandt operators. Let $\mathfrak{O} \subset B_{p,\infty}$ be an Eichler order of level N coprime to p, and let

$$\mathrm{Cl}(\mathfrak{O}) = \{[I_1], \ldots, [I_h]\}$$

be its left ideal class set, in bijection with $\mathcal{S}_p(N)$. The ℓ-th Brandt matrix $B_\ell = (b_{ij})$ is defined by

$$b_{ij} = \#\{\, J \subseteq I_j : \mathrm{Nrd}(J) = \ell\, \mathrm{Nrd}(I_j) \text{ and } [J] = [I_i]\,\},$$

or equivalently, for $f : \mathrm{Cl}(\mathfrak{O}) \to \mathbb{C}$,

$$(B_\ell f)([I]) = \sum_{\substack{J \subseteq I \\ \mathrm{Nrd}(J) = \ell\, \mathrm{Nrd}(I)}} f([J]).$$

Under Deuring's correspondence, B_ℓ agrees with the adjacency operator A_ℓ.

Because of extra automorphisms at $j = 0$ and $j = 1728$, A_ℓ need not be symmetric for the standard inner product. Writing

$$w(E, G) = |\operatorname{Aut}(E, G)/\{\pm 1\}| \in \{1, 2, 3\}, \qquad D = \operatorname{diag}\big(w(E, G)\big)_{(E,G)\in\mathcal{S}_p(N)},$$

the conjugated operator

$$A'_\ell = D^{1/2} B_\ell D^{-1/2} \tag{1}$$

is real symmetric; see [Voi21, §§40.1.9–40.1.10]. Hence A'_ℓ and B_ℓ are similar and have the same spectrum. Moreover, the operators $\{A'_\ell\}$ commute and therefore admit a common eigenbasis.

In particular, if $p \equiv 1 \pmod{12}$, then no supersingular curve has $j = 0$ or 1728, so $w(E, G) = 1$ for all (E, G), and thus $A_\ell = B_\ell$. This assumption is often imposed in isogeny-based cryptography to avoid exceptional automorphisms, loops, and multiple edges. See also Appendix A of [PW24] for more details on (1).

3 Automorphic Forms

In this section, we review several aspects of the theory of automorphic forms. We focus in particular on quaternionic automorphic forms, viewed both as finite-dimensional vector spaces of functions on quaternionic ideal class sets and as functions on 3-spheres. We then recall the classical correspondence between these two perspectives and establish the sup-norm bounds and the ε-separation of Hecke eigenvalues required for our applications. These results will play a fundamental role in our subsequent analysis.

3.1 Space of Quaternionic forms

For background, see [DV13, Marar]. Let $B = B_{p,\infty}$ be the definite quaternion algebra over $\mathbb{Q}$, let $\mathcal{O}$ be an Eichler order of level N coprime to p, and let

$$\mathrm{Cl}(\mathcal{O}) = \{[I_1], \ldots, [I_h]\}$$

be the set of invertible left $\mathcal{O}$-ideal classes, where $h = |\mathrm{Cl}(\mathcal{O})|$.

Let $\hat{\mathbb{Z}} = \prod_p \mathbb{Z}_p$, and let $\mathcal{A}_f = \prod_p' \mathbb{Q}_p$ be the finite adeles of $\mathbb{Q}$. Write

$$\hat{B} = B \otimes_{\mathbb{Q}} \mathcal{A}_f, \qquad \hat{\mathcal{O}} = \mathcal{O} \otimes_{\mathbb{Z}} \hat{\mathbb{Z}},$$

and define

$$H = H(N) := (1 + N\hat{\mathcal{O}}) \cap \hat{\mathcal{O}}^\times \subset \hat{B}^\times.$$

Then the double coset space

$$B^\times \backslash \hat{B}^\times / H$$

is finite and naturally identified with $\mathrm{Cl}(\mathcal{O})$. Choosing representatives $x_1, \ldots, x_h$, we identify functions on $\mathrm{Cl}(\mathcal{O})$ with functions on $\hat{B}^\times$ that are left $B^\times$-invariant and right H-invariant.

The space of automorphic forms of level H is

$$A(\mathcal{O}) := L^2(B^\times \backslash \hat{B}^\times / H) \cong \{\varphi : \mathrm{Cl}(\mathcal{O}) \to \mathbb{C}\},$$

a finite-dimensional Hilbert space equipped with the Petersson inner product

$$\langle \varphi, \psi \rangle = \int_{B^\times \backslash \hat{B}^\times / H} \varphi(x)\overline{\psi(x)}\, d\mu(xH),$$

where μ is a Haar measure on $\hat{B}^\times$.

Let ℓ be a prime unramified in B, so that $\mathcal{O}_\ell \simeq M_2(\mathbb{Z}_\ell)$. Let $a_\ell \in \hat{B}^\times$ be the adele equal to 1 at all places away from ℓ, and equal to $\begin{pmatrix} \ell & 0 \\ 0 & 1 \end{pmatrix}$ at ℓ. If $Ha_\ell H = \bigsqcup_i x_i H$, then the Hecke operator $T_\ell : A(\mathcal{O}) \to A(\mathcal{O})$ is defined by

$$T_\ell(\varphi)(x) = \sum_{x_i H \in Ha_\ell H/H} \varphi(B^\times x x_i H),$$

and its adjoint is

$$T_\ell^*(\varphi)(x) = \sum_{x_i H \in Ha_\ell^{-1} H/H} \varphi(B^\times x x_i H).$$

For $\ell \nmid N$, the operators T_ℓ form a commuting family of self-adjoint operators with respect to the Petersson inner product.

These Hecke operators are the adelic counterparts of the Brandt operators arising from supersingular isogeny graphs; in fact, as the next proposition shows, the two notions coincide.

Proposition 2 ([CGL08, Lemma 3.7])**.** *The Hecke operator T_ℓ defined with respect to the indicator δ-functions is equal to the Brandt matrix B_ℓ.*

3.2 The Sup-Norm problem of Hecke eigenforms

Here we study the so-called sup-norm problem. Let T_ℓ be a Hecke operator on

$$A(\mathcal{O}) = L^2\Big(B^\times \backslash \hat{B}^\times / H\Big).$$

Since T_ℓ is self-adjoint with respect to the Petersson inner product, it admits an orthonormal eigenbasis (of *eigenforms*) consisting of functions $\varphi : \mathrm{Cl}(\mathcal{O}) \to \mathbb{C}$. A classical question in analytic number theory concerns the *sup-norm* of such forms. Simply put, for a normalized Hecke eigenform

$$\varphi : \mathrm{Cl}(\mathcal{O}) \to \mathbb{C}, \qquad [I_i] \longmapsto \varphi([I_i]),$$

or equivalently a function on $B^\times \backslash \hat{B}^\times / H$, the sup-norm problem is to obtain an upper bound for

$$\|\varphi\|_\infty \;=\; \max_{1 \le i \le h} |\varphi([I_i])|.$$

The sup-norm problem fits into the broader circle of ideas known as *arithmetic quantum chaos*: one expects high-energy Hecke eigenfunctions to behave "randomly" and, in particular, to avoid strong concentration. Quantitatively, this is reflected both in bounds for the L^∞-norm and in statements about the distribution of the L^2-mass ("energy") of eigenfunctions [Sar93, IS95]. In this direction, Rudnick and Sarnak initiated a systematic study of quantum limits for arithmetic eigenstates and formulated the *quantum unique ergodicity* (QUE) philosophy in the arithmetic setting [RS94].
On arithmetic surfaces, Lindenstrauss proved arithmetic QUE for compact congruence surfaces (via measure classification) [Lin06], and Soundararajan ruled out "escape of mass" for the modular surface; together these establish the Rudnick–Sarnak QUE conjecture in that case [Sou10]. In the holomorphic setting, Holowinsky and Soundararajan proved the Rudnick–Sarnak mass equidistribution conjecture for Hecke eigenforms [HS10].

In the definite quaternionic setting, the operators T_ℓ are realized by Brandt operators on the finite set $\mathrm{Cl}(\mathcal{O})$, so an eigenform $\varphi : \mathrm{Cl}(\mathcal{O}) \to \mathbb{C}$ may be viewed as an eigenvector of a regular "Brandt graph." From this perspective, bounding $\|\varphi\|_\infty$ is a quantitative *delocalization* statement: it prevents the eigenvector from concentrating on a small subset of vertices. This viewpoint parallels the delocalization and quantum ergodicity theory for eigenvectors on large regular graphs [ALM15].

Some of the first sup-norm bounds in our setting of automorphic forms on

$$X \;=\; B^\times \backslash \hat{B}^\times / H,$$

are due to Blomer and Michel [BM11, Theorem 1]. Specifically, for an L^2-normalized Hecke eigenform φ they prove that there exists $\delta < 1/60$ such that

$$\|\varphi\|_\infty \ll \mathrm{Vol}(X)^{-\delta} \;=\; |\mathrm{disc}(\mathcal{O})|^{-\delta/2} \asymp (pN)^{-\delta},$$

using that $\mathrm{Vol}(X) = \mathrm{disc}^{1/2}(\mathcal{O}) = pN$ where $\mathcal{O} \subset B_{p,\infty}$ again is an Eichler order of level N. This was subsequently strengthened in [BM13, Theorem 1], where they refined the exponent and obtained, for any $\varepsilon > 0$,

$$\|\varphi\|_\infty \ll_\varepsilon \mathrm{Vol}(X)^{-1/6+\varepsilon} \;=\; (pN)^{-1/6+\varepsilon}.$$

The sharpest sup-norm bound currently known in this setting is due to Khayutin, Nelson, and Steiner [KNS24], which we record for later use.

Theorem 2. [KNS24, Corollary 2.3] *Let $\varphi \in A(\mathcal{O})$ be an L^2-normalized Hecke eigenform. Then for any $\varepsilon > 0$,*

$$\|\varphi\|_\infty \ll_\varepsilon \mathrm{Vol}(X)^{-1/4+\varepsilon}.$$

Remark. Norms in [KNS24] are taken with respect to a *probability measure* (i.e. total mass 1). In our setting $X = B^\times \backslash \hat{B}^\times / H$ is finite, and the Haar-induced measure μ satisfies

$$\mathrm{Vol}(X) = \int_X 1 \, d\mu.$$

We set the probability measure $d\mu_{\mathrm{prob}} := \mathrm{Vol}(X)^{-1} d\mu$, so that $\int_X 1 \, d\mu_{\mathrm{prob}} = 1$. Consequently,

$$\|\varphi\|_{2,\mathrm{prob}}^2 = \int_X |\varphi|^2 \, d\mu_{\mathrm{prob}} = \frac{1}{\mathrm{Vol}(X)} \int_X |\varphi|^2 \, d\mu = \frac{1}{\mathrm{Vol}(X)} \|\varphi\|_2^2,$$

and switching between these normalizations rescales φ (hence $\|\varphi\|_\infty$) by a factor $\mathrm{Vol}(X)^{1/2}$.

We now prove our main delocalization theorem of eigenvectors of supersingular ℓ-isogeny graph. For the sake of emphasizing the characteristic p we are working over and the structure of the quaternion algebra, we will denote by $\mathcal{G}(p, \ell, N)$ the supersingular ℓ-isogeny graph with N-level structure where again $(N, \ell p) = 1$.

Theorem 3. *Let A_ℓ be the weighted adjacency matrix of the supersingular ℓ-isogeny graphs with level structure $\mathcal{G}(p, \ell, N)$ given by 1. If v is a normalized eigenvector of $\{A_\ell\}_\ell$, then*

$$\|v\|_\infty \;=\; \max_{E \in \mathcal{G}(p,\ell)} |v(E)| \;\ll_\varepsilon\; (pN)^{-1/4+\varepsilon},$$

where $v(E)$ denotes the $(E, G)^{th}$ entry of v.

Proof. Write v in the δ-indicator basis indexed by the vertices of $\mathcal{G}(p, \ell, N)$. By the identity (1), the vector $w := D^{-1/2} v$ is an eigenvector of the Brandt matrix B_ℓ in the indicator basis of $\mathrm{Cl}(\mathcal{O})$, where D is the weighted diagonal matrix. By Proposition 2, w corresponds to a Hecke eigenform $\varphi \in A(\mathcal{O})$ satisfying $\varphi([I_i]) = w([I_i])$ for all i. Hence, by Theorem 2,

$$\|w\|_\infty \;=\; \|\varphi\|_\infty \;=\; \max_{[I_i] \in \mathrm{Cl}(\mathcal{O})} |\varphi([I_i])| \;\ll_\varepsilon\; (pN)^{-1/4+\varepsilon}.$$

Finally, since $v = Dw$ and $\|D\|_\infty = 3$, we obtain

$$\|v\|_\infty \le \|D\|_\infty^{1/2} \, \|w\|_\infty \ll_\varepsilon (pN)^{-1/4+\varepsilon},$$

as claimed. □

This theorem establishes a strong delocalization property of the isogeny graphs: an L^2-normalized eigenvector of the adjacency operator cannot concentrate a large proportion of its mass on a small set of vertices. Equivalently, the eigenmodes of the random walk on $\mathcal{G}(p, \ell)$ are forced to be spread across many vertices rather than exhibiting sharp spikes.

However, we expect more from these isogeny graphs. As pointed out in [BM13, p. 4], the strongest conceivable upper bound for the sup-norm is

$$\|\varphi\|_\infty \ll_\varepsilon \mathrm{Vol}(X)^{-1/2+\varepsilon} = (pN)^{-1/2+\varepsilon}.$$

From a graph-theoretic perspective, such a bound would imply not merely strong delocalization, but *complete delocalization* of eigenvectors. This so-called *sup-norm conjecture*, highlighted in [BM13] and subsequently echoed in [KNS24], remains a long-standing and widely believed open problem in analytic number theory. While no proof is currently known, we provide further supporting evidence for the sup-norm conjecture in the case of the regular ℓ-isogeny graph, that is when $N = 1$, which is the case of interest.

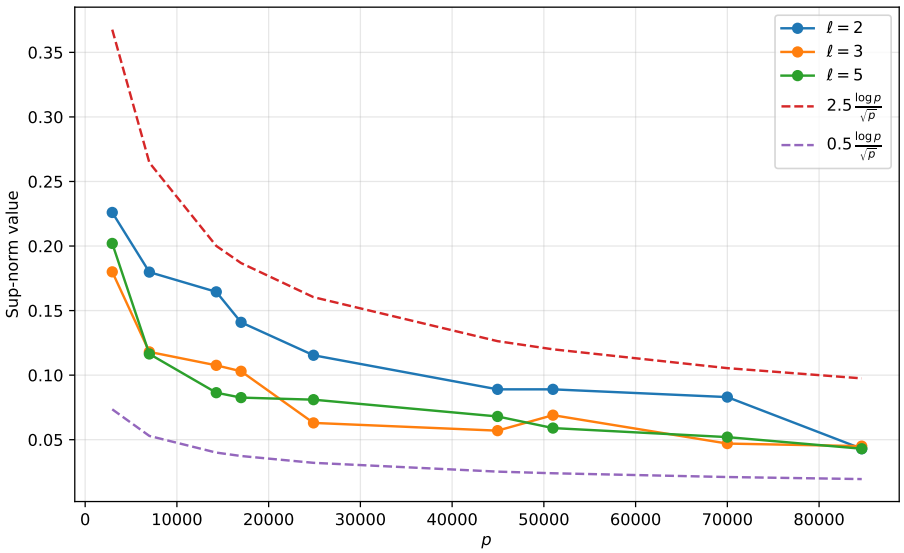


Figure 1: Sup-norm values versus $\frac{\log p}{\sqrt{p}}$ for $\ell = 2, 3, 5$.

Conjecture 1 (Complete delocalization)**.** *Let A_ℓ denote the weighted adjacency matrix of the supersingular ℓ-isogeny graph $\mathcal{G}(p, \ell)$ for all $\ell \neq p$. Then any L^2-normalized simultaneous eigenvector v of A_ℓ satisfies*

$$\|v\|_\infty \ll \frac{\log p}{\sqrt{p}}.$$

In Figure 1, we computed the sup-norm of each simultaneous orthonormal eigenvector of A_ℓ for $\ell = 2, 3, 5$, then took the maximum over all of them and plotted the results. The matrices were computed using Sage's built-in functions `BrandtModule(`p, ℓ`)` and `hecke_matrix`, and subsequently analyzed in MATLAB for improved stability and efficiency. As shown, the computations indicate that the sup-norm satisfies the conjecture, with implicit constants ranging roughly between 0.5 and 2.5.

Remark 1. We emphasize that the above bound is not merely a cryptographic heuristic, but a well-studied conjecture in analytic number theory, which experts in the area have indicated is widely expected to hold.

We now record the following useful theorem, which asserts that eigenvectors of the weighted adjacency matrices A_ℓ are completely delocalized, *on average*. For a vector v, we write $v(x)$ for its x-th coordinate.

Theorem 4 (Fourth Moment bound)**.** *Let A_ℓ denote the weighted adjacency matrices of the graph $\mathcal{G}(p, \ell, N)$, and let n be the number of vertices. For any two vertices $x, y \in \mathcal{G}(p, \ell, N)$, we have*

$$\sum_{i=1}^{n} \left(v_i(x)^2 - v_i(y)^2\right)^2 \ll_\varepsilon (pN)^{-1+\varepsilon},$$

where $\{v_i\}_{i=1}^n$ is a simultaneous orthonormal basis of eigenvectors of $\{A_\ell\}$.

Proof. This follows directly from Theorem 2.1 of [KNS24], since the weighted adjacency matrices A_ℓ satisfy the same hypotheses in the elliptic curve setting. The proof proceeds exactly as in Theorem 3.

□

3.3 ε-Separation of Hecke Eigenvalues

Fix a prime p and let $A_{\ell_1}, \ldots, A_{\ell_r}$ be the weighted adjacency matrices of the supersingular ℓ_i-isogeny graphs over $\mathbb{F}_{p^2}$ as in (1). These matrices are real and symmetric. For each prime ℓ, define

$$U_\ell := e^{iA_\ell/\sqrt{\ell}}.$$

Let v_k, v_j be simultaneous normalized eigenvectors of all A_{ℓ_i} (and hence of all U_{ℓ_i}). For each eigenvector v_j, define the eigenvalue vector

$$\boldsymbol{\lambda}_j := (\lambda_{1,j}, \ldots, \lambda_{r,j}),$$

where $\lambda_{i,j}$ is the eigenvalue of A_{ℓ_i} corresponding to v_j.

Definition 1. The operators $\{A_\ell\}$ are *ε-separated* if for any two distinct simultaneous eigenvectors v_k, v_j,

$$\|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_j\|_2 \geq \varepsilon.$$

Our goal is to obtain explicit lower bounds on $\|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_j\|_2$. Readers interested only in the cryptographic relevance to applications can consult Theorem 5 and Proposition 3. The operators U_ℓ are introduced for applications: since A_ℓ is symmetric, U_ℓ is unitary and can be implemented as a quantum gate.

By the Jacquet–Langlands correspondence, the Brandt operators on the supersingular module coincide with the Hecke operators T_ℓ acting on $S_2(\Gamma_0(p))$ with the Petersson inner product. Thus the eigenvectors of $\{A_\ell\}$ correspond to normalized weight-2 Hecke eigenforms $f_j \in S_2(\Gamma_0(p))$, with

$$A_\ell f_j = a_\ell(f_j) f_j, \qquad U_\ell f_j = e^{ia_\ell(f_j)/\sqrt{\ell}} f_j.$$

We associate to f_j the *Hecke tag vector*

$$\mathbf{t}_j := (\lambda_{\ell_1}(f_j), \ldots, \lambda_{\ell_r}(f_j)), \qquad \lambda_\ell(f) := \frac{a_\ell(f)}{\sqrt{\ell}} \in [-2, 2],$$

where the bound follows from Deligne's theorem [Del80, Piz98].

The *eigenvalue separation problem* asks for r and $\varepsilon > 0$ such that

$$f_j \neq f_k \;\Rightarrow\; \|\mathbf{t}_j - \mathbf{t}_k\|_2 \geq \varepsilon.$$

This enables identification of one-dimensional eigenspaces from approximate measurements, a requirement in many schemes (see e.g. [Kan18, KSS22, BBD+24, Dol23]), which currently rely on heuristic separation assumptions.

At the level of uniqueness, analytic multiplicity-one results show that if $f \neq g$ are Hecke eigenforms then $a_n(f) \neq a_n(g)$ for some $n \ll p \log p$, and under GRH one may take

$$r = O\big((\log p)^2 (\log\log p)^4\big)$$

[GH93, Thms. 2–3]. However, these results give only non-equality and do not yield a quantitative lower bound on $\|\mathbf{t}_f - \mathbf{t}_g\|_2$. Our main result provides such a quantitative eigenvalue separation.

Theorem 5 (GRH). *Let p be a prime, and let f_k and f_j be two distinct normalized and simultaneous cuspidal Hecke eigenforms in $S_2(\Gamma_0(p))$. Assume the Generalized Riemann Hypothesis for the relevant Rankin–Selberg L-functions. If Y and X satisfy $X \gg Y \gg (\log p)^4$, and $\ell_1, \cdots, \ell_r$ are the primes in the interval $[Y, X]$, then*

$$\|\mathbf{t}_k - \mathbf{t}_j\|_2 \geq 1.$$

In order to prove Theorem 5, we introduce several tools from analytic number theory. We refer the reader to sections 1.6 and 1.8 of [Bum09]. We note that the following discussion can also be examined from a representation-theoretic perspective. We let f_k and f_j be normalized cuspidal Hecke eigenforms of weight 2 and level p (with trivial nebentypus).

The Rankin–Selberg L-function $L(s, f_k \times \overline{f_j})$ is the degree 4 L-function whose Euler product (for $\Re(s) > 1$) is given by

$$L(s, f_k \times \overline{f_j}) = \prod_{\ell} \prod_{i,j=1}^{2} \Big(1 - \alpha_{\ell,i}(f_k)\,\overline{\alpha_{\ell,j}(f_j)}\,\ell^{-s}\Big)^{-1},$$

where $\{\alpha_{\ell,1}(f_k), \alpha_{\ell,2}(f_k)\}$ and $\{\alpha_{\ell,1}(f_j), \alpha_{\ell,2}(f_j)\}$ are the Satake parameters at $\ell \nmid p$ (so $\alpha_{\ell,1}(f_k) + \alpha_{\ell,2}(f_k) = \lambda_\ell(f_k)$ and $\alpha_{\ell,1}(f_k)\alpha_{\ell,2}(f_k) = 1$, and similarly for f_j). Equivalently, the logarithmic derivative of $L(s, f_k \times \overline{f_j})$ encodes the correlations $\lambda_\ell(f_k)\lambda_\ell(f_j)$:

$$-\frac{L'}{L}(s, f_k \times \overline{f_j}) = \sum_{n \geq 1} \frac{\Lambda(n)\, b_{k,j}(n)}{n^s}, \qquad b_{k,j}(\ell) = \lambda_\ell(f_k)\lambda_\ell(f_j) \quad (\ell \nmid p), \tag{2}$$

where Λ is the von Mangoldt function.

$L(s, f_k \times \overline{f_j})$ admits analytic continuation and a functional equation [IK04, chapter 5]. Moreover, if $f_k \neq f_j$ then $L(s, f_k \times \overline{f_j})$ is entire, whereas

$$L(s, f_k \times \overline{f_k}) = \zeta(s)\, L(s, \mathrm{sym}^2 f_k)$$

has a simple pole at $s = 1$, and similarly for f_j. These facts imply a sharp distinction between *diagonal* and *off-diagonal* prime sums: the diagonal terms have a main term of size X, while the off-diagonal correlation has no main term. We are now ready to prove Theorem 5.

Proof. Fix parameters $1 \leq Y < X$ and define the interval second moment

$$S(Y, X) := \sum_{Y < \ell \leq X} \big(\lambda_\ell(f_k) - \lambda_\ell(f_j)\big)^2 \log \ell.$$

Expanding gives

$$S(Y, X) = \sum_{Y < \ell \leq X} \lambda_\ell(f_k)^2 \log \ell + \sum_{Y < \ell \leq X} \lambda_\ell(f_j)^2 \log \ell - 2 \sum_{Y < \ell \leq X} \lambda_\ell(f_k)\lambda_\ell(f_j) \log \ell.$$

Under the Generalized Riemann Hypothesis for the relevant Rankin–Selberg L-functions, the corresponding prime number theorem is proved in full generality in [IK04, Theorem 5.15] (see also [Qu07]). A similar unconditional version is obtained in [LWY05]. By the discussion above, this implies that (for $i \in \{j, k\}$)

$$\sum_{\ell \leq T} \lambda_\ell(f_i)^2 \log \ell = T + O\big(T^{1/2} \log^2(QT)\big), \quad \sum_{\ell \leq T} \lambda_\ell(f_k)\lambda_\ell(f_j) \log \ll T^{1/2} \log^2(QT),$$

where Q is the analytic conductor (here $Q \asymp p^2$). Subtracting the estimates at $T = X$ and $T = Y$ yields

$$\sum_{Y < \ell \leq X} \lambda_\ell(f_i)^2 \log \ell = (X - Y) + O\big(X^{1/2} \log^2(QX) + Y^{1/2} \log^2(QY)\big),$$

and similarly

$$\sum_{Y < \ell \leq X} \lambda_\ell(f_k)\lambda_\ell(f_j) \log \ell = O\big(X^{1/2} \log^2(QX) + Y^{1/2} \log^2(QY)\big).$$

Consequently,

$$S(Y,X) = 2(X-Y) + O\big(X^{1/2}\log^2(QX)\big).$$

In particular, if $X \gg Y \gg (\log p)^4$, then the main term dominates and we may ensure

$$S(Y,X) \geq c(X-Y)$$

with $c > 1$ by taking the implicit constants sufficiently large.

On the other hand, if $|\lambda_\ell(f_k) - \lambda_\ell(f_j)| < c_0 < \sqrt{c}$ for all primes $Y < \ell \leq X$, where c_0 is picked such that

$$S(Y,X) < c_0^2 \sum_{Y<\ell\leq X} \log \ell < c(X-Y),$$

we get a contradiction. In other words, if the individual pointwise distance of eigenvalues is less than a constant, then the average must also be smaller. Hence there exists a prime $\ell \in (Y,X]$ with $\ell \nmid p$ such that

$$|\lambda_\ell(f_k) - \lambda_\ell(f_j)| \geq 1.$$

Since the tag set $\{\ell_1,\ldots,\ell_r\}$ contains all primes in $(Y,X]$, then this implies $\|\mathbf{t}_k - \mathbf{t}_j\|_2 \geq 1$. □

Since our applications concern the unitary operators $\{U_\ell\}_\ell$, we will also need the following proposition. Unlike the previously proposed separation guarantees in [KSS22], which are heuristic, our result is not only rigorous under GRH but yields a strictly stronger bound.

Proposition 3 (GRH). *Fix a prime p, and let $\ell_1,\cdots,\ell_r$ be the primes $\in [Y,X]$ with $X \asymp Y \asymp \log^4 p$ and $\ell_i \nmid p$. Then the unitary operators $\{U_{\ell_i}\}$ are ε-separated with $\varepsilon \gg 1$.*

Proof. Assuming the content of the proposition, the proof of theorem 5 shows that there exists a prime $\ell^* \in [Y,X]$ such that

$$|\lambda_{\ell^*}(f_j) - \lambda_{\ell^*}(f_k)| \geq 1 \quad \Rightarrow \quad |a_{\ell^*}(f_j) - a_{\ell^*}(f_k)| \geq \sqrt{\ell^*} \gg \log^2 p.$$

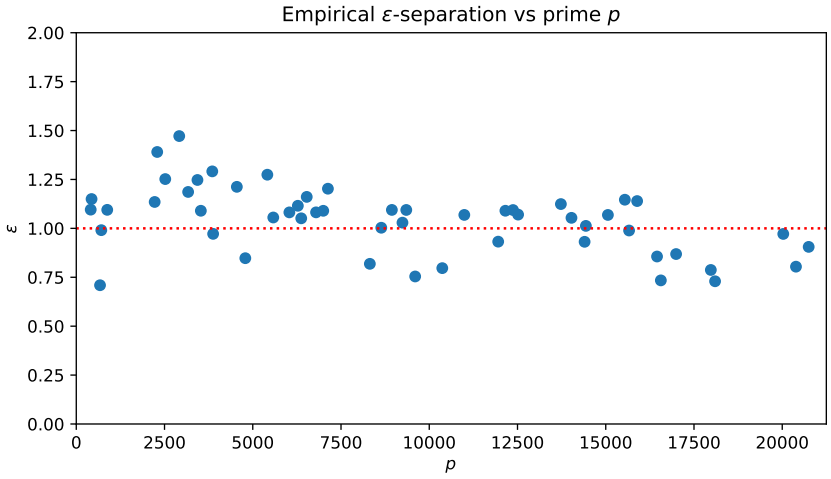
But the $a_\ell(f_i)'s$ coincide with the eigenvalues $\lambda_{\ell,i}$ of the corresponding eigenvector v_i under A_ℓ. Hence, for any simultaneous eigenvectors v_j and v_k of the $\{A_\ell\}$, we obtain

$$||\boldsymbol{\lambda}_j - \boldsymbol{\lambda}_k||_2 \gg \log^2 p,$$

i.e. $\{A_{\ell_i}\}$ are $(\log^2 p)$-separated. Proposition 6.15 of [KSS22] states that if $\{A_\ell\}$ are ε-separated then the $\{U_\ell\}$ are ε'-separated for $\varepsilon' = \varepsilon/\sqrt{\max \ell}$. Our theorem then follows. □ □

Remark 2. Using the publicly available GitHub implementation of Sharif [Sha25], we computed the vectors of eigenvalues $\boldsymbol{\lambda}_i = (\lambda_{i,\ell_1},\ldots,\lambda_{i,\ell_r})$ associated with the operators U_{ℓ_i}, where the primes satisfy $\log p < \ell_i < 4\log p$. Figure 2 displays the minimum Euclidean distance $\|\boldsymbol{\lambda}_i - \boldsymbol{\lambda}_j\|_2$ over all $i \neq j$.

The empirical results deserve a brief discussion. In [KSS22], the authors computed the ε-separation for primes p up to approximately 20,000, at all primes $\ell_i < \log p$. They conjecture that the resulting ε value should be on the order of $(4\log p)^{-1}$. By contrast, Proposition 3 predicts a separation, on average, bounded below by a positive constant, uniformly in p . This prediction is fully consistent with our computations, as Figure 2 shows that the observed values cluster around $\varepsilon = 1$. We suspect that the interval $[2,\log p]$ considered in [KSS22] was too narrow to capture the constant gap. In contrast, enlarging the range of primes to $\log p < \ell_i < 4\log p$ allows this behavior to be clearly observed. In particular, our results suggest that $\varepsilon = 0.25$ is a reasonable lower bound for the separation.



at different primes p

Figure 2: ε-separation for $\log p < \ell_i < 4 \log p$

Remark 3. We note a minor gap in the heuristic argument of [Dol23, Lemma 5.2]. Under the heuristic independence assumption on the tag vectors $\mathbf{t}_i$, the estimate there bounds the *pairwise* collision probability $\Pr[\|\mathbf{t}_i - \mathbf{t}_j\| \le \varepsilon]$ for any fixed pair (i, j). However, in their setting one needs a *global* (worst-case) guarantee that no collision occurs among *any* pair, i.e., a bound on $\Pr[\exists\, i \neq j : \|\mathbf{t}_i - \mathbf{t}_j\| \le \varepsilon]$. Applying the union bound yields $\Pr[\exists\, i \neq j : \|\mathbf{t}_i - \mathbf{t}_j\| \le \varepsilon] \le \sum_{1 \le i < j \le N} \Pr[\|\mathbf{t}_i - \mathbf{t}_j\| \le \varepsilon]$, which introduces an additional $\binom{N}{2}$ factor and therefore changes the quantitative bound stated in [Dol23, Lemma 5.2].

4 Quantum Sampling of Secure Curves

In this section, we present the main application of our theory: generating a random supersingular elliptic curve with unknown endomorphism ring. The main results are Theorems 6 and 8 for sampling from the full graph and Theorem 9 for oriented curves.

We briefly recall the quantum notions needed below; see [KLM07, Ch. 3–7].

Qubits and measurement. A qubit is a unit vector $|x\rangle = \alpha|0\rangle + \beta|1\rangle \in \mathbb{C}^2$ with $|\alpha|^2 + |\beta|^2 = 1$. An n-qubit register is a state

$$|x\rangle = \sum_{i=0}^{2^n-1} \alpha_i |i\rangle \in (\mathbb{C}^2)^{\otimes n}, \qquad \sum_i |\alpha_i|^2 = 1,$$

where $\{|0\rangle, \ldots, |2^n - 1\rangle\}$ is the computational basis. Measuring $|x\rangle$ in this basis yields $|i\rangle$ with probability $|\alpha_i|^2$.

Quantum algorithms. A quantum algorithm is a sequence of unitary operations. An oracle $O : X \to Y$ is implemented reversibly by

$$U_O : |x\rangle|y\rangle \mapsto |x\rangle|y \oplus O(x)\rangle.$$

Quantum Phase Estimation. Given a unitary U and an eigenstate $|\psi\rangle$ with $U|\psi\rangle = e^{2\pi i\lambda}|\psi\rangle$, for $\lambda \in [0, 1)$, Quantum Phase Estimation (QPE) outputs an m-bit estimate $\tilde{\lambda}$ of λ with error $O(2^{-m})$ with high probability, using controlled powers of U and an

inverse quantum Fourier transform. For a state $|\psi\rangle$, we write $\mathrm{QPE}_U^{(\varepsilon)}(|\psi\rangle)$ for the resulting entangled state after applying QPE to $U|\psi\rangle$ at precision ε.

4.1 Sampling from the Full Isogeny Graph

Throughout, we restrict to supersingular elliptic curves without level structure, although the same ideas extend naturally. Our algorithm 1 outputs a provably secure curve without generating an explicit isogeny path, in the sense that any QPT adversary has only negligible advantage in computing its endomorphism ring. Moreover, using interactive protocols to verify the quantum computations, e.g. [FK17, ABOEM17], one can verify if a curve is generated by an honest execution of the algorithm and thus the curve is verifiably correct and secure.

In what follows, our quantum states live in the Hilbert space $\mathbb{C}^K$, where K is the number of elliptic curves over $\mathbb{F}_{p^2}$ (including ordinary curves). Let V_m denote the subspace spanned by supersingular curves over $\mathbb{F}_{p^2}$. Using Proposition 1, we index the computational basis by pairs (j, b), where $j \in \mathbb{F}_{p^2}$ is the j-invariant and b denotes the twisting data, giving the basis $\{|(j, b)\rangle\}$.

It was noted to us that Montgomery and Sharif [MS25] previously used such a parametrization for curves over $\mathbb{F}_p$; here we adapt it to $\mathbb{F}_{p^2}$. Alternatively, one could encode curves by Weierstrass coefficients (a, b), but this would require roughly twice as many qubits.

We now outline the sampling algorithm. Let $|E_0\rangle$ be an initial vertex state corresponding to a supersingular curve in V_n, which can be efficiently prepared using [Brö09]. For a prime ℓ, let A_ℓ denote the weighted symmetric adjacency matrix of the supersingular ℓ-isogeny graph, and define

$$H_\ell := \frac{A_\ell}{\sqrt{\ell}}.$$

The matrices $\{A_\ell\}$ admit a common orthonormal eigenbasis $\{|\phi_i\rangle\}_{i=1}^n$, and consequently the unitaries $\{e^{itH_\ell}\}$ share this eigenbasis. Expanding the initial state, we write

$$|E_0\rangle = \sum_{i=1}^{n} \alpha_i |\phi_i\rangle, \qquad \alpha_i := \langle \phi_i | E_0 \rangle.$$

We first apply Quantum Phase Estimation (QPE) with respect to $e^{iH_{\ell_1}}$, obtaining

$$|\varphi\rangle = \sum_{i=1}^{n} \alpha_i |\tilde{\lambda}_{1,i}\rangle |\phi_i\rangle,$$

where $\tilde{\lambda}_{1,i}$ is an approximation of the eigenvalue $\lambda_{1,i}$ of H_{ℓ_1}. Measuring the eigenvalue register yields an outcome $\tilde{\lambda}_1$ and projects the state onto a corresponding subspace, resulting in a post-measurement state $|\psi_1\rangle$.

We then iterate this procedure: applying QPE with respect to $e^{iH_{\ell_k}}$ to $|\psi_{k-1}\rangle$ and measuring the eigenvalue register produces an approximation $\tilde{\lambda}_k$ and further refines the state to $|\psi_k\rangle$, supported on a progressively smaller subset of the eigenbasis.

By Proposition 3 (or its heuristic form in Remark 2), after sufficiently many iterations the tuple $(\tilde{\lambda}_1, \ldots, \tilde{\lambda}_r)$ uniquely identifies an eigenvector. At this stage, the state is approximately supported on a single eigenstate $|\phi\rangle$. Finally, measuring the vertex register in the basis $\{|(j, b)\rangle\}$ yields a curve $E \in \mathcal{G}(p, \ell)$ with probability $|\langle E|\phi\rangle|^2$.

Algorithm 1 QuantumSampling$(p; X, Y)$

Input: An odd prime p, an initial curve $E_0 \in \mathcal{S}$, and parameters $1 \le X \le Y$.
Output: A supersingular elliptic curve E encoded by (j, b).

Let $\{\ell_1, \dots, \ell_r\}$ be the primes in $[X, Y]$ (in any fixed order)
Initialize $|\psi_0\rangle \leftarrow |E_0\rangle$
$c \leftarrow \max(0.25, c_0)$, where c_0 is the ε-separation constant from Proposition 3
for $k = 1$ **to** r **do**
    Set $H_k \leftarrow A_{\ell_k}/\sqrt{\ell_k}$
    $|\psi_k\rangle \leftarrow \mathrm{QPE}_{e^{iH_k}}^{(c/2\sqrt{r})}(|\psi_{k-1}\rangle)$
    $|\psi_k\rangle \leftarrow$ post-measurement state after measuring the eigenvalue register of $|\psi_k\rangle$
end for
Measure the vertex register of $|\psi_r\rangle$ in the basis $\{|(j, b)\rangle\}$ to obtain E
return E

The remainder of the section analyzes (i) the gate complexity of QuantumSampling, in particular the number of controlled-e^{iH_k} calls required for the desired precision, and (ii) the distribution and security of the output curve. We note that the phase of the trivial eigenvalue may overlap modulo 2π with other phases, however we treat this issue in the security analysis rather than the running time of the algorithm.

Proposition 4. *Given primes p and $\ell \nmid p$, the gate complexity to execute one call to Quantum Phase Estimation $\mathrm{QPE}_{e^{iH}}^{\varepsilon}$ with failure probability η, where $H = A_\ell/\sqrt{\ell}$, is*

$$\tilde{O}\left(\log p \cdot \log\left(\frac{1}{\eta}\right)\left[\ell^{11/4}\left(\frac{1}{\varepsilon}\right)^{3/2} + \ell^{9/4}\left(\frac{1}{\varepsilon}\right)^{1/2}\left(\log\frac{1}{\eta} + \log\log\frac{1}{\varepsilon}\right)\right]\right).$$

Proof. Since the non-trivial eigenvalues of H lie in $[-2, 2]$ (Deligne's bound [Del80]), the eigenphases of U are in $[-2, 2]$ (as $2 < \pi$), hence there is no modular ambiguity between an eigenphase and the corresponding eigenvalue. For the trivial eigenvalue $\ell + 1$ this might indeed wrap $\bmod 2\pi$, however this doesn't change the cost of running QPE or Hamiltonian simulation.

Fix a target additive phase accuracy ε and an overall failure probability η. Standard QPE uses

$$q \;:=\; \left\lceil \log_2\left(\frac{4}{\varepsilon}\right)\right\rceil$$

phase bits, and thus requires controlled applications of the powers $U^{2^j} = e^{iH2^j}$ for $j = 0, 1, \dots, q-1$. Equivalently, one execution of QPE requires implementing controlled Hamiltonian evolutions for times

$$t_j \;=\; 2^j \qquad (0 \le j \le q-1).$$

To reduce the failure probability to at most η, we repeat QPE $R = \Theta(\log(1/\eta))$ times and output the median estimate.

We implement each e^{iHt} using sparse-Hamiltonian simulation as in Berry–Childs–Kothari [BCK17]. Their cost is expressed in terms of

$$\tau(t) \;:=\; d\,\|H\|_{\max}\,t,$$

where d is the sparsity and $\|H\|_{\max} = \max_{x,y}|H_{x,y}|$. For $H = A_\ell/\sqrt{\ell}$, we have $d = \ell + 1$ and $\|H\|_{\max} = 1/\sqrt{\ell}$, hence

$$\tau(t) \;=\; (\ell+1)\frac{t}{\sqrt{\ell}} \;=\; \kappa\,t, \qquad \kappa \;:=\; \frac{\ell+1}{\sqrt{\ell}}.$$

We allocate simulation error per controlled evolution as

$$\delta_{\text{sim}} \;:=\; \frac{\eta}{10Rq},$$

so that over the q controlled evolutions and R repetitions, a union bound ensures that the total probability of any simulation failure is at most $\eta/10$. Thus

$$\log \frac{1}{\delta_{\text{sim}}} \;=\; \log\Big(\Theta(Rq/\eta)\Big) \;=\; O\Big(\log(1/\eta) + \log q\Big).$$

Specializing [BCK17, Theorem 3] with $\alpha = 1$, each evolution e^{iHt} can be implemented within error δ_{sim} using $O\Big(\tau(t)^{3/2} + \tau(t)^{1/2}\log(1/\delta_{\text{sim}})\Big)$ oracle queries. Therefore, the j-th controlled power costs

$$Q_j \;=\; O\Big((\kappa 2^j)^{3/2} + (\kappa 2^j)^{1/2}\log(1/\delta_{\text{sim}})\Big).$$

Summing over $j = 0, \ldots, q-1$ and using geometric-series bounds yields

$$Q_{\text{QPE}} \;=\; O\Big(\kappa^{3/2}\, 2^{3q/2} \;+\; \kappa^{1/2}\, \log(1/\delta_{\text{sim}})\, 2^{q/2}\Big)\,.$$

Since $q = \lceil \log_2(4/\varepsilon) \rceil$ implies $2^{3q/2} = O((1/\varepsilon)^{3/2})$, we obtain

$$Q_{\text{QPE}} \;=\; O\left(\kappa^{3/2}\left(\frac{1}{\varepsilon}\right)^{3/2} \;+\; \kappa^{1/2}\left(\frac{1}{\varepsilon}\right)^{1/2} \log\frac{1}{\delta_{\text{sim}}}\right).$$

Finally, repeating $R = \Theta(\log(1/\eta))$ times gives the total query complexity

$$Q_{\text{total}} \;=\; O\left(\log\frac{1}{\eta}\left[\kappa^{3/2}\left(\frac{1}{\varepsilon}\right)^{3/2} \;+\; \kappa^{1/2}\left(\frac{1}{\varepsilon}\right)^{1/2}\left(\log\frac{1}{\eta} + \log\log\frac{1}{\varepsilon}\right)\right]\right).$$

At this point, we need to estimate the gate cost of one coherent call to compute the isometry

$$|E\rangle \mapsto \frac{1}{\sqrt{\ell+1}}\sum_{i=1}^{\ell+1}|E\rangle|E_i\rangle,$$

where the curves E_i are the ℓ-neighbors of E. If $\ell \nmid p-1$, the kernel generating our ℓ-isogenies are not defined over $\mathbb{F}_{p^2}$ and thus the most efficient way requires going to field extension of degree ℓ. A detailed analysis is given in [BBC+25, section 4.4]; we obtain a running time of $O(\ell^2 \log p)$ to compute the above isometry. Then the gate complexity to execute one QPE is $T = O\Big(Q_{\text{total}} \cdot \ell^2 \log p\Big)$; the result then follows. □

We now specialize our previous result to Algorithm 1.

Theorem 6. *Given a prime p and the primes $\ell_i \in [Y, X]$ where $\ell_i \neq p$, Algorithm 1 samples a random supersingular elliptic curve from V_m and has heuristic quantum gate complexity of $O(\log^{6.5} p)$ with $Y = \log p$ and $X = 4\log p$. Assuming the GRH, the algorithm has provable running time of $O(\log^{20} p)$ with $Y \asymp X \asymp \log^4 p$.*

Proof. First notice that the tags $\mathbf{t_i} = (\lambda_{1,i}, \cdots, \lambda_{r,i})$ uniquely identify the eigenvectors ϕ_i and are ε-separated for $\varepsilon = c_0$ or 0.25 depending on whether we use the GRH or the heuristics presented in remark 2. Hence if we use approximation tags $\tilde{\mathbf{t}}_i$ with accuracy $\varepsilon = c/(2\sqrt{r})$ each $\tilde{\mathbf{t}}_i$ will also uniquely correspond to an eigenvector (see [Dol23, lemma 3.3] for more details).

Specializing Proposition 4 to the heuristic regime, we choose $\eta = 1/p$, $\varepsilon = 1/2\sqrt{r}$, and $\ell_i = O(\log p)$, and hence the running time of one QPE call is $\tilde{O}(\log^{11/2} p)$. By the prime number theorm, we have

$$\gg \frac{X}{\log X} - \frac{Y}{\log Y},$$

many primes in $[Y, X]$, thus the heuristic running time becomes $\tilde{O}(\log^{6.5} p)$. The GRH-based regime can be analyzed similarly and we obtain a gate complexity of $\tilde{O}(\log^{20} p)$. □

We now turn to the problem of studying the distribution of the output curve. Since the proof is straightforward and won't play a crucial role in the security proof of the output curve, we defer it to appendix A.

Theorem 7. *Let p be an odd prime and let $E'/\mathbb{F}_{p^2}$ be any supersingular elliptic curve. For any $\varepsilon > 0$, the output E of Algorithm 1 satisfies*

$$0 < \Pr[E = E'] \ \ll_\varepsilon \ p^{-1+\varepsilon}$$

Security of the output curves. We now study the security of the output curve by Algorithm 1. The main challenge is that it is unclear whether the sequence of measured eigenvalues help an adversary compute the endomorphism ring of the output curve. This issue was raised in [BBD+24, Dol23], but to the best of our knowledge it has not been addressed in any prior work. Thanks to our delocalization results, the following theorem shows that, assuming the sup-norm conjecture 1, the output curve remains secure whenever EndRing is quantumly hard on average.

Theorem 8. *Let (E, Λ) be the output of Algorithm 1, where $E/\mathbb{F}_{p^2}$ is the output curve and Λ is the sequence of approximate eigenvalues measured. Let $n = \log p$. Assume Conjecture 1 holds, and let $\alpha := O(\log^2 p)$. Then for every QPT adversary $\mathcal{A}$ against (E, Λ), there exists a QPT adversary $\mathcal{B}$ against EndRing on a uniformly random supersingular curve E' such that*

$$\Pr[\mathcal{A}(E, \Lambda) \text{ outputs } \mathrm{End}(E)] \le \alpha \cdot \Pr_{E' \leftarrow U}[\mathcal{B}(E') \text{ outputs } \mathrm{End}(E')] + O(\alpha/\sqrt{p})$$

Proof. Let $m = |\mathcal{S}_p|$, and let U denote the uniform distribution on $\mathcal{S}_p$. Let P be the joint distribution of (E, Λ) in the real execution of Algorithm 1, initialized at $|E_0\rangle$. Let Q be the analogous distribution for the ideal projected experiment initialized at

$$|E_0^\perp\rangle = \frac{\Pi_{\mathrm{nt}}|E_0\rangle}{\|\Pi_{\mathrm{nt}}|E_0\rangle\|},$$

where Π_{nt} projects onto the orthogonal complement of the trivial eigenspace.

Since

$$|E_0\rangle = \frac{1}{\sqrt{m}}|\mathbf{1}\rangle + \sqrt{1 - \frac{1}{m}}|E_0^\perp\rangle,$$

the two initial states have trace distance $1/\sqrt{m} = O(1/\sqrt{p})$. As the whole algorithm, including QPE and the final measurement, is a quantum channel i.e. doesn'y increase the trace distance, the resulting output distributions satisfy

$$\Delta(P, Q) \le \delta_{\mathrm{triv}}, \qquad \delta_{\mathrm{triv}} = O(1/\sqrt{p}).$$

Now define, for any QPT adversary $\mathcal{A}$,

$$s_\lambda(y) := \Pr[\mathcal{A}(y, \lambda) \text{ outputs } \mathrm{End}(y)],$$

where the probability is over the internal randomness and measurements of $\mathcal{A}$.

In the projected experiment, the spectral analysis is carried out entirely on the nontrivial subspace. Each measured tag λ identifies a unique simultaneous eigenvector ϕ_λ, and hence $\Pr_{(E,\Lambda)\leftarrow Q}[E = y \mid \Lambda = \lambda] = |\langle y \mid \phi_\lambda \rangle|^2$. By Conjecture 1,

$$\|\phi_\lambda\|_\infty \leq \frac{C \log p}{\sqrt{p}}. \tag{3}$$

Therefore, for all $y \in \mathcal{S}_p$,

$$Q[E = y \mid \Lambda = \lambda] \leq \frac{C^2 \log^2 p}{p} \leq \frac{\alpha}{m}, \qquad \alpha = O(\log^2 p).$$

Conditioned on $\Lambda = \lambda$, its success probability is

$$\Pr[\mathcal{A} \text{ succeeds} \mid \Lambda = \lambda] = \sum_{y \in \mathcal{S}_p} \Pr[E = y \mid \Lambda = \lambda]\ s_\lambda(y).$$

Using inequality (3), we get

$$\Pr[\mathcal{A} \text{ succeeds} \mid \Lambda = \lambda] \leq \sum_{y \in \mathcal{S}_p} \frac{\alpha}{m}\, s_\lambda(y) = \alpha \cdot \frac{1}{m} \sum_{y \in \mathcal{S}_p} s_\lambda(y) = \alpha \cdot \mathbb{E}_{y \leftarrow U}[s_\lambda(y)], \tag{4}$$

where U is the uniform distribution on $\mathcal{S}_p$.

Averaging over Λ gives

$$\Pr_Q[\mathcal{A}(E, \Lambda) \text{ succeeds}] = \mathbb{E}_{\lambda \leftarrow \Lambda}\big[\Pr[\mathcal{A} \text{ succeeds} \mid \Lambda = \lambda]\big] \leq \alpha \cdot \mathop{\mathbb{E}}_{\lambda \leftarrow \Lambda} \mathop{\mathbb{E}}_{y \leftarrow U}[s_\lambda(y)].$$

Since the marginals P_Λ and Q_Λ also differ by at most δ_{triv} in total variation, and since $\mathbb{E}_{y \leftarrow U}[s_\lambda(y)] \in [0, 1]$, we have

$$\mathbb{E}_{\lambda \leftarrow Q_\Lambda} \mathbb{E}_{y \leftarrow U}[s_\lambda(y)] \leq \mathbb{E}_{\lambda \leftarrow P_\Lambda} \mathbb{E}_{y \leftarrow U}[s_\lambda(y)] + \delta_{\text{triv}}.$$

Now define a QPT adversary $\mathcal{B}$ for uniform EndRing as follows. On input $E' \leftarrow U$, sample $\lambda \leftarrow P_\Lambda$ by running Algorithm 1 once and discarding its curve output. Then run $\mathcal{A}(E', \lambda)$ and output whatever $\mathcal{A}$ outputs. By construction,

$$\Pr_{E' \leftarrow U}[\mathcal{B}(E') \text{ outputs } \operatorname{End}(E')] = \mathbb{E}_{\lambda \leftarrow P_\Lambda} \mathbb{E}_{y \leftarrow U}[s_\lambda(y)].$$

Finally, using $\Delta(P, Q) \leq \delta_{\text{triv}}$,

$$\Pr_P[\mathcal{A} \text{ succeeds}] \leq \Pr_Q[\mathcal{A} \text{ succeeds}] + \delta_{\text{triv}}.$$

Combining the previous inequalities gives

$$\Pr_P[\mathcal{A} \text{ succeeds}] \leq \alpha \Pr_{E' \leftarrow U}[\mathcal{B}(E') \text{ outputs } \operatorname{End}(E')] + (\alpha + 1)\delta_{\text{triv}}.$$

Since $\alpha = O(\log^2 p)$, $\delta_{\text{triv}} = O(1/\sqrt{p})$, and $n = \log p$, the claimed security bound now follows. □

4.2 Sampling Secure Oriented Curves

In this subsection, we give a quantum algorithm for sampling uniformly random secure $\mathfrak{O}$-oriented supersingular elliptic curves, for any imaginary quadratic order $\mathfrak{O}$, assuming the class group action is regular (that is, transitive and free). This is particularly relevant for protocols where the public key includes both a curve and an orientation, while the

endomorphism ring must remain hidden; see, e.g., [DFMPS19, LGDdSG21, BDF21]. The main result here is Theorem 9.
We first establish the necessary background and notation. We note that isogenies can be abstracted by considering only the underlying group action.

Let G be a finite abelian group. Its characters are the homomorphisms $\chi(a,\cdot) : G \to \mathbb{C}^\times$, indexed by $a \in G$. If

$$G \cong \mathbb{Z}_{N_1} \oplus \cdots \oplus \mathbb{Z}_{N_k},$$

then $\chi(a,x) = \omega_{N_1}^{a_1 x_1} \cdots \omega_{N_k}^{a_k x_k}$, where $\omega_L = \exp(2\pi i/L)$. The Fourier transform of a function $f : G \to \mathbb{C}$ is

$$\widehat{f}(a) = \frac{1}{\sqrt{|G|}} \sum_{x \in G} \chi(a,x) f(x),$$

and the quantum Fourier transform maps $\sum_{g\in G} f(g)|g\rangle$ to $\sum_{x\in G} \widehat{f}(x)|x\rangle$.

Now let $(G, \mathcal{X}, *)$ be a regular group action. For any subset $S \subseteq G$, any $y \in \mathcal{X}$, and any $h \in G$, define

$$|S^{(h)} * y\rangle = \frac{1}{\sqrt{|S|}} \sum_{g \in S} \chi(g,h)|g * y\rangle. \tag{5}$$

The space $\mathbb{C}^{\mathcal{X}}$ admits two natural orthonormal bases: the standard basis $\{|x\rangle : x \in \mathcal{X}\}$, and, for any fixed $x \in \mathcal{X}$, the Fourier basis

$$|G^{(h)} * x\rangle = \frac{1}{\sqrt{|G|}} \sum_{g \in G} \chi(g,h)|g * x\rangle, \qquad h \in G.$$

These states are simultaneous eigenstates of the action operators. Indeed, if $U_k : |y\rangle \mapsto |k * y\rangle$, then

$$U_k|G^{(h)} * x\rangle = \chi(-k,h)|G^{(h)} * x\rangle.$$

The following procedure, which appears in [Zha25, Theorem 3.3, 3.5] and [Dol25], recovers h from $|G^{(h)} * x\rangle$ via phase kickback by implementing

$$|G^{(h)} * x\rangle|0\rangle \mapsto |G^{(h)} * x\rangle|h\rangle.$$

Algorithm 2 CompIndex

Require: A state $|G^{(h)} * x\rangle|0\rangle$ for some $h \in G$
Ensure: The state $|G^{(h)} * x\rangle|h\rangle$
1: Apply the quantum Fourier transform to the second register.
2: Apply $\sum_{k\in G} U_k \otimes |k\rangle\langle k|$.
3: Apply the inverse quantum Fourier transform to the second register.

To prove our main theorem 9, we need a simple Fourier-analytic fact showing that quadratic refinements have flat Fourier magnitude.

Lemma 1. *Let G be a finite abelian group, let $\chi : G \times G \to \mathbb{C}^\times$ be a bicharacter, and let $\kappa : G \to \mathbb{T}$ be a quadratic refinement of χ, i.e.,*

$$\kappa(g+h) = \kappa(g)\kappa(h)\chi(g,h) \qquad \text{for all } g,h \in G.$$

Then, for every $g \in G$, $\widehat{\kappa}(g) = \frac{1}{\sqrt{|G|}}\kappa(g)^{-1} \sum_{h\in G} \kappa(h)$. In particular, $|\widehat{\kappa}(g)|$ is independent of g.

Proof. From the defining identity of a quadratic refinement,

$$\chi(g,h)\kappa(h) = \kappa(g)^{-1}\kappa(g+h).$$

Hence

$$\widehat{\kappa}(g) = \frac{1}{\sqrt{|G|}}\sum_{h\in G}\chi(g,h)\kappa(h) = \frac{1}{\sqrt{|G|}}\kappa(g)^{-1}\sum_{h\in G}\kappa(g+h).$$

Since translation by g permutes G, this equals $\frac{1}{\sqrt{|G|}}\kappa(g)^{-1}\sum_{h\in G}\kappa(h)$, which proves the claim. □

We now apply this observation to obtain a uniform sampler. We denote the complex unit circle by $\mathbb{T}$.

Algorithm 3 OrientedSampling(x,κ)

1: **Input:** $x=(E,\iota)\in\mathcal{X}$, $\kappa : G=\mathrm{Cl}(\mathfrak{O})\to\mathbb{T}$
2: **Output:** Uniform $(E,\iota)\in\mathcal{X}$
3: Prepare $|x\rangle|0\rangle = \frac{1}{\sqrt{|G|}}\sum_{h\in G}|G^{(h)}*x\rangle|0\rangle$
4: Apply CompIndex to $|x\rangle|0\rangle$
5: Apply the phase oracle $|y\rangle|h\rangle\mapsto\kappa(h)|y\rangle|h\rangle$
6: Apply CompIndex^{-1} and discard the second register
7: Measure the first register and output the result

Theorem 9. *Let $(\mathrm{Cl}(\mathfrak{O}),\mathcal{X},*)$ be a regular group action on the set $\mathcal{X}$ of $\mathfrak{O}$-oriented supersingular elliptic curves over $\mathbb{F}_{p^2}$. Then algorithm 3 samples a uniformly secure random curve (E,ι) from $\mathcal{X}$.*

Proof. Let $G=\mathrm{Cl}(\mathfrak{O})$ and any $x=(E,\iota)\in\mathcal{X}$ is an oriented curve along with an orientation. Starting from $|x\rangle = \frac{1}{\sqrt{|G|}}\sum_{h\in G}|G^{(h)}*x\rangle|0\rangle$, and applying CompIndex to this state we obtain the state

$$\frac{1}{\sqrt{|G|}}\sum_{h\in G}|G^{(h)}*x\rangle|h\rangle.$$

Now, we can compute any efficiently computable function $\kappa : G\to\mathbb{T}$ into the phase to obtain the state

$$\frac{1}{\sqrt{|G|}}\sum_{h\in G}\kappa(h)|G^{(h)}*x\rangle|h\rangle.$$

We can uncompute the second register by applying CompIndex^{-1} to obtain

$$\frac{1}{\sqrt{|G|}}\sum_{h\in G}\kappa(h)|G^{(h)}*x\rangle = \frac{1}{|G|}\sum_{g\in G}\left(\sum_{h\in G}\chi(g,h)\kappa(h)\right)|g*x\rangle = \frac{1}{\sqrt{|G|}}\sum_{g\in G}\widehat{\kappa}(g)|g*x\rangle.$$

Thus it suffices to choose κ so that $|\widehat{\kappa}(g)|$ is constant in g. Write $\mathrm{Cl}(\mathfrak{O})\cong\mathbb{Z}_{N_1}\times\cdots\times\mathbb{Z}_{N_k}$. For each j, let $\chi_j(a,h)=\exp(2\pi iah/N_j)$ be the standard bicharacter on $\mathbb{Z}_{N_j}$, and define

$$\kappa_j(h) = \begin{cases} \exp\left(\dfrac{\pi ih^2}{N_j}\right), & N_j \text{ even},\\ \exp\left(\dfrac{2\pi it_jh^2}{N_j}\right), & N_j \text{ odd},\end{cases}$$

where $2t_j\equiv 1 \pmod{N_j}$ in the odd case. Then each κ_j is a well-defined quadratic refinement of χ_j, so $\kappa(h_1,\ldots,h_k)=\prod_{j=1}^k\kappa_j(h_j)$ is a quadratic refinement of χ on $G=$

$\mathrm{Cl}(\mathfrak{O})$. By Lemma 1, $|\widehat{\kappa}(g)|$ is independent of g, and hence measuring $\frac{1}{\sqrt{|G|}}\sum_{g\in G}\widehat{\kappa}(g)|g*x\rangle$ produces a uniformly random element of $\mathcal{X}$.
The security of the curve clearly hinges on the VECTORIZATION problem as by construction the algorithm outputs the curve plus an $\mathfrak{O}$-orientation. However, as is the case in algorithm 1, there is no path or ideal used to construct the curve and hence it is not vulnerable to a path-storing attack. Additionally, for any QPT adversary storing a set P of polynomially many $\mathfrak{O}$-oriented curves E with known $\mathrm{End}(E)$ the probability that the output curve lies in this set is $\frac{|P|}{|\mathrm{Cl}(\mathfrak{O})|}$ which is negligible for cryptographic size $\mathrm{Cl}(\mathfrak{O})$. □ □

Remark 4. With more work one can analyze the running time of algorithm 3, as the most-costly step is evaluating the group action in superposition. Using the CLAPOTI algorithm [PR23] the running time is provably polynomial-time by using 8-dimensional isogenies.

Acknowledgement.

The authors would like to thank Raphael Steiner for valuable feedback on Section 3. We also thank Alessandro Sferlazza, Lorenz Panny, and the research group at the Technical University of Munich for pointing out a previous error in the running-time analysis of Proposition 4.
David Jao and Maher Mamah were supported by the NSERC Alliance Consortia Quantum Grant "Accelerating the Transition to Quantum-Resistant Cryptography" (ALLRP 578463–2022) and by the NSERC Discovery Grants Program. Jake Doliskani was supported by the NSERC Discovery Accelerator Supplement and Discovery Grants programs.

References

[ABOEM17] Dorit Aharonov, Michael Ben-Or, Elad Eban, and Urmila Mahadev. Interactive proofs for quantum computations. 2017. arXiv:1704.04487.

[ALM15] Nalini Anantharaman and Etienne Le Masson. Quantum ergodicity on large regular graphs. *Duke Mathematical Journal*, 164(4):723–765, 2015. doi:10.1215/00127094-2881592.

[Arp24] Sarah Arpin. Adding level structure to supersingular elliptic curve isogeny graphs. *Journal de Théorie des Nombres de Bordeaux*, 36(2):405–443, 2024. doi:10.5802/jtnb.1283.

[Bas24] Andrea Basso. A post-quantum round-optimal oblivious prf from isogenies. In Claude Carlet, Kalikinkar Mandal, and Vincent Rijmen, editors, *Selected Areas in Cryptography – SAC 2023*, volume 14201 of *Lecture Notes in Computer Science*, pages 147–168, Cham, August 2024. Springer.

[BBC+25] Andrea Basso, Giacomo Borin, Wouter Castryck, Maria Corte-Real Santos, Riccardo Invernizzi, Antonin Leroux, Luciano Maino, Frederik Vercauteren, and Benjamin Wesolowski. PRISM: Simple and Compact Identification and Signatures from Large Prime Degree Isogenies. In *Public-Key Cryptography – PKC*, volume 15676 of *Lecture Notes in Computer Science*, pages 300–332. Springer, 2025.

[BBD+24] 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. *The Computer Journal*, 67(8):2702–2718, 2024.

[BCC+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*, volume 14005 of *Lecture Notes in Computer Science*, pages 405–437. Springer, 2023. doi:10.1007/978-3-031-30617-4_14.

[BCK17] Dominic W. Berry, Andrew M. Childs, and Robin Kothari. Hamiltonian simulation with nearly optimal dependence on all parameters. *SIAM Journal on Computing*, 46(1):243–280, 2017. arXiv:1405.0093, doi:10.1137/16M106203X.

[BDF21] Jonas Burdges and Luca De Feo. Delay encryption. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021*, volume 12696 of *Lecture Notes in Computer Science*, pages 302–326. Springer, 2021.

[BFLS20] Daniel J. Bernstein, Luca De Feo, Antonin Leroux, and Benjamin Smith. Faster computation of isogenies of large prime degree. In *Advances in Cryptology – ASIACRYPT 2020*, volume 12491 of *Lecture Notes in Computer Science*, pages 311–339. Springer, 2020. doi:10.1007/978-3-030-64837-4_11.

[BJS14] Jean-François Biasse, David Jao, and Anirudh Sankar. A quantum algorithm for computing isogenies between supersingular elliptic curves. In *Progress in Cryptology – INDOCRYPT 2014*, volume 8885 of *Lecture Notes in Computer Science*, pages 428–442. Springer, 2014. doi:10.1007/978-3-319-13039-2_25.

[BKV19] Ward Beullens, Thorsten Kleinjung, and Frederik Vercauteren. Csi-fish: Efficient isogeny based signatures through class group computations. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology – ASIACRYPT 2019*, volume 11921 of *Lecture Notes in Computer Science*, pages 227–247. Springer, 2019.

[BM11] Valentin Blomer and Philippe Michel. Sup-norms of eigenfunctions on arithmetic ellipsoids. *International Mathematics Research Notices*, 2014(16):4323–4356, 2011. doi:10.1093/imrn/rnt045.

[BM13] Valentin Blomer and Philippe Michel. Hybrid bounds for automorphic forms on ellipsoids over number fields. *Journal of the Institute of Mathematics of Jussieu*, 14(2):253–300, 2013. doi:10.1017/S1474748014000141.

[Brö09] Jeroen Bröker. Constructing supersingular elliptic curves. *Mathematics of Computation*, 78(265):377–390, 2009. doi:10.1090/S0025-5718-08-02160-5.

[Bum09] Daniel Bump. *Automorphic Forms and Representations*, volume 55 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 2009.

[CGL08] Denis X. Charles, Eyal Z. Goren, and Kristin E. Lauter. Families of Ramanujan graphs and quaternion algebras. In *CRM Proceedings and Lecture Notes*, volume 47, pages 53–80. American Mathematical Society, 2008.

[CGL09] Denis X. Charles, Eyal Z. Goren, and Kristin E. Lauter. Cryptographic hash functions from expander graphs. *Journal of Cryptology*, 22(1):93–113, 2009. doi:10.1007/s00145-007-9002-x.

[CLM+18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. Csidh: An efficient post-quantum commutative group action. In *Advances in Cryptology – ASIACRYPT 2018*, volume 11274 of *Lecture Notes in Computer Science*, pages 395–427. Springer, 2018. doi:10.1007/978-3-030-03332-3_15.

[DEF+25] Pierrick Dartois, Jonathan Komada Eriksen, Tako Boris Fouotsa, Arthur Herlédan Le Merdy, Riccardo Invernizzi, Damien Robert, Ryan Rueger, Frederik Vercauteren, and Benjamin Wesolowski. Pegasis: Practical effective class group action using 4-dimensional isogenies. In *Advances in Cryptology – CRYPTO 2025*, volume 16000 of *Lecture Notes in Computer Science*, pages 67–99. Springer, 2025. doi:10.1007/978-3-032-01855-7_3.

[Del80] Pierre Deligne. La conjecture de weil. II. *Publications Mathématiques de l'IHÉS*, 52:137–252, 1980.

[DFMPS19] Luca De Feo, Simon Masson, Christophe Petit, and Antonio Sanso. Verifiable delay functions from supersingular isogenies and pairings. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology – ASIACRYPT 2019*, volume 11921 of *Lecture Notes in Computer Science*, pages 248–277, Cham, December 2019. Springer.

[Dol23] Javad Doliskani. How to sample from the limiting distribution of a continuous-time quantum walk. *IEEE Transactions on Information Theory*, 69(11):7149–7159, 2023. doi:10.1109/TIT.2023.3287556.

[Dol25] Jake Doliskani. Public-key quantum money from standard assumptions (in the generic model). *Cryptology ePrint Archive*, 2025.

[DV13] Lassina Dembélé and John Voight. Explicit methods for hilbert modular forms. In *Elliptic Curves, Hilbert Modular Forms and Galois Deformations*, Advanced Courses in Mathematics – CRM Barcelona, pages 135–198. Birkhäuser/Springer, Basel, 2013.

[EHL+18] Kirsten Eisentraäger, Sean Hallgren, Kristin Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: Reductions and solutions. In *Advances in Cryptology – EUROCRYPT 2018*, volume 10822 of *Lecture Notes in Computer Science*, pages 329–368. Springer, Cham, 2018. doi:10.1007/978-3-319-78372-7_11.

[FK17] Joseph F. Fitzsimons and Elham Kashefi. Unconditionally verifiable blind quantum computation. *Physical Review A*, 96(1):012303, 2017. arXiv:1203.5217, doi:10.1103/PhysRevA.96.012303.

[GH93] Dorian Goldfeld and Jeffrey Hoffstein. On the number of fourier coefficients that determine a modular form. In *A Tribute to Emil Grosswald: Number Theory and Related Analysis*, volume 143 of *Contemporary Mathematics*, pages 385–393. American Mathematical Society, 1993.

[GS25] Elif Özbay Gürler and Patrick Struck. How (not) to build identity-based encryption from isogenies. In *Selected Areas in Cryptography (SAC 2025)*, 2025.

[HS10] Roman Holowinsky and Kannan Soundararajan. Mass equidistribution for Hecke eigenforms. *Annals of Mathematics*, 172(2):1517–1528, 2010.

[IK04] Henryk Iwaniec and Emmanuel Kowalski. *Analytic Number Theory*, volume 53 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2004.

[IS95] Henryk Iwaniec and Peter Sarnak. l^∞ norms of eigenfunctions of arithmetic surfaces. *Annals of Mathematics*, 141(2):301–320, 1995. doi:10.2307/2118522.

[Kan18] Daniel M. Kane. Quantum money from modular forms. arXiv preprint, 2018. arXiv:1809.05925.

[KLM07] Phillip Kaye, Raymond Laflamme, and Michele Mosca. *An Introduction to Quantum Computing*. Oxford University Press, 2007.

[KNS24] Ilya Khayutin, Paul D. Nelson, and Raphael S. Steiner. Theta functions, fourth moments of eigenforms and the sup-norm problem II. *Forum of Mathematics, Pi*, 12:e11, 2024. doi:10.1017/fmp.2024.9.

[Koh96] David R. Kohel. *Endomorphism Rings of Elliptic Curves over Finite Fields*. PhD thesis, University of California, Berkeley, 1996.

[KSS22] Daniel M. Kane, Shahed Sharif, and Alice Silverberg. Quantum money from quaternion algebras. *Transactions on Mathematical Cryptology*, 2(1):60–83, 2022.

[Kup05] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM Journal on Computing*, 35(1):170–188, 2005. doi:10.1137/S0097539703436345.

[LGDdSG21] Yifan Lai, Steven D. Galbraith, and Caroline Delpech de Saint Guilhem. Compact, efficient and uc-secure isogeny-based oblivious transfer. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021*, volume 12696 of *Lecture Notes in Computer Science*, pages 213–241. Springer, 2021.

[Lin06] Elon Lindenstrauss. Invariant measures and arithmetic quantum unique ergodicity. *Annals of Mathematics*, 163(1):165–219, 2006. doi:10.4007/annals.2006.163.165.

[LP25] Shai Levin and Robi Pedersen. Faster proofs and vrfs from isogenies. In *Advances in Cryptology – ASIACRYPT 2025, Part IV*, Lecture Notes in Computer Science, pages 307–340, Singapore, December 2025. Springer.

[LWY05] Jianya Liu, Yonghui Wang, and Yangbo Ye. A proof of selberg's orthogonality for automorphic L-functions. *Manuscripta Mathematica*, 118:135–149, 2005. doi:10.1007/s00229-005-0561-8.

[Mam24] Maher Mamah. The supersingular ℓ-isogeny path and endomorphism ring problems: Tighter unconditional reductions. IACR Cryptology ePrint Archive, 2024. Report 2024/1569.

[Marar] Kimball Martin. The basis problem revisited. *Transactions of the American Mathematical Society*, to appear. Preprint available at arXiv:1804.04234.

[MJ23] Yassine Mokrani and David Jao. Generating supersingular elliptic curves over $\mathbb{F}_p$ with unknown endomorphism ring. In *Progress in Cryptology – INDOCRYPT 2023*, volume 14404 of *Lecture Notes in Computer Science*, pages 179–200. Springer, 2023. doi:10.1007/978-3-031-48228-1_9.

[MMP24] Marzio Mula, Nadir Murru, and Federico Pintore. On random sampling of supersingular elliptic curves. *Annali di Matematica Pura ed Applicata*, 204(3):1293–1335, 2024. doi:10.1007/s10231-024-01528-x.

[MS25] Hart Montgomery and Shahed Sharif. Quantum money from class group actions on elliptic curves. In Kang Min Chung and Yu Sasaki, editors, *Advances in Cryptology – ASIACRYPT 2024*, volume 15492 of *Lecture Notes in Computer Science*, pages 33–64. Springer, Singapore, 2025. doi:10.1007/978-981-96-0947-5_2.

[Onu21] Hiroshi Onuki. On oriented supersingular elliptic curves. *Finite Fields and Their Applications*, 69:101777, 2021. URL: https://www.sciencedirect.com/science/article/pii/S1071579720301465, doi:10.1016/j.ffa.2020.101777.

[Piz98] Arnold K. Pizer. Ramanujan graphs. In *The Legacy of Srinivasa Ramanujan*, volume 7 of *IAS/IP Studies in Advanced Mathematics*, pages 159–178. American Mathematical Society, 1998.

[PR23] Aurel Page and Damien Robert. Introducing clapoti(s): Evaluating the isogeny class group action in polynomial time. Cryptology ePrint Archive, 2023. Report 2023.

[PW24] Aurel Page and Benjamin Wesolowski. The supersingular endomorphism ring and one endomorphism problems are equivalent. *arXiv preprint*, page arXiv:2309.10432, 2024. Available at `arXiv:2309.10432`.

[Qu07] Yan Qu. The prime number theorem for automorphic L-functions for GL_m. *Journal of Number Theory*, 2007.

[RS94] Zeév Rudnick and Peter Sarnak. The behaviour of eigenstates of arithmetic hyperbolic manifolds. *Communications in Mathematical Physics*, 161(1):195–213, 1994. doi:10.1007/BF02099418.

[Sar93] Peter Sarnak. Arithmetic quantum chaos, 1993. Blyth Lectures, University of Toronto (May 1993). URL: https://web.math.princeton.edu/sarnak/Arithmetic%20Quantum%20Chaos.pdf.

[Sha25] Shahed Sharif. QuantumMoneyCode. https://github.com/ssharif/QuantumMoneyCode, 2025. Sage code for generating epsilon-separation data for quantum money serial numbers.

[Sil09] Joseph H. Silverman. *The Arithmetic of Elliptic Curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, 2 edition, 2009. doi:10.1007/978-0-387-09494-6.

[Sou10] Kannan Soundararajan. Quantum unique ergodicity for $\mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}$. *Annals of Mathematics*, 172(2):1529–1538, 2010.

[Ste22] Bruno Sterner. Commitment schemes from supersingular elliptic curve isogeny graphs. *Transactions on Mathematical Cryptology*, 1(2):40–51, 2022.

[Vél71] Jacques Vélu. Isogénies entre courbes elliptiques. *Comptes Rendus de l'Académie des Sciences de Paris*, 273:238–241, 1971.

[Voi21] John Voight. *Quaternion Algebras*, volume 288 of *Graduate Texts in Mathematics*. Springer, 2021. doi:10.1007/978-3-030-56694-4.

[Wes22] Benjamin Wesolowski. The supersingular isogeny path and endomorphism ring problems are equivalent. In *Proceedings of the 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, 2022.

[Zha25] Mark Zhandry. Quantum money from abelian group actions. *TheoretiCS*, 4, 2025.

A Appendix

A.1 Proof of Theorem 7

Here we give the proof of Theorem 7, which states that for any supersingular curve $E'/\mathbb{F}_{p^2}$, the output E of Algorithm 1 satisfies

$$\Pr[E = E'] \ll_\varepsilon p^{-1+\varepsilon} \qquad \text{for any } \varepsilon > 0.$$

Proof. By the separation assumption, the joint QPE tag $\tilde{\mathbf{t}}_i = (\tilde{\lambda}_{i,1}, \ldots, \tilde{\lambda}_{i,r})$ uniquely identifies the common eigenvector $|\phi_i\rangle$.

Let $|E'\rangle$ be any vertex basis state. Measuring only the vertex register sequentially yields

$$\Pr[E = E'] = \sum_{i=1}^{n} |\alpha_i|^2 \left|\langle E'|\phi_i\rangle\right|^2, \qquad \alpha_i = \langle E_0|\phi_i\rangle. \tag{6}$$

In particular, (6) implies $\Pr[E = E'] > 0$ for every vertex E'. Indeed, write $\rho := \mathrm{Trd}_{\mathrm{tags}}(|\Psi\rangle\langle\Psi|) = \sum_i |\alpha_i|^2 |\phi_i\rangle\langle\phi_i|$, so that $\Pr[E = E'] = \langle E'|\rho|E'\rangle$. Let

$$K := \mathrm{span}\{A_\ell^k |E_0\rangle : k \geq 0\}$$

be the Krylov subspace generated by A_ℓ for some prime $\ell \nmid p$. By spectral decomposition, $K = \mathrm{span}\{|\phi_i\rangle : \alpha_i \neq 0\}$ and hence $\ker(\rho) = K^\perp$. If $\langle E'|\rho|E'\rangle = 0$, then $|E'\rangle \in \ker(\rho) = K^\perp$, so $\langle E'|A_\ell^k|E_0\rangle = 0$ for all $k \geq 0$. But note that

$$\langle E'|A_\ell^k|E_0\rangle = \#\{\text{paths in } \mathcal{G}(p,\ell) \text{ from } E_0 \text{ to } E' \text{of length } k\},$$

and since $\mathcal{G}(p,\ell)$ is connected and A_ℓ has nonnegative weights, there exists k such that $(A_\ell^k)_{E',E_0} = \langle E'|A_\ell^k|E_0\rangle > 0$, a contradiction.
If we apply theorem 3 in equation (6) we obtain $\Pr[E = E'] \ll_\varepsilon p^{-1/2+\varepsilon}$ (since $\sum |\alpha_i|^2 = 1$); however we can do much better.
Recall that $m := |\mathcal{S}_p| \asymp p$, and define

$$r_E := \left(|\phi_1(E)|^2, \ldots, |\phi_m(E)|^2\right) \in \mathbb{R}^n, \qquad u := \left(\tfrac{1}{m}, \ldots, \tfrac{1}{m}\right).$$

Then (6) may be written as

$$r_{E_0} \cdot r_{E'} = \big(u + (r_{E_0} - u)\big) \cdot \big(u + (r_{E'} - u)\big),$$

and the cross terms cancel, yielding

$$\Pr[E = E'] = \|u\|_2^2 + (r_{E_0} - u) \cdot (r_{E'} - u).$$

By Cauchy–Schwarz inequality,

$$|(r_{E_0} - u) \cdot (r_{E'} - u)| \leq \|r_{E_0} - u\|_2 \, \|r_{E'} - u\|_2,$$

and Theorem 4 (fourth-moment bound), which states

$$\sup_{x,y} \sum_{i=1}^{m} (\phi_i^2(x) - \phi^2(y))^2 \ll_\varepsilon p^{-1+\varepsilon},$$

implies $\|r_E - u\|_2^2 \ll_\varepsilon p^{-1+\varepsilon}$ uniformly in E. Since $\|u\|_2^2 = m^{-1}$, we obtain

$$\Pr[E = E'] = m^{-1} + O_\varepsilon(p^{-1+\varepsilon}) \ll_\varepsilon p^{-1+\varepsilon}.$$

□

Remark 5. As indicated by conjecture 1, the exponent ε could be plausibly replaced by a logarithmic factor (e.g. $\log^2 p$). At present, however, unconditional results only achieve this bound up to a power loss.