Elliptic Extraction and Pushable Hints for Higher-Dimensional SQIsign

Anis Bkakria[1]

IRT SystemX, France

Abstract. Higher-dimensional SQIsign variants use product and Kani-style isogeny representations to make response isogenies compact, checkable, or compatible with non-smooth degrees. These representations raise a foundations question: do they introduce genuinely higher-dimensional hardness, or do they preserve elliptic extraction while moving the additional assumptions into simulation? This paper gives a proof-theoretic separation between these two layers. We identify a recoverability condition under which an accepted higher-dimensional response publicly induces an elliptic Hom element; under the usual challenge-separation hypotheses, special soundness then extracts an elliptic OneEnd witness. The remaining higher-dimensional assumptions are simulation assumptions, and we formulate them as non-interactive hint distributions.

The separation yields two concrete theorem packages. First, SQIsign2D–West admits an exact Fiat–Shamir-with-hints EUF-CMA theorem: the security loss is expressed through a simulation-hint OneEnd assumption, the statistical simulator loss, the challenge-space term, and the commitment min-entropy term; a pushable uniform-hint corollary gives the EndRing-compatible interpretation under a named hint-indistinguishability assumption. Second, for SQIPrime2D we do not claim a proof for the published auxiliary sampler. We define and analyze SQIPrime2D-UAux, a foundations variant that replaces the published Algorithm-context AIO layer by a context-independent uniform auxiliary hint distribution. UAux is an ideal proof-model distribution unless an efficient represented sampler is separately supplied; an approximate sampler instantiates the theorems with an explicit additive sampling loss.

For this SQIPrime2D-UAux foundations variant, UAux gives a context-independent and exactly pushable replacement for the auxiliary AIO layer. We then introduce challenge-carrying uniform hints that additionally reveal the non-smooth challenge isogeny, the residual response isogeny, and the UAux auxiliary isogeny. These hints are stronger than ordinary SQIPrime auxiliary information, but they are fully pushable through the smooth self-reduction isogenies used in EndRing-style reductions. Consequently SQIPrime2D-UAux admits a fully pushable EndRing-compatible challenge-carrying hint layer, with a q-HintDist assumption separating simulation-shaped hints from the pushable distribution. None of these SQIPrime-side statements is a claim of plain EndRing security or a claim that the published SQIPrime2D AIO sampler itself is pushable.

1 Introduction

SQIsign-style signatures derive their compactness from the Deuring correspondence between supersingular elliptic curves and quaternion orders. Their core knowledge claim is elliptic: the prover demonstrates knowledge associated with the endomorphism ring of a public supersingular elliptic curve, and the special-soundness extractor obtains a non-scalar endomorphism, or equivalently an instance of the elliptic OneEnd problem that is compatible with EndRing under the appropriate reductions [De +20; Ård+25]. At the same time, modern SQIsign-family constructions increasingly rely on higher-dimensional representations. SQIsignHD, SQIsign2D–West, SQIsign2D–East, SQIPrime, and recent ideal-to-isogeny algorithms use products of elliptic curves, Kani-style decompositions, level structures, and higher-dimensional evaluation data to represent isogenies whose degrees are awkward for direct elliptic chains [Dar+24; Bas+24; NO24; DF24; ON24; Bor+25; Bor+26].

This higher-dimensional turn creates a foundations problem. When verification accepts a response represented by an isogeny between products or higher-dimensional abelian varieties, it is not immediate whether the proof is still extracting an elliptic endomorphism, or whether the protocol has introduced a genuinely higher-dimensional hard relation. It is also not immediate how to compare oracle-assisted simulations used by different variants: some proofs rely on uniform-isogeny oracles, some on constrained isogeny oracles, and others on auxiliary isogenies generated by

protocol-specific Kani constructions. Treating all of these objects as a single "higher-dimensional assumption" hides the structure of the proof.

This paper separates the proof into two layers:

$$\text{elliptic extraction} \qquad \text{and} \qquad \text{higher-dimensional simulation.}$$

The extraction layer concerns what a special-soundness extractor obtains from two accepting transcripts with the same commitment and different challenges. The simulation layer concerns how accepting transcripts are sampled without the secret key. The main structural message is that higher-dimensional representation is not, by itself, higher-dimensional hardness. Whenever an accepted response publicly recovers an elliptic Hom element, the extractor remains elliptic. The additional assumptions introduced by higher-dimensional techniques appear in the simulation layer, where one must sample response representations of the correct distribution.

1.1 Central Question

The paper addresses the following question.

> For SQIsign-family protocols using higher-dimensional response representations, when does special soundness still extract an elliptic witness, and what additional assumptions are needed to make simulation and Fiat–Shamir security statements non-interactive and EndRing-compatible?

The answer is a recoverability-and-hints principle. If verification certifies a higher-dimensional morphism

$$\Phi_D : A \to B$$

together with public maps

$$\iota : E \to A, \qquad \pi : B \to F,$$

and the induced composite

$$\sigma_D = \pi \circ \Phi_D \circ \iota$$

is a nonzero efficiently evaluable elliptic map, then the accepted response datum D is elliptically recoverable. In SQI-style protocols, two accepting recoverable transcripts yield recovered elliptic response maps; composing them with the corresponding challenge maps produces an endomorphism of the public-key elliptic curve. The scheme-specific challenge-separation argument then excludes scalar endomorphisms, and the extractor outputs a OneEnd witness.

The simulation layer is different. The simulator must sample response representations, auxiliary isogenies, or constrained isogenies without knowing the secret key. Existing analyses express this through interactive oracles or equivalent auxiliary information. We use the Fiat–Shamir-with-hints viewpoint of the complete SQIsign proof: simulation-shaped hints $\mathcal{H}_{\mathsf{sim}}$ explain how to simulate transcripts, while pushable uniform hints $\mathcal{H}_{\mathsf{unif}}$ are the distributions that can be transported through self-reductions and therefore support EndRing-compatible interpretations [Ård+25]. A named $\mathsf{HintDist}$ assumption relates the simulation-shaped distribution to the pushable distribution when such a replacement is needed.

1.2 Scope of the SQIPrime2D-UAux Contribution

The SQIPrime-side contribution is intentionally a foundations-variant statement. Published SQIPrime2D uses an auxiliary-isogeny distribution generated by its response algorithm; this distribution is tied to Algorithm-context data and is modeled through AIO. This paper does not prove that the published AIO sampler is curve-local, uniform, or pushable. Instead, it separates two statements that should not be conflated:

$$\text{published SQIPrime2D AIO} \qquad \neq \qquad \mathsf{SQIPrime2D\text{-}UAux}.$$

The UAux variant keeps the SQIPrime2D response and dimension-two verification shape but replaces only the auxiliary AIO layer by a curve-local uniform auxiliary hint distribution. The

stronger challenge-carrying theorem then equips this foundations variant with a fully pushable hint distribution by giving the hinted hardness game explicit representations of the q-challenge isogeny, the residual response isogeny, and the UAux auxiliary isogeny. Thus the SQIPrime-side theorem is a hinted foundations theorem for SQIPrime2D-UAux, not a proof that the published SQIPrime2D auxiliary sampler itself has a fully pushable EndRing-compatible hint layer.

1.3 Contributions

The paper makes four contributions.

Recoverable higher-dimensional responses and elliptic extraction. We define HD-recoverability as an interface for accepted higher-dimensional response data that publicly induces an elliptic Hom element. The definition covers product-component and Kani-style response representations: the verifier checks higher-dimensional data, but the extractor uses the recovered elliptic component. We prove that HD-recoverability transfers special soundness to the elliptic setting. Under dual closure and challenge separation, two accepting transcripts with the same commitment and different challenges yield a non-scalar endomorphism of the public-key elliptic curve. The extracted object is therefore OneEnd; EndRing enters only through the applicable OneEnd-to-EndRing reduction and a compatible hint distribution.

Exact hint-based security for SQIsign2D–West. We instantiate the framework on SQIsign2D–West. Its accepted two-dimensional responses are recoverable as elliptic response maps, and its special-soundness proof extracts an elliptic OneEnd witness under the parameter condition governing challenge and response degrees [Bas+24]. We package the UTO/FIDIO-style simulator outputs into a simulation-shaped hint distribution and obtain an exact ROM EUF-CMA theorem for the Fiat–Shamir transform. The bound contains the $\mathcal{H}_{\mathsf{sim}}$-HintOneEnd advantage, the SQIsign2D–West hint-HVZK loss, the challenge-space term, and the commitment min-entropy term. We then state the EndRing-compatible corollary under a pushable 2D-West uniform-hint distribution and a corresponding HintDist assumption.

Context-independent auxiliary hints for a SQIPrime2D foundations variant. The published SQIPrime2D proof uses an auxiliary-isogeny oracle whose distribution is tied to the protocol's Algorithm-10 Kani/endomorphism construction and therefore to hidden context such as the base curve, the secret path, and the sampled endomorphism [DF24]. We isolate this as an AIO boundary: pushability of that auxiliary layer requires a context-independence property that is not part of the published model. We then define SQIPrime2D-UAux, a foundations variant which replaces the protocol-shaped auxiliary sampler by

$$\delta \leftarrow \mathsf{UAux}(E, 2^\alpha - q \cdot d),$$

the uniform distribution on admissible cyclic auxiliary isogenies of the required degree from the visible curve. We prove that UAux is context-independent and exactly pushable through isogenies of degree prime to the auxiliary degree. UAux therefore gives a context-independent, exactly pushable replacement for SQIPrime2D's auxiliary AIO layer, but it is a proof-model distribution unless an efficient represented sampler is separately instantiated.

Fully pushable challenge-carrying hints for SQIPrime2D-UAux. UAux alone fixes the auxiliary layer, but it does not make the full SQIPrime2D response hint EndRing-compatible: a hint of the form (L, κ, δ), where $L \subset E[q]$ is the challenge line and $L \subseteq \ker(\kappa)$, cannot be pushed through the q-challenge quotient by a bijection on q-torsion. We therefore introduce a stronger challenge-carrying uniform hint distribution. On a levelled public curve (E, P, Q), the hint samples a scalar a, the explicit q-isogeny

$$\phi_a : E \to F_a, \qquad \ker(\phi_a) = \langle P + [a]Q \rangle,$$

a residual response isogeny $\sigma : F_a \to E_1$ of degree d, and an independent UAux isogeny $\delta : E \to E_\delta$ of degree $2^\alpha - q \cdot d$. The transcript response is reconstructed as $\kappa = \sigma \circ \phi_a$. We prove that this challenge-carrying distribution is fully pushable through the smooth self-reduction isogenies whose

degrees are prime to $q \cdot d \cdot (2^\alpha - q \cdot d)$. Consequently SQIPrime2D-UAux admits a fully pushable EndRing-compatible challenge-carrying hint layer, with a corresponding q-HintDist assumption relating simulation-shaped hints to the pushable distribution. This hinted hardness assumption is stronger than ordinary SQIPrime assumptions, because the adversary receives explicit challenge-isogeny, residual-response, and UAux auxiliary information.

1.4 Resulting Perspective

The resulting picture is summarized by the following separation.

$$\begin{aligned} \text{higher-dimensional response representation} &\not\Rightarrow \text{higher-dimensional extraction hardness}, \\ \text{elliptic recoverability} &\Rightarrow \text{elliptic OneEnd extraction}, \\ \text{pushable hint distribution} &\Rightarrow \text{EndRing-compatible hint assumptions}. \end{aligned}$$

SQIsign2D–West gives the clean positive instantiation for exact hint-based Fiat–Shamir security. SQIPrime2D shows why an auxiliary layer can be context-shaped and hence unsuitable for pushable reductions. SQIPrime2D-UAux repairs that auxiliary layer, and challenge-carrying hints repair the remaining non-smooth challenge transport issue by carrying the q-isogeny itself inside the pushable distribution.

2 Technical Overview

This section describes the proof architecture and the role of the two concrete theorem packages. It is organized around one structural principle and two instantiations: recoverable higher-dimensional responses extract elliptic witnesses; SQIsign2D–West gives an exact Fiat–Shamir-with-hints theorem; and SQIPrime2D-UAux, a foundations variant rather than the published SQIPrime2D protocol, gives a fully pushable challenge-carrying hint layer after replacing the published auxiliary AIO distribution.

2.1 Representation Dimension versus Hardness Dimension

The first abstraction is HD-recoverability. An accepted response datum D may be represented by a morphism

$$\Phi_D : A \to B$$

between products of elliptic curves or other higher-dimensional objects. The response is nevertheless elliptic for extraction if the verifier's accepted data also gives public maps

$$\iota : E \to A, \qquad \pi : B \to F$$

such that the composite

$$\sigma_D = \pi \circ \Phi_D \circ \iota$$

is a nonzero efficiently evaluable element of $\mathrm{Hom}(E, F)$. This condition is deliberately representation-agnostic. It covers product projections, Kani matrices, and level-structure variants whenever the accepted transcript exposes a public elliptic component.

The extraction theorem then has a simple form. Suppose two accepting transcripts share the same commitment and have distinct challenges. Recover their elliptic response maps, compose them with the two challenge maps in the usual SQI-style cycle, and obtain an endomorphism of the public-key elliptic curve. If the scheme's challenge-separation condition excludes scalar cycles, this endomorphism is non-scalar. The output is an elliptic OneEnd witness. No higher-dimensional endomorphism-ring problem is used in this step.

This abstraction does not trivialize higher-dimensional protocols. It identifies where the higher-dimensional machinery matters: in representation and simulation. The verifier may need Kani data, torsion images, product maps, or auxiliary isogenies to check that Φ_D is valid and that the recovered elliptic component has the required degree and kernel properties. The simulator must sample such data with the right distribution. These are not extraction assumptions; they are distributional assumptions.

2.2 Simulation Hints and Pushable Hints

The second abstraction is the hint layer. A simulation-shaped hint distribution $\mathcal{H}_{\mathsf{sim}}$ packages exactly the auxiliary information that an honest-verifier simulator would otherwise request from interactive isogeny-sampling oracles. If the original simulator calls stateless simulation oracles and receives polynomial-size representations, the oracle outputs can be sampled first and supplied as one non-interactive hint. The simulator then executes the same deterministic transcript-construction steps, and the transcript distribution is unchanged. This gives an exact oracle-to-hint normal form.

However, $\mathcal{H}_{\mathsf{sim}}$ is usually transcript-shaped. It may depend on the sampled challenge, on the simulator's particular path through the protocol, or on auxiliary choices that do not transport cleanly across curves. EndRing-compatible reductions require a different property: hints must be pushable. A pushable uniform hint distribution $\mathcal{H}_{\mathsf{unif}}$ is one for which, whenever a self-reduction transports an instance through an isogeny $\rho : E \to E'$, the pushed hint has the same distribution as a fresh hint for E'. This is the role played by pushable hints in the complete SQIsign proof [Ård+25]. Simulation uses $\mathcal{H}_{\mathsf{sim}}$; EndRing compatibility uses $\mathcal{H}_{\mathsf{unif}}$; a HintDist assumption bridges the two distributions.

The paper applies this philosophy in two different ways. SQIsign2D–West already has a clean two-dimensional simulation story, so the main work is to package the UTO/FIDIO outputs into $\mathcal{H}_{\mathsf{sim}}$, instantiate Fiat–Shamir with hints, and then isolate the additional HintDist assumption needed to move to a pushable uniform distribution. SQIPrime2D is more delicate: its published auxiliary isogeny is generated by a Kani/endomorphism algorithm tied to hidden context. UAux repairs this auxiliary layer, while challenge-carrying hints are needed to make the entire SQIPrime2D-UAux hint distribution pushable.

2.3 SQIsign2D–West as the Clean Fiat–Shamir Instantiation

SQIsign2D–West is the principal positive instantiation. Its responses are represented in dimension two, but accepted responses recover elliptic isogenies between the commitment and challenge curves. The special-soundness proof uses two accepting responses to construct a non-scalar endomorphism of the public-key curve; the relevant scalar-exclusion argument is a degree and cyclic-factor comparison under the parameter condition controlling challenge and response exponents [Bas+24].

For simulation, SQIsign2D–West uses oracle-assisted sampling: the simulator samples challenge data and obtains the corresponding response and auxiliary representations through isogeny-sampling oracles. The hint normal form packages these oracle answers into $\mathcal{H}_{\mathsf{sim}}^{2D\text{-}W}$. The Fiat–Shamir-with-hints theorem then gives a ROM EUF-CMA bound of the form

$$\mathrm{Adv}^{\mathsf{EUF\text{-}CMA}} \leq \mathrm{Adv}^{\mathsf{HintOneEnd}^{\mathcal{H}_{\mathsf{sim}}^{2D\text{-}W}}} + \varepsilon_{\mathsf{HVZK}} + \varepsilon_{\mathsf{FS}} + \varepsilon_{\mathsf{minent}}.$$

Here $\varepsilon_{\mathsf{HVZK}}$ is the statistical distance inherited from the SQIsign2D–West simulator, $\varepsilon_{\mathsf{FS}}$ contains the challenge-space and random-oracle terms, and $\varepsilon_{\mathsf{minent}}$ is the commitment min-entropy term. The exact theorem is stated at the hint-OneEnd layer because this is the distribution the simulator actually uses.

The EndRing-compatible form is obtained only after replacing the simulation-shaped hints by pushable uniform hints. The paper therefore states a separate corollary under a 2D-West HintDist assumption. This separation prevents an overclaim: the theorem does not assert plain EndRing security without hints, and it does not identify $\mathcal{H}_{\mathsf{sim}}$ with a pushable distribution unless the explicit indistinguishability assumption is made.

2.4 The SQIPrime2D AIO Boundary

SQIPrime changes the structure of the response by using non-smooth challenge isogenies. In the two-dimensional variant, the verifier checks a response isogeny

$$\kappa : E_A \to E_1$$

of degree $q \cdot d$ whose kernel contains the verifier's q-challenge line, together with auxiliary data enabling a dimension-two Kani verification. The auxiliary isogeny has degree

$$2^\alpha - q \cdot d,$$

so that the response and auxiliary degrees add to 2^α. Published SQIPrime2D obtains this auxiliary isogeny through a protocol-shaped Kani/endomorphism construction rather than through a context-independent uniform sampler [DF24].

This creates a pushability boundary. A public pushability theorem for the published auxiliary oracle would require the distribution of the auxiliary isogeny to depend only on the visible curve and the auxiliary degree. The published AIO distribution is instead defined by matching the auxiliary isogeny produced inside the response algorithm, which is tied to hidden context. The paper therefore does not claim that the published AIO is non-pushable; it proves the sharper statement that pushability requires a context-independence property absent from the published oracle model.

UAux is the corresponding proof-model repair. It is first a distribution used by the hinted experiment; an efficient sampler is a separate instantiation issue accounted for by an explicit sampler-distance loss. For a visible curve E and admissible auxiliary degree n, $\mathsf{UAux}(E, n)$ samples uniformly from the admissible cyclic degree-n auxiliary isogenies from E. Because the sample space is defined only from (E, n), the distribution is context-independent. Because an isogeny $\rho : E \to E'$ of degree prime to n induces a bijection on n-torsion, pushforward maps uniform auxiliary kernels from E to uniform auxiliary kernels from E'. Hence

$$\rho_* \mathsf{UAux}(E, n) = \mathsf{UAux}(E', n).$$

This is an exact distributional statement, not a computational assumption.

2.5 Why Challenge-Carrying Hints Are Needed

UAux solves only the auxiliary layer. It does not make a line-constrained SQIPrime2D hint fully pushable. A natural hint would contain

$$(L, \kappa, \delta), \qquad L \subset E[q], \qquad L \subseteq \ker(\kappa),$$

where L is the challenge line and δ is the UAux auxiliary isogeny. This form is not suitable for full EndRing-compatible transport through the q-challenge quotient: the quotient map kills L, and its restriction to q-torsion is not a bijection. Thus one cannot push such hints by the same kernel-bijection argument.

The fully pushable distribution carries the challenge isogeny itself. The public key is treated as a levelled curve (E, P, Q), and the hint samples

$$a \leftarrow \mathbb{Z}_q, \qquad \phi_a : E \to F_a, \qquad \ker(\phi_a) = \langle P + [a]Q \rangle,$$

a residual response isogeny

$$\sigma : F_a \to E_1 \qquad \text{of degree } d,$$

and an independent UAux auxiliary isogeny

$$\delta : E \to E_\delta \qquad \text{of degree } 2^\alpha - q \cdot d.$$

The transcript response is reconstructed as

$$\kappa = \sigma \circ \phi_a.$$

The hint is stronger than UAux-only because it contains an efficient representation of the non-smooth q-challenge isogeny, the residual response isogeny, and the UAux auxiliary isogeny. It therefore defines a stronger hinted-hardness assumption than ordinary SQIPrime or UAux-only. This stronger distribution is precisely what makes full pushability possible.

Let $\rho : (E, P, Q) \to (E', P', Q')$ be a smooth self-reduction isogeny of degree prime to $q \cdot d \cdot (2^\alpha - q \cdot d)$, with $P' = \rho(P)$ and $Q' = \rho(Q)$. Since ρ is invertible on q-torsion, the same scalar a determines the transported challenge line

$$\rho(\langle P + [a]Q \rangle) = \langle P' + [a]Q' \rangle.$$

The pushout of ϕ_a along ρ is therefore a uniformly distributed q-isogeny from E'. The residual response σ pushes through the induced isogeny on the codomain because the transport degree is

prime to d. The auxiliary δ pushes by UAux pushability because the transport degree is prime to $2^\alpha - q \cdot d$. The three component maps are distribution-preserving bijections, so the joint pushed hint is distributed exactly as a fresh challenge-carrying hint for (E', P', Q').

This proves the paper's SQIPrime-side flagship statement:

SQIPrime2D-UAux admits a fully pushable EndRing-compatible challenge-carrying hint layer.

The corresponding EUF-CMA theorem is necessarily conditional on a q-HintDist assumption between the simulation-shaped distribution and the challenge-carrying uniform distribution, and on the hint-EndRing hardness assumption for this stronger hint family. If UAux is instantiated by an approximate sampler rather than by the ideal-hint experiment, the theorem gains the sampler-loss term of Proposition 10. The theorem does not claim plain EndRing security without hints and does not claim security for the published SQIPrime2D AIO sampler.

2.6 Proof Roadmap

The proof development follows the order suggested by the separation above, and each step is tied to a later formal statement. First, Section 4 defines HD-recoverability and proves Theorem 1: an accepted recoverable higher-dimensional response exposes an efficiently evaluable elliptic Hom element. Section 5 then turns this recovered Hom element into an extraction theorem. The central result there is Theorem 2, which says that recoverability, dual closure, and a protocol-specific scalar-exclusion property yield an elliptic OneEnd witness.

Second, Section 6 supplies the proof infrastructure used by the security theorems. Theorem 3 converts oracle-assisted simulators into simulation-shaped hints; Theorem 4 gives the Fiat–Shamir-with-hints reduction template; Lemma 2, Proposition 6, and Theorem 5 separate simulation hints from pushable EndRing-compatible hints.

Third, Section 7 instantiates this infrastructure for SQIsign2D–West. Proposition 7, Theorem 6, and Lemmas 3 and 4 provide the recoverability, special-soundness, simulation, and min-entropy inputs. These inputs are assembled in Theorem 7, then translated to the pushable-hint layer in Corollary 4 and to commitment-recoverable signatures in Corollary 5.

Fourth, Section 8 isolates the SQIPrime2D auxiliary-isogeny obstruction. Proposition 8 records that SQIPrime accepted responses remain elliptic at extraction, while Proposition 9 and Corollary 6 show why the published Algorithm-context AIO cannot be treated as a curve-local pushable distribution without an additional distributional claim.

Fifth, Section 9 gives the auxiliary repair. Theorems 8 and 9 prove that UAux is context-independent and exactly pushable, and Theorem 10 and Proposition 11 show that it preserves the dimension-two SQIPrime verification degree relation while replacing the published AIO layer.

Finally, Section 10 adds the full pushable layer. Theorem 11 reconstructs SQIPrime2D-UAux transcripts from challenge-carrying hints, Theorem 12 proves exact pushability of the whole challenge-carrying distribution, Corollary 7 records the hinted EndRing consequence, and Theorem 13 gives the corresponding ROM EUF-CMA theorem under Q-HintDist.

3 Preliminaries and Representation Interfaces

This section fixes the algebraic notation and the representation interfaces used throughout the paper. The purpose of the interface language is to keep the extraction and pushability theorems independent of a particular coordinate system. A response may be represented by chains of product isogenies, theta-coordinate evaluators, Kani matrices, or ideal-to-isogeny output; the reductions use only validity, evaluation, composition, dualization when declared by the protocol, and pushforward operations for the distributions that are explicitly called pushable.

3.1 Elliptic Curves, Hom Spaces, and Isogenies

Unless stated otherwise, all curves are supersingular elliptic curves over $\mathbb{F}_{p^2}$, with $p > 3$, and all morphisms are defined over the indicated base field or over the extension field specified by their evaluator. For elliptic curves E and F, we write

$$\mathrm{Hom}(E, F)$$

for the abelian group of origin-preserving morphisms $E \to F$, and

$$\mathrm{End}(E) = \mathrm{Hom}(E, E)$$

for the endomorphism ring. Composition is written with $\circ$. The zero homomorphism is denoted by $0_{E,F}$, or by 0 when the source and target are clear. For $m \in \mathbb{Z}$, the multiplication-by-m endomorphism is denoted by $[m]_E$, or simply by $[m]$.

An isogeny is a nonzero element of $\mathrm{Hom}(E, F)$. The degree of an isogeny $\phi : E \to F$ is denoted by $\deg(\phi)$. Its dual isogeny is denoted by

$$\widehat{\phi} : F \to E,$$

and satisfies

$$\widehat{\phi} \circ \phi = [\deg(\phi)]_E, \qquad \phi \circ \widehat{\phi} = [\deg(\phi)]_F.$$

When n is prime to p, the n-torsion group scheme is etale and is denoted by $E[n]$. If $\rho : E \to F$ is an isogeny with $\gcd(n, \deg(\rho)) = 1$, then ρ restricts to an isomorphism

$$\rho|_{E[n]} : E[n] \xrightarrow{\sim} F[n].$$

Indeed, $\widehat{\rho} \circ \rho = [\deg(\rho)]$, and multiplication by $\deg(\rho)$ is invertible on $E[n]$. The inverse is

$$[\deg(\rho)]^{-1} \cdot \widehat{\rho}|_{F[n]}.$$

This elementary observation is the kernel-bijection behind the pushability theorems used later.

Proposition 1 (Nonzero elliptic homomorphisms are isogenies). *Let E and F be elliptic curves over a field k. If $\sigma \in \mathrm{Hom}(E, F)$ and $\sigma \neq 0$, then $\sigma : E \to F$ is an isogeny.*

Proof. The image $\sigma(E)$ is a closed irreducible subgroup variety of F. Since $\sigma \neq 0$, this image is not the identity point. The elliptic curve F is a one-dimensional smooth projective curve, so its irreducible closed subvarieties are points and the whole curve. Hence $\sigma(E) = F$, and σ is surjective.

The kernel $\ker(\sigma)$ is a closed subgroup scheme of E. If it had positive dimension, its identity component would be all of E, because E is connected of dimension one. Then σ would vanish on all of E, contradicting $\sigma \neq 0$. Therefore $\ker(\sigma)$ has dimension zero, and a zero-dimensional closed subgroup scheme of a projective curve is finite. Thus σ is a surjective homomorphism with finite kernel, which is precisely an isogeny. □

3.2 OneEnd, EndRing, and Hint Relations

The paper keeps two hardness layers separate. The extraction theorems output a single non-scalar endomorphism, which is a OneEnd witness. EndRing enters later, after a pushable hint distribution has been defined so that hinted OneEnd instances can be transported through the standard self-reductions used in EndRing-style security statements. This distinction prevents the extraction theorem from being read as a claim that a full endomorphism-ring basis is recovered from two transcripts.

The extraction theorems in this paper output elliptic OneEnd witnesses. EndRing-compatible statements are obtained only when the hint distribution used by the hardness problem is compatible with the standard self-reductions from EndRing to OneEnd.

Definition 1 (OneEnd and EndRing). *Let $\mathcal{D}_\lambda$ be a distribution on supersingular elliptic curves over $\mathbb{F}_{p^2}$. The* OneEnd *problem for $\mathcal{D}_\lambda$ asks, given $E \leftarrow \mathcal{D}_\lambda$, to output an efficient evaluator for an endomorphism*

$$\alpha \in \mathrm{End}(E) \setminus \mathbb{Z},$$

where $\mathbb{Z}$ is identified with the scalar endomorphisms $[m]_E$. The EndRing *problem asks, given $E \leftarrow \mathcal{D}_\lambda$, to output a $\mathbb{Z}$-basis of $\mathrm{End}(E)$, together with efficient evaluators for the basis elements.*

Definition 2 (Relation with hints). *Let $R \subseteq X \times W$ be a relation, and let $\mathcal{H} = \{\mathcal{H}_x\}_{x \in X}$ be a family of polynomial-size auxiliary distributions indexed by statements. The hinted relation $R^{\mathcal{H}}$ gives the adversary the statement x, together with polynomially many independent hints sampled from $\mathcal{H}_x$, and asks it to output a witness w such that $(x, w) \in R$. We write* $\mathsf{HintOneEnd}^{\mathcal{H}}$ *and* $\mathsf{HintEndRing}^{\mathcal{H}}$ *for the corresponding hinted versions of* OneEnd *and* EndRing.

This definition is intentionally semantic: a hint may contain isogeny representations, auxiliary-isogeny metadata, torsion images, or challenge-carrying data. Later sections define the concrete distributions used for SQIsign2D–West and SQIPrime2D-UAux. The important distinction is that a simulation-shaped hint distribution is designed to reproduce a simulator's transcript distribution, whereas a uniform pushable hint distribution is designed to be transported through self-reductions.

3.3 Efficient Evaluators

The formal statements use evaluator interfaces rather than committing to one representation technology. This convention is only a correctness interface: it records which morphism an accepted datum evaluates and which operations the verifier certifies. It does not hide a hardness assumption or assert that arbitrary higher-dimensional data admits the required operations.

The reductions use evaluators rather than a fixed representation language. This is necessary because the schemes considered in the paper use different representation technologies: elliptic chains, two-dimensional product isogenies, theta coordinates, level structures, and ideal-to-isogeny outputs.

Definition 3 (Efficient elliptic evaluator). *Let E and F be elliptic curves over $\mathbb{F}_{p^2}$. An efficient evaluator for $\sigma \in \mathrm{Hom}(E, F)$ is a string D_σ, together with deterministic polynomial-time algorithms*

$$\mathsf{Valid}_{\mathrm{Hom}}(D_\sigma), \qquad \mathsf{Eval}_{\mathrm{Hom}}(D_\sigma, P),$$

such that, when $\mathsf{Valid}_{\mathrm{Hom}}(D_\sigma) = (E, F)$, the evaluation algorithm outputs $\sigma(P)$ for every point $P \in E(K)$ and every extension field $K/\mathbb{F}_{p^2}$ whose degree is polynomial in $\log p$. Polynomial time is measured in $\log p$, the extension degree, and the bitlength of D_σ.

Definition 4 (Efficient higher-dimensional evaluator). *Let A and B be abelian varieties over $\mathbb{F}_{p^2}$. An efficient higher-dimensional evaluator for a morphism $\Phi : A \to B$ is a string D_Φ, together with deterministic polynomial-time algorithms*

$$\mathsf{Valid}_{\mathsf{HD}}(D_\Phi), \qquad \mathsf{Eval}_{\mathsf{HD}}(D_\Phi, X),$$

such that, when $\mathsf{Valid}_{\mathsf{HD}}(D_\Phi) = (A, B)$, the evaluation algorithm outputs $\Phi(X)$ for every point $X \in A(K)$ and every extension field $K/\mathbb{F}_{p^2}$ of polynomial degree.

Definition 5 (Public embedding-projection pair). *Let E and F be elliptic curves, and let A and B be abelian varieties. A public elliptic embedding-projection pair for (E, F, A, B) is a pair of efficiently evaluable morphisms*

$$\iota : E \to A, \qquad \pi : B \to F.$$

In a product representation one often has $A = E \times A'$, $B = F \times B'$, $\iota(P) = (P, 0)$, and $\pi(Q, Q') = Q$. The definitions do not require product structure; they require only that ι and π are public morphisms with efficient evaluation.

3.4 Representation Interfaces for Extraction and Pushability

The operations listed below are theorem hypotheses for the representation layer. In particular, a PushForward operation is not automatic for every way of encoding an isogeny. When a pushability theorem invokes this operation, it assumes that the concrete representation supplies sound pushout and transport metadata for the morphisms being moved.

The following interface records the operations that a protocol may expose. Not every scheme exposes every operation. When a theorem requires an operation, it states that requirement explicitly.

Definition 6 (Admissible representation interface). *An admissible representation interface for a family of response data consists of polynomial-time algorithms with the following semantics.*

(i) $\mathsf{DomCod}(D)$ *returns the source and target objects of the represented morphism.*
(ii) $\mathsf{Valid}(D)$ *verifies that* D *is a valid evaluator for the declared morphism.*
(iii) $\mathsf{Eval}(D, X)$ *evaluates the represented morphism on points over polynomial-degree extension fields.*
(iv) $\mathsf{Compose}(D_2, D_1)$*, when declared, outputs an evaluator for the composed morphism represented by* $D_2 \circ D_1$*.*
(v) $\mathsf{Dual}(D)$*, when declared, outputs an evaluator for the dual isogeny of the morphism represented by* D*.*
(vi) $\mathsf{Degree}(D)$ *and* $\mathsf{KernelCheck}(D, C)$*, when declared, certify degree and kernel-membership claims used by the verifier.*
(vii) $\mathsf{PushForward}(\rho, D)$*, when declared for an isogeny* ρ*, outputs a representation of the pushforward of the morphism represented by* D *along* ρ*, together with the induced codomain transport data.*

The interface is sound if accepted outputs represent the mathematical morphisms specified by the semantics above.

The extraction results use validity, evaluation, composition, and dualization. The challenge-carrying SQIPrime2D-UAux pushability theorem additionally uses pushforward and transport of level structures. The interface is deliberately stated as an abstraction because concrete schemes instantiate it differently: a chain of elliptic isogenies gives a direct evaluator; a product-isogeny representation gives a product evaluator plus public projections; a Kani representation gives matrix or theta data with verified component behavior.

Definition 7 (Levelled elliptic curve and smooth transport). *Let* $q \neq p$ *be prime. A* q*-levelled elliptic curve is a triple*

$$\mathbf{E} = (E, P, Q),$$

where (P, Q) *is a public ordered basis of* $E[q]$*. A smooth transport from* $\mathbf{E}$ *to* $\mathbf{E}' = (E', P', Q')$ *is an isogeny* $\rho : E \to E'$ *whose degree is prime to* q*, together with the condition*

$$P' = \rho(P), \qquad Q' = \rho(Q).$$

Since ρ *is invertible on* q*-torsion, the transported pair* (P', Q') *is again a basis of* $E'[q]$*.*

This levelled notation is used only for the SQIPrime2D-UAux challenge-carrying hint distribution. It makes explicit why a scalar challenge $a \in \mathbb{Z}/q\mathbb{Z}$ transports cleanly under smooth self-reductions:

$$\rho(\langle P + [a]Q \rangle) = \langle P' + [a]Q' \rangle.$$

3.5 Pushforward of Uniform Kernel Distributions

The later UAux theorem relies on the following general kernel-bijection fact. It is stated here so that the scheme-specific sections can refer to a fixed algebraic principle.

Lemma 1 (Kernel-bijection pushforward). *Let* $\rho : E \to F$ *be an isogeny of elliptic curves over* $\mathbb{F}_{p^2}$*, and let* n *be an integer with* $\gcd(n, p \cdot \deg(\rho)) = 1$*. Then the map*

$$K \longmapsto \rho(K)$$

is a bijection from cyclic subgroup schemes $K \subset E[n]$ *of order* n *to cyclic subgroup schemes* $K' \subset F[n]$ *of order* n*. Consequently, if* K *is uniform over the first set, then* $\rho(K)$ *is uniform over the second set.*

Proof. Because $\gcd(n, \deg(\rho)) = 1$, multiplication by $\deg(\rho)$ is an automorphism of $E[n]$ and $F[n]$. The dual identity gives

$$\widehat{\rho} \circ \rho = [\deg(\rho)]_E, \qquad \rho \circ \widehat{\rho} = [\deg(\rho)]_F.$$

Restricting to n-torsion and multiplying by the inverse of $\deg(\rho)$ modulo n shows that $\rho|_{E[n]}$ is an isomorphism with inverse

$$[\deg(\rho)]^{-1} \cdot \widehat{\rho}|_{F[n]}.$$

An isomorphism of finite group schemes maps cyclic order-n subgroup schemes to cyclic order-n subgroup schemes and has an inverse that maps cyclic order-n subgroup schemes back. Hence $K \mapsto \rho(K)$ is a bijection. A bijection maps the uniform distribution on a finite set to the uniform distribution on the image finite set, and the image set is the entire target set by surjectivity of the bijection. □

4 Recoverable Higher-Dimensional Responses

This section gives the structural abstraction underlying the paper. A response may be verified by a higher-dimensional computation, but it is elliptically recoverable when the accepted data publicly induces an elliptic Hom element. The abstraction is intentionally phrased as an interface: it does not assert that every higher-dimensional isogeny relation contains an elliptic component, and it does not replace the verification theorem of any concrete scheme. It isolates the condition under which higher-dimensional representation does not change the extraction target.

4.1 The HD-Recoverability Interface

The interface below is the first formal expression of the paper's dimension separation. A response can be encoded, checked, or evaluated in dimension two or higher, yet still expose an elliptic component after the verifier accepts. Recoverability is exactly the condition that the accepted public data determines an elliptic composite

$$\sigma_D = \pi \circ \Phi_D \circ \iota : E \to F.$$

When this condition holds, a special-soundness extractor need not solve a higher-dimensional endomorphism problem; it can work with the recovered elliptic map. The definition is deliberately silent about scalar exclusion and about simulation: those are separate proof obligations handled in Sections 5 and 6.

Definition 8 (HD-recoverable response). *Let E and F be elliptic curves over $\mathbb{F}_{p^2}$. An accepted response datum D is* $\mathsf{HDRecoverable}$ *for (E, F) if the accepted transcript publicly determines abelian varieties A and B, a valid efficient evaluator for a morphism*

$$\Phi_D : A \to B,$$

and public efficiently evaluable morphisms

$$\iota : E \to A, \qquad \pi : B \to F,$$

such that the verifier's acceptance predicate certifies the following statements.

(i) D is valid as a higher-dimensional evaluator for $\Phi_D : A \to B$.
(ii) ι and π are the public embedding and projection maps associated with the transcript.
(iii) The composite

$$\sigma_D = \pi \circ \Phi_D \circ \iota : E \to F$$

is not the zero homomorphism.

The recovered elliptic evaluator is

$$\mathsf{RecEval}_D(P) = \pi\big(\mathsf{Eval}_{\mathsf{HD}}(D, \iota(P))\big).$$

Scope of the definition. The definition is mathematical rather than syntactic. In a direct product representation, A and B may be products of elliptic curves and ι, π may be the standard inclusion and projection. In a Kani or theta-coordinate representation, Φ_D may be represented by matrix data or by an evaluator in a non-elliptic coordinate model; recoverability requires that the public verification data nevertheless identifies the elliptic composite $\pi \circ \Phi_D \circ \iota$. If no such public composite is available, the response is outside this definition.

4.2 Recoverability Gives Elliptic Hom Elements

The next theorem is intentionally elementary. Its role is to pin down the exact mathematical object exposed by an accepted recoverable response. It is not a theorem about how to generate responses, and it does not prove special soundness. It says that once the verifier has accepted the higher-dimensional datum and the public embedding-projection maps are fixed, the response has already determined an efficiently evaluable elliptic Hom element. Later extraction arguments use precisely that object and nothing stronger.

Theorem 1 (HD recoverability gives an elliptic Hom element). *Let E and F be elliptic curves over $\mathbb{F}_{p^2}$. Let D be an accepted* HDRecoverable *response datum for (E, F). Then*

$$\sigma_D = \pi \circ \Phi_D \circ \iota$$

belongs to $\mathrm{Hom}(E, F)$*. Moreover,* $\mathsf{RecEval}_D$ *is an efficient evaluator for σ_D. Since acceptance certifies $\sigma_D \neq 0$, the recovered homomorphism is an isogeny $E \to F$.*

Proof. By Definition 8, acceptance certifies that D is a valid higher-dimensional evaluator for a morphism $\Phi_D : A \to B$, and that $\iota : E \to A$ and $\pi : B \to F$ are public morphisms. Morphisms of abelian varieties are closed under composition. Therefore

$$\pi \circ \Phi_D \circ \iota : E \to F$$

is a morphism. Each of ι, Φ_D, and π preserves the identity element, so their composite also preserves the identity. Hence

$$\sigma_D = \pi \circ \Phi_D \circ \iota$$

is an element of $\mathrm{Hom}(E, F)$.

We now prove efficient evaluability. Let $K/\mathbb{F}_{p^2}$ be an extension field whose degree is polynomial in $\log p$, and let $P \in E(K)$. The public evaluator for ι computes

$$X = \iota(P) \in A(K)$$

in polynomial time. Since D is valid as a higher-dimensional evaluator, $\mathsf{Eval}_{\mathsf{HD}}(D, X)$ computes

$$Y = \Phi_D(X)$$

in polynomial time. The public evaluator for π computes

$$Q = \pi(Y) \in F(K)$$

in polynomial time. By definition of composition,

$$Q = \pi(\Phi_D(\iota(P))) = \sigma_D(P).$$

Thus $\mathsf{RecEval}_D$ evaluates σ_D on all polynomial-degree extension points in polynomial time.

Finally, acceptance certifies $\sigma_D \neq 0$. By Proposition 1, a nonzero homomorphism between elliptic curves is an isogeny. Therefore $\sigma_D : E \to F$ is an elliptic isogeny. □

Corollary 1 (Representation dimension is irrelevant to elliptic extraction). *Let D and D' be two accepted response data for the same elliptic pair (E, F). Suppose both are* HDRecoverable *and their recovered composites are equal in* $\mathrm{Hom}(E, F)$*:*

$$\pi \circ \Phi_D \circ \iota = \pi' \circ \Phi_{D'} \circ \iota'.$$

Then every extraction argument that uses only the recovered elliptic Hom element receives the same mathematical input from D and from D'.

Proof. Let

$$\sigma_D = \pi \circ \Phi_D \circ \iota, \qquad \sigma_{D'} = \pi' \circ \Phi_{D'} \circ \iota'.$$

The hypothesis states that $\sigma_D = \sigma_{D'}$ as elements of $\mathrm{Hom}(E, F)$. Equality in $\mathrm{Hom}(E, F)$ is equality as morphisms $E \to F$. Therefore, for every extension field $K/\mathbb{F}_{p^2}$ and every point $P \in E(K)$, one has

$$\sigma_D(P) = \sigma_{D'}(P).$$

An extraction argument that uses only the recovered elliptic Hom element can use the response datum only through the recovered map, its domain and codomain, and protocol-certified composition operations involving that map. Those mathematical objects are identical for D and D'. Hence the representation dimension, coordinate model, or higher-dimensional evaluator used to certify the response does not affect such an extraction argument. □

4.3 Product-Component Recovery

The simplest source of recoverability is a product evaluator whose elliptic component is obtained by public inclusion and projection.

Proposition 2 (Product-component recovery). *Let E_0, E_1, E_2, E_{12} be elliptic curves. Suppose an accepted datum D is a valid evaluator for a product morphism*

$$\Phi : E_0 \times E_{12} \to E_1 \times E_2$$

of the form

$$\Phi(P, R) = \big(\phi_1(P) + \psi_1(R),\ \phi_2(P) + \psi_2(R)\big),$$

where

$$\phi_1 \in \mathrm{Hom}(E_0, E_1), \quad \phi_2 \in \mathrm{Hom}(E_0, E_2), \quad \psi_1 \in \mathrm{Hom}(E_{12}, E_1), \quad \psi_2 \in \mathrm{Hom}(E_{12}, E_2).$$

Let $\iota_0 : E_0 \to E_0 \times E_{12}$ be $P \mapsto (P, 0)$, and let π_i be the product projection to the i-th target factor. Then

$$\pi_1 \circ \Phi \circ \iota_0 = \phi_1, \qquad \pi_2 \circ \Phi \circ \iota_0 = \phi_2.$$

Consequently an efficient evaluator for Φ gives efficient evaluators for both elliptic components ϕ_1 and ϕ_2.

Proof. Let $K/\mathbb{F}_{p^2}$ be any extension field and let $P \in E_0(K)$. The embedding ι_0 sends P to $(P, 0)$. Therefore

$$\Phi(\iota_0(P)) = \Phi(P, 0) = \big(\phi_1(P) + \psi_1(0),\ \phi_2(P) + \psi_2(0)\big).$$

The maps ψ_1 and ψ_2 are homomorphisms, hence preserve the identity:

$$\psi_1(0) = 0, \qquad \psi_2(0) = 0.$$

Thus

$$\Phi(\iota_0(P)) = (\phi_1(P), \phi_2(P)).$$

Applying the product projections gives

$$\pi_1(\Phi(\iota_0(P))) = \phi_1(P), \qquad \pi_2(\Phi(\iota_0(P))) = \phi_2(P).$$

Since these equalities hold on points over every extension field, they are equalities of morphisms. Efficient evaluability follows by composing the evaluator for Φ with the public inclusion ι_0 and the public projections π_1, π_2. □

Corollary 2 (Kani matrix component recovery). *Let a product isogeny be represented in matrix form by*

$$\Phi = \begin{pmatrix} \phi_1 & \widehat{\phi_2'} \\ -\phi_2 & \widehat{\phi_1'} \end{pmatrix} : E_0 \times E_{12} \to E_1 \times E_2.$$

With $\iota_0(P) = (P, 0)$ and product projections π_1, π_2, one has

$$\pi_1 \circ \Phi \circ \iota_0 = \phi_1, \qquad \pi_2 \circ \Phi \circ \iota_0 = -\phi_2.$$

Thus an efficient evaluator for Φ gives efficient evaluators for ϕ_1 and ϕ_2.

Proof. The matrix expression means that, for every point (P, R),

$$\Phi(P, R) = \big(\phi_1(P) + \widehat{\phi_2'}(R),\ -\phi_2(P) + \widehat{\phi_1'}(R)\big).$$

Applying Proposition 2 gives the first component ϕ_1 and the second component $-\phi_2$. The negation map $[-1]_{E_2} : E_2 \to E_2$ is public and efficiently evaluable. Hence an evaluator for $-\phi_2$ immediately gives an evaluator for ϕ_2 by post-composition with $[-1]_{E_2}$. □

4.4 Kani-Style Recoverability and Verification Metadata

Kani-style verification usually checks more than the degree of a single elliptic map. It checks that two maps with a common source satisfy compatibility conditions that define a product or quotient representation. The recoverability interface does not replace those checks; it records the elliptic map that is exposed after the checks have succeeded.

Proposition 3 (Recoverability from certified Kani component data). *Let an accepted response certify valid evaluation metadata for elliptic maps*

$$\kappa : E \to E_1, \qquad \delta : E \to E_\delta,$$

together with higher-dimensional Kani data proving that (κ, δ) satisfies the scheme's verification relation. Suppose the accepted data contains, or publicly induces, an evaluator for a higher-dimensional morphism $\Phi_D : A \to B$ and public maps $\iota : E \to A$, $\pi : B \to E_1$ such that

$$\pi \circ \Phi_D \circ \iota = \kappa.$$

If acceptance certifies $\kappa \neq 0$, then the response datum is HDRecoverable *for (E, E_1), and the recovered elliptic Hom element is κ.*

Proof. The hypotheses instantiate exactly the data required by Definition 8: the accepted transcript provides a valid higher-dimensional evaluator D for $\Phi_D : A \to B$, public maps $\iota : E \to A$ and $\pi : B \to E_1$, and the equality

$$\pi \circ \Phi_D \circ \iota = \kappa.$$

Since acceptance certifies $\kappa \neq 0$, the recovered composite is nonzero. Therefore D is HDRecoverable for (E, E_1), and Theorem 1 gives an efficient evaluator for the recovered elliptic Hom element. The displayed equality identifies that recovered element with κ.

The Kani metadata is not ignored. It is the reason the verifier accepts the higher-dimensional representation and the pair (κ, δ). The proposition uses it only through its certified consequence: the public evaluator exposes κ as the elliptic component needed by extraction. □

4.5 What Recoverability Does Not Assert

Theorem 1 is a transfer theorem for accepted response data. It does not claim that every higher-dimensional isogeny between abelian varieties has a public elliptic component. It does not claim that a native relation between non-split principally polarized abelian surfaces, Jacobians, or general abelian varieties reduces to elliptic OneEnd. It applies only when the accepted transcript includes public morphisms ι and π, or equivalent public evaluation data, that recover an elliptic Hom element.

Recoverability also does not provide a simulator. The theorem starts from an accepted response and recovers the elliptic object that extraction uses. The problem of sampling accepted responses with the correct distribution is separate. That problem is handled by simulation-shaped and pushable hint distributions in the following sections. This separation is the organizing principle of the paper: recoverability governs extraction, whereas hint distributions govern simulation and EndRing-compatible transport.

5 Elliptic Extraction from Recoverable Transcripts

This section proves the extraction theorem that underlies the paper. The theorem is deliberately independent of the representation used by the verifier. A response may be a chain of elliptic isogenies, a product-isogeny representation, a Kani matrix, or a theta-coordinate evaluator. Once the accepted response is HDRecoverable and the protocol supplies the dual maps needed by its special-soundness argument, the extractor outputs an elliptic OneEnd witness. Thus the extracted object is controlled by the recovered elliptic Hom element, not by the dimension of the representation used to transmit or verify it.

5.1 Recoverable Transcript Systems

We first isolate the transcript-level data needed by extraction. The abstraction is SQI-style because it follows the common pattern of SQIsign and its higher-dimensional variants: the public key is an elliptic curve, a commitment determines another elliptic curve, a challenge determines an isogeny out of the public key, and an accepting response recovers an isogeny out of the commitment.

Definition 9 (Extractable challenge datum). *Let E_{pk} and E_{chl} be elliptic curves. A challenge datum is extractable if it publicly determines efficient evaluators for a nonzero isogeny*

$$\phi : E_{\mathsf{pk}} \to E_{\mathsf{chl}}$$

and for its dual

$$\widehat{\phi} : E_{\mathsf{chl}} \to E_{\mathsf{pk}}.$$

The evaluator may be given directly, by a chain of small-degree isogenies, by a public kernel representation, or by any scheme-specified representation interface whose validity implies correct evaluation of ϕ and $\widehat{\phi}$.

Definition 10 (Dual-closed recoverable SQI-style transcript system). *A transcript system Σ is a dual-closed recoverable SQI-style transcript system if every accepting transcript*

$$\tau = (\mathsf{com}, \mathsf{chl}, \mathsf{rsp})$$

for a public key curve E_{pk} satisfies the following conditions.

(i) The commitment com *publicly determines an elliptic curve E_{com}.*

(ii) The challenge chl *is an extractable challenge datum and publicly determines an elliptic curve E_{chl} and an isogeny*

$$\phi_{\mathsf{chl}} : E_{\mathsf{pk}} \to E_{\mathsf{chl}}.$$

(iii) The accepted response datum is HDRecoverable *for $(E_{\mathsf{com}}, E_{\mathsf{chl}})$. We write the recovered nonzero elliptic isogeny as*

$$\sigma_{\mathsf{rsp}} : E_{\mathsf{com}} \to E_{\mathsf{chl}}.$$

(iv) The response representation is dual-closed for accepted transcripts: from rsp *and the public transcript data one can compute an efficient evaluator for*

$$\widehat{\sigma}_{\mathsf{rsp}} : E_{\mathsf{chl}} \to E_{\mathsf{com}}.$$

The dual-closure condition is not a hidden assumption about arbitrary higher-dimensional maps. It records a property that concrete special-soundness proofs use: the extractor composes one recovered response with the dual of the other. If a protocol exposes an equivalent inverse relation rather than an explicit dual evaluator, the definition can be instantiated through the corresponding evaluator supplied by the protocol's representation interface.

Definition 11 (Extractable accepting pair). *Let Σ be a dual-closed recoverable SQI-style transcript system. An ordered pair of accepting transcripts*

$$\tau_1 = (\mathsf{com}, \mathsf{chl}_1, \mathsf{rsp}_1), \qquad \tau_2 = (\mathsf{com}, \mathsf{chl}_2, \mathsf{rsp}_2)$$

for the same public key curve E_{pk} is an extractable accepting pair if the commitments are identical and the challenges are distinct according to the challenge equality relation of Σ. The common commitment determines E_{com}. For $i \in \{1, 2\}$, the challenge determines

$$\phi_i : E_{\mathsf{pk}} \to E_{\mathsf{chl},i}, \qquad \widehat{\phi}_i : E_{\mathsf{chl},i} \to E_{\mathsf{pk}},$$

and the response recovers

$$\sigma_i : E_{\mathsf{com}} \to E_{\mathsf{chl},i}, \qquad \widehat{\sigma}_i : E_{\mathsf{chl},i} \to E_{\mathsf{com}}.$$

The only scheme-specific part of elliptic extraction is the argument that two different challenges cannot produce a scalar cycle. In SQIsign-style protocols this is proved by degree, kernel, ideal, or level-structure arguments. The transfer theorem below treats that argument as a property of the transcript system, and the concrete instantiation sections prove it for the schemes considered in the paper.

Definition 12 (Cycle scalar-exclusion property). *A dual-closed recoverable SQI-style transcript system has the cycle scalar-exclusion property if, for every extractable accepting pair (τ_1, τ_2), the endomorphism*

$$\alpha(\tau_1, \tau_2) = \widehat{\phi}_2 \circ \sigma_2 \circ \widehat{\sigma}_1 \circ \phi_1 \in \mathrm{End}(E_{\mathsf{pk}})$$

is not a scalar multiplication map $[m]_{E_{\mathsf{pk}}}$ for any $m \in \mathbb{Z}$.

5.2 The Extraction Transfer Theorem

The transfer theorem is the formal bridge from response recoverability to cryptographic special soundness. Given two accepting transcripts with the same commitment, recoverability supplies two elliptic response maps from the commitment curve to the two challenge curves. Dual closure lets the extractor reverse the first response and the second challenge. The extractor therefore obtains the elliptic cycle

$$E_{\mathsf{pk}} \xrightarrow{\phi_1} E_{\mathsf{chl},1} \xrightarrow{\widehat{\sigma}_1} E_{\mathsf{com}} \xrightarrow{\sigma_2} E_{\mathsf{chl},2} \xrightarrow{\widehat{\phi}_2} E_{\mathsf{pk}}.$$

The theorem does not hide the protocol-specific non-scalarity argument. It assumes a scalar-exclusion property and proves that, once that property has been established for the concrete protocol, the extracted witness is an elliptic OneEnd witness.

Theorem 2 (Elliptic special-soundness transfer). *Let Σ be a dual-closed recoverable SQI-style transcript system satisfying the cycle scalar-exclusion property. Given any extractable accepting pair (τ_1, τ_2) for the same public key curve E_{pk}, there is a deterministic polynomial-time extractor that outputs an efficient evaluator for an endomorphism*

$$\alpha \in \mathrm{End}(E_{\mathsf{pk}}) \setminus \mathbb{Z}.$$

Consequently, the extractor outputs a valid OneEnd witness for E_{pk}.

Proof. Let

$$\tau_i = (\mathsf{com}, \mathsf{chl}_i, \mathsf{rsp}_i) \qquad (i \in \{1, 2\})$$

be an extractable accepting pair. By Definition 11, the common commitment determines an elliptic curve E_{com}. For each i, the challenge datum provides efficient evaluators for nonzero isogenies

$$\phi_i : E_{\mathsf{pk}} \to E_{\mathsf{chl},i} \quad \text{and} \quad \widehat{\phi}_i : E_{\mathsf{chl},i} \to E_{\mathsf{pk}}.$$

The accepted response is $\mathsf{HDRecoverable}$ for $(E_{\mathsf{com}}, E_{\mathsf{chl},i})$. Hence, by Theorem 1, it gives an efficient evaluator for a nonzero elliptic isogeny

$$\sigma_i : E_{\mathsf{com}} \to E_{\mathsf{chl},i}.$$

Dual-closure gives an efficient evaluator for

$$\widehat{\sigma}_i : E_{\mathsf{chl},i} \to E_{\mathsf{com}}.$$

The extractor defines

$$\alpha = \widehat{\phi}_2 \circ \sigma_2 \circ \widehat{\sigma}_1 \circ \phi_1.$$

The domains and codomains match in the chain

$$E_{\mathsf{pk}} \xrightarrow{\phi_1} E_{\mathsf{chl},1} \xrightarrow{\widehat{\sigma}_1} E_{\mathsf{com}} \xrightarrow{\sigma_2} E_{\mathsf{chl},2} \xrightarrow{\widehat{\phi}_2} E_{\mathsf{pk}}.$$

Therefore α is an element of $\mathrm{End}(E_{\mathsf{pk}})$.

We now construct the evaluator. On input a point $P \in E_{\mathsf{pk}}(K)$, for an extension field $K/\mathbb{F}_{p^2}$ of polynomial degree, the extractor computes

$$P_1 = \phi_1(P), \qquad P_2 = \widehat{\sigma}_1(P_1), \qquad P_3 = \sigma_2(P_2), \qquad P_4 = \widehat{\phi}_2(P_3).$$

It outputs P_4. By the definition of composition,

$$P_4 = \widehat{\phi}_2\big(\sigma_2(\widehat{\sigma}_1(\phi_1(P)))\big) = \alpha(P).$$

Each component evaluator is polynomial-time by extractability, HD-recoverability, and dual-closure. A constant number of polynomial-time evaluations is polynomial-time, so the extractor obtains an efficient evaluator for α.

The endomorphism α is nonzero. Each of ϕ_1, $\widehat{\sigma}_1$, σ_2, and $\widehat{\phi}_2$ is a nonzero isogeny between elliptic curves. If $\eta : E \to F$ and $\theta : F \to G$ are nonzero isogenies, then both are surjective with finite kernels by Proposition 1. The composite $\theta \circ \eta$ is surjective, and

$$\ker(\theta \circ \eta) = \eta^{-1}(\ker \theta)$$

is finite because it is a finite union of cosets of $\ker \eta$, one union component for each point of the finite group scheme $\ker \theta$. Thus $\theta \circ \eta$ is an isogeny and is not zero. Repeating this argument over the four factors shows $\alpha \neq 0$.

Finally, the pair has distinct challenges and Σ satisfies the cycle scalar-exclusion property. Therefore

$$\alpha \notin \mathbb{Z},$$

where $\mathbb{Z}$ is identified with scalar endomorphisms. The extractor outputs an efficient evaluator for a non-scalar element of $\mathrm{End}(E_{\mathsf{pk}})$. This is precisely a OneEnd witness in the sense of Definition 1. □

Corollary 3 (HD responses do not change the extracted witness). *Let Σ satisfy the hypotheses of Theorem 2. Suppose the response maps in an extractable accepting pair are obtained from accepted higher-dimensional data D_1, D_2 by*

$$\sigma_i = \pi_i \circ \Phi_{D_i} \circ \iota_i \qquad (i \in \{1, 2\}).$$

Then the extracted witness is the elliptic endomorphism

$$\widehat{\phi}_2 \circ (\pi_2 \circ \Phi_{D_2} \circ \iota_2) \circ (\pi_1 \circ \widehat{\Phi_{D_1}} \circ \iota_1) \circ \phi_1 \in \mathrm{End}(E_{\mathsf{pk}}) \setminus \mathbb{Z},$$

where the hat denotes the dual evaluator supplied by the transcript system. The ambient dimension of the accepted response data affects verification and representation, but not the type of witness extracted by special soundness.

Proof. By Theorem 1, each accepted datum D_i recovers the elliptic Hom element

$$\sigma_i = \pi_i \circ \Phi_{D_i} \circ \iota_i.$$

The extractor of Theorem 2 uses the response data only through σ_1, $\widehat{\sigma}_1$, and σ_2. Substituting the recovered formulas for σ_1 and σ_2 gives the displayed endomorphism. The transfer theorem proves that the resulting endomorphism is efficiently evaluable and non-scalar. No step of the extraction proof outputs a basis of a higher-dimensional endomorphism algebra or a native endomorphism of an abelian surface; the output is an elliptic endomorphism of E_{pk}. □

5.3 Scalar-Exclusion Criteria

The transfer theorem deliberately separates representation recovery from the protocol-specific proof that the extracted cycle is non-scalar. The following criteria cover the cases most often used later in the paper.

Proposition 4 (Degree-based scalar exclusion). *Let E be an elliptic curve and let $\alpha \in \mathrm{End}(E)$ be a nonzero endomorphism with degree $N = \deg(\alpha)$. If N is not a square integer, then $\alpha \notin \mathbb{Z}$.*

Proof. Assume for contradiction that $\alpha \in \mathbb{Z}$. Then $\alpha = [m]_E$ for some $m \in \mathbb{Z}$. For every integer m, scalar multiplication on an elliptic curve has degree

$$\deg([m]_E) = m^2.$$

For $m = 0$ this is immediate. For $m \neq 0$, the kernel of $[m]_E$ is the finite group scheme $E[m]$, whose rank is m^2; the degree of the separable-inseparable multiplication morphism equals the rank of this kernel group scheme. Therefore

$$N = \deg(\alpha) = \deg([m]_E) = m^2,$$

which is a square integer. This contradicts the hypothesis that N is not a square. Hence $\alpha \notin \mathbb{Z}$. □

Proposition 5 (Kernel-line scalar exclusion). *Let $q \neq p$ be prime, let E be an elliptic curve, and let $L_1, L_2 \subset E[q]$ be distinct lines. Suppose an extracted endomorphism $\alpha \in \mathrm{End}(E)$ has the following property on q-torsion: there exists a nonzero scalar $c \in \mathbb{F}_q^\times$ such that*

$$\alpha(L_1) = 0 \quad \textit{and} \quad \alpha(P) = [c]P \textit{ for every } P \in L_2.$$

Then $\alpha \notin \mathbb{Z}$.

Proof. Assume, for contradiction, that $\alpha = [m]_E$ for some $m \in \mathbb{Z}$. Restricting to $E[q]$, the scalar endomorphism acts as multiplication by $\bar{m} \in \mathbb{F}_q$.

Because $\alpha(L_1) = 0$, every point $P \in L_1$ satisfies $[\bar{m}]P = 0$. Since L_1 is one-dimensional over $\mathbb{F}_q$ and contains a nonzero point, this implies $\bar{m} = 0$. Thus $[m]$ acts as the zero map on all of $E[q]$.

But on L_2, the same endomorphism is assumed to act as multiplication by $c \in \mathbb{F}_q^\times$. Since L_2 contains a nonzero point, this action is nonzero. This contradicts $\bar{m} = 0$. Therefore α cannot be scalar. □

Use in concrete schemes. SQIsign2D–West uses a more refined degree and cyclic-factor argument, because the extracted cycle includes response degrees and challenge degrees. SQIPrime-style arguments use the fact that two accepting responses must factor through distinct challenge lines while yielding a cycle on the public-key curve. In both cases the scheme-specific proof instantiates Definition 12; once that property is established, Theorem 2 supplies the elliptic OneEnd extractor.

6 Fiat–Shamir with Hints and Pushable Distributions

The extraction theorem outputs OneEnd witnesses. A signature proof, however, also needs simulation. Higher-dimensional SQIsign variants often obtain simulation through auxiliary oracles that sample represented isogenies, auxiliary paths, torsion images, or constrained response data. Following the hint methodology used in complete SQIsign security proofs [Ård+25], this section separates two roles of auxiliary information. Simulation-shaped hints reproduce the simulator's transcript distribution. Pushable uniform hints are designed to survive the self-reductions that connect OneEnd-style extraction to EndRing-style hardness.

6.1 Hinted Relations and Simulation Hints

A hint is auxiliary information supplied by the security game, not a value sent by an honest signer and not a value that the verifier must compute. The purpose of a simulation-shaped hint is to make an oracle-assisted simulator non-interactive: the hint contains exactly the represented isogenies, torsion images, or auxiliary paths that the simulator would otherwise request from sampling oracles. When the oracle-assisted simulator is adaptive, the hint also contains the simulator coins that determine the adaptive query path. Thus the hinted simulator replays the same path and consumes the stored oracle answers in the same order. A security statement relative to such hints therefore says that the underlying relation remains hard even after the adversary receives polynomially many samples from the specified hint distribution.

Two hint distributions appear throughout the paper. The distribution $\mathcal{H}_{\mathsf{sim}}$ is chosen to match the simulator's transcript distribution. A pushable distribution $\mathcal{H}_{\mathsf{unif}}$, by contrast, is chosen so that hints can be transported across the self-reductions used to relate OneEnd-style extraction to EndRing-style hardness. These roles need not be served by the same distribution; when they differ, the gap is recorded by an explicit HintDist assumption.

Definition 13 (Simulation-shaped hint family). *Let Σ be an identification protocol with public statement space X. A simulation-shaped hint family for Σ is a family*

$$\mathcal{H}_{\mathsf{sim}} = \{\mathcal{H}_{\mathsf{sim},x}\}_{x\in X}$$

of polynomial-size distributions, together with a deterministic polynomial-time algorithm

$$\mathsf{Sim}^{\mathcal{H}_{\mathsf{sim}}}(x, h; r)$$

that takes a statement x, a hint $h \leftarrow \mathcal{H}_{\mathsf{sim},x}$, and simulation randomness r, and outputs a transcript of Σ. The family has $\varepsilon_{\mathsf{zk}}$-hint-assisted weak honest-verifier zero knowledge if the transcript distribution of $\mathsf{Sim}^{\mathcal{H}_{\mathsf{sim}}}$ is within distinguishing advantage $\varepsilon_{\mathsf{zk}}$ of an honest accepting transcript for every valid statement.

The definition is intentionally distributional. It does not require the hint to be efficiently samplable by the verifier or by an external party. In a security reduction, hints are sampled by the hardness game. This is the same modeling role played by hint distributions in complete SQIsign proofs: the adversary is allowed polynomially many hints, and hardness is asserted relative to that auxiliary information.

Theorem 3 (Oracle-to-simulation-hint normal form). *Let $\mathsf{Sim}^{\mathcal{O}}$ be a polynomial-time simulator that makes at most $t(\lambda)$ calls to stateless oracle algorithms*

$$\mathcal{O}_1, \ldots, \mathcal{O}_m,$$

where each oracle output has polynomial bitlength. Suppose the simulator's next message and next oracle query are deterministic polynomial-time functions of the public statement, a simulator-coin string r_{sim}, and the oracle outputs already received. Then there is a simulation-shaped hint family $\mathcal{H}_{\mathsf{sim}}$ and a hint-assisted simulator $\mathsf{Sim}^{\mathcal{H}_{\mathsf{sim}}}$ whose output distribution is identical to the output distribution of $\mathsf{Sim}^{\mathcal{O}}$.

Proof. Fix a statement x. The hint distribution $\mathcal{H}_{\mathsf{sim},x}$ first samples the simulator-coin string r_{sim} with exactly the same distribution used by $\mathsf{Sim}^{\mathcal{O}}$. It then runs the stateless oracle algorithms with fresh oracle randomness along the adaptive query path determined by x, r_{sim}, and the oracle answers already sampled. Because the simulator is polynomial-time and makes at most $t(\lambda)$ oracle calls, the list of queries and answers has polynomial total bitlength. A hint is defined to be

$$h = (r_{\mathsf{sim}}, (i_1, Q_1, y_1), \ldots, (i_s, Q_s, y_s)),$$

where i_j identifies the oracle queried at step j, Q_j is the query issued at that step, and y_j is the answer sampled from $\mathcal{O}_{i_j}(Q_j)$. The query entries are not needed by the replay when parsing is unambiguous, but including them makes the coupling explicit. The distribution of h is exactly the joint distribution induced by one oracle-assisted execution of $\mathsf{Sim}^{\mathcal{O}}$.

The hint-assisted simulator $\mathsf{Sim}^{\mathcal{H}_{\mathsf{sim}}}$ ignores any external simulator coins except for transcript-independent formatting randomness explicitly allowed by the protocol, reads r_{sim} from h, and runs the same deterministic transition function as $\mathsf{Sim}^{\mathcal{O}}$. Whenever the oracle-assisted simulator would issue its j-th query, the replay checks that the query equals the stored Q_j, reads the stored answer y_j, and continues. The equality of the stored and replayed queries follows by induction on j: it is true for $j = 1$ because the same public statement and r_{sim} are used, and if it is true up to $j-1$, then the deterministic transition function has seen the same prior answers and therefore issues the same j-th query.

For every fixed choice of r_{sim} and oracle randomness, the transcript output by the replay is equal to the transcript output by the original oracle-assisted simulator. Averaging over r_{sim} and the oracle randomness gives equality of output distributions. The replay is polynomial-time because it performs the same polynomial-time computation as $\mathsf{Sim}^{\mathcal{O}}$ and reads a polynomial-size hint. Thus $\mathcal{H}_{\mathsf{sim}}$ and $\mathsf{Sim}^{\mathcal{H}_{\mathsf{sim}}}$ have the required identical distribution. □

Table 1. Alignment between the complete SQIsign Fiat–Shamir-with-hints framework and the notation used in this paper. The table is a bookkeeping device for the reductions below.

Complete SQIsign hint-framework object	Object used in this paper
public-coin identification protocol	Σ, $\Sigma_{2D\text{-}W}$, or $\Sigma_{\mathsf{SQIP2D\text{-}UAux}}$
hard relation with hints	$\mathsf{HintOneEnd}^{\mathcal{H}}$ or $\mathsf{HintEndRing}^{\mathcal{H}}$
hint family	$\mathcal{H}_{\mathsf{sim}}^{2D\text{-}W}$, $\mathcal{H}_{\mathsf{unif}}^{2D\text{-}W}$, or $\mathcal{H}_{q\text{-}\mathsf{unif}}^{\mathsf{SQIP2D\text{-}UAux}}$
random-oracle query bound	Q_{RO}, distinct from the SQIPrime challenge degree q
signing-query bound and number of hints	s
challenge-guessing term	$\lvert\mathcal{C}\rvert^{-1}$
commitment min-entropy term	$(Q_{\mathsf{RO}}+s+1)\cdot s\cdot \mathsf{MEnt}(\Sigma)$
simulation-to-uniform replacement	s-sample HintDist, $2D$-West-HintDist, or QHintDist

6.2 Fiat–Shamir Security with Hints

The next theorem records the reduction shape used later. It follows the complete SQIsign Fiat–Shamir-with-hints framework with the notation of Table 1: a forger either violates the hinted relation extracted by special soundness, distinguishes simulated signing transcripts, guesses a challenge, or causes a commitment collision/min-entropy failure. The displayed bound deliberately uses the conservative $(Q_{\mathsf{RO}}+1)$-multiplied form for simulation and hint-replacement losses, so later instantiations remain aligned with the complete SQIsign proof template.

Definition 14 (Commitment min-entropy loss). *Let Σ be an identification protocol with commitment space $\mathcal{C}_x$ for statement x. For a commitment distribution Com_x, define*

$$\mathsf{MEnt}(\Sigma,x)=\max_{c\in\mathcal{C}_x}\Pr[c\leftarrow\mathsf{Com}_x].$$

For a statement distribution $\mathcal{X}$, write

$$\mathsf{MEnt}(\Sigma)=\max_{x\in\mathrm{supp}(\mathcal{X})}\mathsf{MEnt}(\Sigma,x).$$

This quantity upper-bounds the probability that a random-oracle query or a simulated signing query hits a preselected honest commitment.

The reduction below is not a new Fiat–Shamir theorem; it is the notation-normalized form of the Fiat–Shamir-with-hints framework used by the complete SQIsign proof, specialized to the recoverable-transcript setting of this paper. The same quantity s denotes both the number of signing queries and the number of hints consumed by the simulated signing experiment.

Theorem 4 (Fiat–Shamir with hints, reduction template). *Let Σ be a public-coin identification protocol with challenge space $\mathcal{C}$, challenge size $\lvert\mathcal{C}\rvert$, a dual-closed recoverable transcript system, and a special-sound extractor as in Theorem 2. Let $\mathcal{H}$ be a hint family for the public statements of Σ. Assume that:*

(i) the hint-assisted simulator for Σ has distinguishing error at most $\varepsilon_{\mathsf{zk}}(s)$ for s simulated signing transcripts;
(ii) the forking extractor succeeds except when the adversary repeats the same challenge, guesses a challenge in $\mathcal{C}$, or hits a commitment event accounted for by $\mathsf{MEnt}(\Sigma)$;
(iii) from two accepting transcripts with the same commitment and distinct challenges, the extractor outputs a witness for $\mathsf{HintOneEnd}^{\mathcal{H}}$.

Then for every random-oracle EUF-CMA adversary $\mathcal{A}$ making at most Q_{RO} random-oracle queries and s signing queries, there is an algorithm $\mathcal{B}$, with running time polynomial in that of $\mathcal{A}$, such that

$$\begin{aligned}\mathrm{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{SIG}[\Sigma]}(\mathcal{A})\le{}&(Q_{\mathsf{RO}}+1)\cdot\left(\mathrm{Adv}^{\mathsf{HintOneEnd}^{\mathcal{H}}}(\mathcal{B})+\varepsilon_{\mathsf{zk}}(s)+\lvert\mathcal{C}\rvert^{-1}\right)\\&+(Q_{\mathsf{RO}}+s+1)\cdot s\cdot\mathsf{MEnt}(\Sigma).\end{aligned}$$

Proof. Let ϵ be the success probability of $\mathcal{A}$ in the real EUF-CMA experiment. Define a simulated signing experiment in which every signing query is answered by drawing the hints prescribed by $\mathcal{H}$, running the hint-assisted simulator, and programming the random oracle at the resulting commitment-message point. By assumption (i), the joint distribution of the s simulated signing transcripts differs from the joint distribution of s honest signing transcripts by at most $\varepsilon_{\mathsf{zk}}(s)$. Hence the adversary succeeds in the simulated experiment with probability at least

$$\epsilon - \varepsilon_{\mathsf{zk}}(s).$$

The displayed theorem uses a conservative form in which this simulation loss is charged inside the same outer $Q_{\mathsf{RO}} + 1$ factor as the forking loss. This is the loss form used by the complete SQIsign hint framework and suffices for all later instantiations.

Consider a successful forgery in the simulated experiment, written as a commitment $c^\star$, a message $m^\star$, a Fiat–Shamir challenge $e^\star \in \mathcal{C}$, and a response $z^\star$. The reduction chooses one candidate fork point among the adversary's Q_{RO} random-oracle queries and the final forgery point. There are at most $Q_{\mathsf{RO}} + 1$ candidates, so guessing the correct point costs the multiplicative factor $Q_{\mathsf{RO}} + 1$.

The reduction rewinds $\mathcal{A}$ with the same coins and the same oracle answers before the fork point, but replaces the oracle answer at that point by a fresh uniform element of $\mathcal{C}$. If the fresh answer equals the old challenge, then the two accepting transcripts have equal challenges; this bad event has probability exactly $|\mathcal{C}|^{-1}$. Otherwise, except for commitment-hit events, the two executions give accepting transcripts with the same commitment and distinct challenges.

We now bound commitment-hit events. In the simulated signing experiment the oracle is programmed at s simulated commitments. For every one of the adversary's Q_{RO} random-oracle queries, the final forgery point, and the signing-query consistency checks, the probability of hitting any particular simulated commitment is at most $\mathsf{MEnt}(\Sigma)$. Taking a union bound over at most $(Q_{\mathsf{RO}} + s + 1) \cdot s$ pairs gives

$$(Q_{\mathsf{RO}} + s + 1) \cdot s \cdot \mathsf{MEnt}(\Sigma).$$

Outside the repeated-challenge and commitment-hit events, assumption (iii) applies to the two accepting transcripts. The extractor therefore outputs a witness for $\mathsf{HintOneEnd}^{\mathcal{H}}$. Thus the success probability in the simulated experiment is at most

$$(Q_{\mathsf{RO}} + 1) \cdot \left(\mathrm{Adv}^{\mathsf{HintOneEnd}^{\mathcal{H}}}(\mathcal{B}) + |\mathcal{C}|^{-1}\right) + (Q_{\mathsf{RO}} + s + 1) \cdot s \cdot \mathsf{MEnt}(\Sigma).$$

Combining this with the simulation replacement inequality and then using

$$\varepsilon_{\mathsf{zk}}(s) \le (Q_{\mathsf{RO}} + 1) \cdot \varepsilon_{\mathsf{zk}}(s)$$

gives the displayed bound. The reduction is polynomial-time because it performs polynomially many executions of $\mathcal{A}$ and invokes only the polynomial-time simulator, extractor, and hint-game interface. □

6.3 Hint Replacement and Pushability

Simulation-shaped hints and pushable hints serve different purposes. A simulator wants the distribution that reproduces transcripts. A reduction from hinted $\mathsf{EndRing}$ to hinted OneEnd wants a distribution that can be transported along self-reduction isogenies. The bridge between them is a hint-distinguishability assumption.

Definition 15 (Hint-distinguishing advantage). *Let $\mathcal{H}^0 = \{\mathcal{H}^0_x\}$ and $\mathcal{H}^1 = \{\mathcal{H}^1_x\}$ be two hint families over the same statement distribution. The q-hint distinguishing advantage of an algorithm $\mathcal{A}$ is*

$$\begin{aligned}\mathrm{Adv}^{q\text{-}\mathsf{HintDist}}_{\mathcal{H}^0,\mathcal{H}^1}(\mathcal{A}) = \big|&\Pr\left[1 \leftarrow \mathcal{A}(x, h_1, \ldots, h_q) : h_i \leftarrow \mathcal{H}^0_x\right] \\ &- \Pr\left[1 \leftarrow \mathcal{A}(x, h_1, \ldots, h_q) : h_i \leftarrow \mathcal{H}^1_x\right]\big|,\end{aligned}$$

where x is sampled from the statement distribution and all hints are sampled independently conditioned on x.

Lemma 2 (Hint replacement). *Let* Sim *be a deterministic polynomial-time algorithm that receives a statement and q hints. For every transcript distinguisher $\mathcal{D}$ and every pair of hint families $\mathcal{H}^0, \mathcal{H}^1$, there is a hint distinguisher $\mathcal{B}$ with running time equal to that of* Sim *plus that of $\mathcal{D}$, up to polynomial overhead, such that*

$$\begin{aligned}&\left|\Pr[1 \leftarrow \mathcal{D}(\mathsf{Sim}(x, h_1^0, \ldots, h_q^0))] - \Pr[1 \leftarrow \mathcal{D}(\mathsf{Sim}(x, h_1^1, \ldots, h_q^1))]\right| \\ &\qquad \leq \mathrm{Adv}^{q\text{-}\mathsf{HintDist}}_{\mathcal{H}^0,\mathcal{H}^1}(\mathcal{B}),\end{aligned}$$

where $h_i^b \leftarrow \mathcal{H}_x^b$ independently for $b \in \{0,1\}$.

Proof. The distinguisher $\mathcal{B}$ receives $(x, h_1, \ldots, h_q)$, computes

$$T = \mathsf{Sim}(x, h_1, \ldots, h_q),$$

and outputs $\mathcal{D}(T)$. If the hints come from $\mathcal{H}_x^0$, then the output distribution of $\mathcal{B}$ is exactly the first probability distribution in the displayed difference. If the hints come from $\mathcal{H}_x^1$, then the output distribution is exactly the second. Thus the distinguishing advantage of $\mathcal{B}$ is the left-hand side. The running-time statement follows from the construction. □

Definition 16 (Pushable hint family). *Let $\mathcal{R}$ be a class of isogenies between public statements, and let $\mathcal{H}_{\mathsf{unif}} = \{\mathcal{H}_{\mathsf{unif},x}\}$ be a hint family. The family is pushable through $\mathcal{R}$ if there is a polynomial-time algorithm*

$$\mathsf{PushHint}(x, h, \rho)$$

such that, for every $\rho : x \to x'$ in $\mathcal{R}$, if $h \leftarrow \mathcal{H}_{\mathsf{unif},x}$, then

$$\mathsf{PushHint}(x, h, \rho)$$

is distributed identically to $\mathcal{H}_{\mathsf{unif},x'}$. The computational version replaces equality of distributions by computational indistinguishability, with the distinguishing loss stated explicitly.

Proposition 6 (Pushability gives EndRing-compatible hinted hardness). *Let $\mathcal{H}_{\mathsf{unif}}$ be a hint family for supersingular elliptic-curve statements. Suppose $\mathcal{H}_{\mathsf{unif}}$ is pushable through the self-reduction isogenies used by a reduction from* EndRing *to* OneEnd*. Then the same reduction can be lifted to the hinted setting: an algorithm that solves $\mathsf{HintOneEnd}^{\mathcal{H}_{\mathsf{unif}}}$ after receiving pushed hints yields an algorithm for $\mathsf{HintEndRing}^{\mathcal{H}_{\mathsf{unif}}}$, with the same algebraic loss as the unhinted reduction and with no additional distributional loss beyond the stated pushability error.*

Proof. Consider the unhinted self-reduction. Starting from an EndRing instance E, it samples or computes isogenies

$$\rho_j : E \to E_j$$

that produce OneEnd instances E_j. From non-scalar endomorphisms output for enough related E_j's, the unhinted reduction transports the endomorphisms back along the ρ_j's and reconstructs a basis of $\mathrm{End}(E)$, with whatever algebraic success probability and loss the unhinted reduction has.

In the hinted experiment, the adversary for $\mathsf{HintEndRing}^{\mathcal{H}_{\mathsf{unif}}}$ receives polynomially many hints sampled from $\mathcal{H}_{\mathsf{unif},E}$. To call the $\mathsf{HintOneEnd}^{\mathcal{H}_{\mathsf{unif}}}$ solver on E_j, the reduction applies $\mathsf{PushHint}(E, h, \rho_j)$ to each available hint h. By pushability, the resulting list has exactly the distribution of independent hints from $\mathcal{H}_{\mathsf{unif},E_j}$ when pushability is exact; in the computational variant the discrepancy is precisely the declared pushability distinguishing error. Therefore the OneEnd solver is invoked with correctly distributed hinted instances. Its outputs are valid non-scalar endomorphisms of the E_j's. The rest of the reduction is the same algebraic transport and reconstruction used in the unhinted proof. Thus the unhinted reduction lifts to the hinted setting, and no extra loss appears except the pushability error already accounted for. □

Theorem 5 (Simulation-to-pushable-hint conversion). *Let Σ satisfy the hypotheses of Theorem 4 with simulation-shaped hints $\mathcal{H}_{\mathsf{sim}}$. Let $\mathcal{H}_{\mathsf{unif}}$ be a pushable hint family for the same public statements. Assume that $\mathcal{H}_{\mathsf{unif}}$ supports the same transcript-formatting map as the $\mathcal{H}_{\mathsf{sim}}$-simulation experiment, or that an explicitly supplied deterministic formatter maps $\mathcal{H}_{\mathsf{unif}}$-hints to the same*

public transcript fields used by the simulator. If $\mathcal{H}_{\mathsf{sim}}$ and $\mathcal{H}_{\mathsf{unif}}$ are s-hint indistinguishable, then the EUF-CMA theorem for Σ can be restated using $\mathcal{H}_{\mathsf{unif}}$, with the additional loss

$$(Q_{\mathsf{RO}} + 1) \cdot \mathrm{Adv}^{s\text{-}\mathsf{HintDist}}_{\mathcal{H}_{\mathsf{sim}},\mathcal{H}_{\mathsf{unif}}}(\mathcal{D})$$

when the theorem is written in the safe forking-bound form of Theorem 4. If $\mathcal{H}_{\mathsf{unif}}$ is pushable through the relevant EndRing*-to-*OneEnd *self-reductions, the resulting theorem is EndRing-compatible in the sense of Proposition 6.*

Proof. Run the proof of Theorem 4 in the experiment where the simulator receives hints from $\mathcal{H}_{\mathsf{unif}}$ instead of $\mathcal{H}_{\mathsf{sim}}$. The transcript-formatting hypothesis ensures that the same deterministic simulation wrapper can be executed after the replacement. The only part of the proof that depends on the hint distribution is the transcript simulation for the s signing queries. By Lemma 2, for every adversary and every final transcript distinguisher derived from its success predicate, replacing s independent $\mathcal{H}_{\mathsf{sim}}$ hints by s independent $\mathcal{H}_{\mathsf{unif}}$ hints changes the success probability by at most the s-hint distinguishing advantage between the two distributions. All other parts of the forking argument, including challenge guessing and commitment min-entropy, are unchanged. Therefore the same EUF-CMA bound holds with the additional HintDist loss charged in the same $(Q_{\mathsf{RO}} + 1)$ forking factor as the simulation error.

If $\mathcal{H}_{\mathsf{unif}}$ is pushable through the self-reduction isogenies used to relate EndRing and OneEnd, then Proposition 6 lifts the hinted OneEnd hardness term in the bound to the corresponding hinted EndRing hardness statement, with only the pushability error stated in the definition of $\mathcal{H}_{\mathsf{unif}}$. This proves both claims. □

6.4 Consequences for the Two Main Instantiations

The next sections instantiate the general framework in two directions. For SQIsign2D–West, the simulation-shaped hints package the oracle outputs used by its simulator, while the pushable distribution and its HintDist assumption yield an EndRing-compatible corollary. For SQIPrime2D-UAux, the auxiliary UAux distribution repairs the context dependence of the published AIO layer, and a stronger challenge-carrying uniform hint distribution supplies the pushability needed for an EndRing-compatible statement. In both cases, the proof structure is the same: recoverability controls extraction, $\mathcal{H}_{\mathsf{sim}}$ controls simulation, and $\mathcal{H}_{\mathsf{unif}}$ controls transport through hardness reductions.

7 Hint-Based EUF-CMA Security of SQIsign2D–West

This section gives the first concrete theorem package. SQIsign2D–West [Bas+24] uses two-dimensional representations to evaluate and transmit response isogenies, but its special-soundness witness is elliptic. The theorem below makes this separation explicit: the extraction term is a hinted OneEnd relation, while the higher-dimensional information used by the zero-knowledge simulator is represented by a non-interactive hint distribution. The statement applies first to the uncompressed Fiat–Shamir transform that contains the commitment. The commitment-recoverable signature format is handled at the end of the section.

7.1 Transcript Profile and Extracted Relation

Let $\Sigma_{2D\text{-}W}$ denote the SQIsign2D–West identification protocol. A public key is a supersingular elliptic curve E_{pk}. An accepting transcript contains a commitment curve E_{com}, a challenge curve E_{chl}, and represented response data from which the verifier can evaluate the response isogeny in the two-dimensional representation. The challenge and response degrees are governed by parameters e_{chl} and e_{rsp}. The special-soundness theorem of SQIsign2D–West applies in the regime

$$e_{\mathsf{chl}} + e_{\mathsf{rsp}} \leq e.$$

Definition 17 (SQIsign2D–West hinted OneEnd relation). *The relation $R^{2D\text{-}W}_{\mathsf{OneEnd}}$ contains pairs $(E_{\mathsf{pk}}, \alpha)$ such that E_{pk} is a SQIsign2D–West public-key curve and α is an efficient evaluator for an endomorphism*

$$\alpha \in \mathrm{End}(E_{\mathsf{pk}}) \setminus \mathbb{Z}.$$

For a hint family $\mathcal{H}$, the hinted relation $\mathsf{HintOneEnd}^{\mathcal{H}}_{2D\text{-}W}$ *is the relation-with-hints experiment of Definition 2 instantiated with* $R^{2D\text{-}W}_{\mathsf{OneEnd}}$.

Proposition 7 (Recoverable SQIsign2D–West responses). *Every accepting SQIsign2D–West response datum that verifies through its two-dimensional representation is an* HDRecoverable *response in the sense of Definition 8. The recovered elliptic map is the response isogeny*

$$\sigma : E_{\mathsf{com}} \to E_{\mathsf{chl}}$$

represented by the two-dimensional datum.

Proof. The two-dimensional datum supplied to the verifier is accepted only if the representation interface can evaluate the relevant higher-dimensional morphism and certify the component that carries the elliptic response. In the notation of Definition 8, the elliptic source is E_{com}, the elliptic target is E_{chl}, and the public representation data supplies an embedding-projection pair

$$E_{\mathsf{com}} \xrightarrow{\iota} A, \qquad B \xrightarrow{\pi} E_{\mathsf{chl}}$$

for the two-dimensional source and target objects used by the verifier. Applying Evalto the represented higher-dimensional morphism $\Phi_D : A \to B$ and then projecting gives

$$\pi \circ \Phi_D \circ \iota : E_{\mathsf{com}} \to E_{\mathsf{chl}}.$$

Acceptance asserts exactly that this component is the response isogeny used in the transcript. Thus the response satisfies the HDRecoverable interface, and the recovered elliptic map is σ. □

Theorem 6 (SQIsign2D–West elliptic special soundness). *Assume $e_{\mathsf{chl}} + e_{\mathsf{rsp}} \le e$. From two accepting SQIsign2D–West transcripts with the same commitment and distinct challenges, there is a polynomial-time extractor that outputs an efficient evaluator for a non-scalar endomorphism of E_{pk}. Hence the extracted witness is a witness for* $R^{2D\text{-}W}_{\mathsf{OneEnd}}$.

Proof. By Proposition 7, each accepting response yields an efficiently evaluable elliptic response isogeny. Let the two recovered response isogenies be

$$\sigma_1 : E_{\mathsf{com}} \to E_{\mathsf{chl},1}, \qquad \sigma_2 : E_{\mathsf{com}} \to E_{\mathsf{chl},2}.$$

Let $\phi_1 : E_{\mathsf{pk}} \to E_{\mathsf{chl},1}$ and $\phi_2 : E_{\mathsf{pk}} \to E_{\mathsf{chl},2}$ be the challenge isogenies determined by the two challenges, with duals $\widehat{\phi}_1$ and $\widehat{\phi}_2$. The usual SQI special-soundness composition forms an endomorphism of the public key curve by composing one challenge path, the two response paths, and a dual challenge path. Up to the orientation convention of the protocol, the extracted endomorphism has the form

$$\alpha = \widehat{\phi}_2 \circ \sigma_2 \circ \widehat{\sigma}_1 \circ \phi_1 \in \mathrm{End}(E_{\mathsf{pk}}).$$

All four maps have efficient evaluators, and the admissible representation interface supports composition and duals. Therefore the extractor computes an efficient evaluator for α.

It remains to show that α is not scalar. This is precisely the scalar-exclusion condition proved in the SQIsign2D–West special-soundness analysis. The argument is degree-theoretic. If $\alpha = [m]$ were scalar, then after moving the challenge factors to the two sides and cancelling the common scalar part, the equality of cyclic isogeny components would force the cyclic component induced by the difference of the two challenges to have degree at most the response bound. The accepting response bound gives

$$\deg(\nu) \le 2^{2e_{\mathsf{rsp}}}$$

for the relevant cyclic response component, while distinct challenges in the SQIsign2D–West challenge space force the corresponding cyclic challenge-separation component to have degree

$$\deg(\psi) > 2^{2(e-e_{\mathsf{chl}})}.$$

Because $e_{\mathsf{chl}} + e_{\mathsf{rsp}} \le e$, we have

$$2^{2(e-e_{\mathsf{chl}})} \ge 2^{2e_{\mathsf{rsp}}},$$

so the two degree constraints are incompatible. Hence $\alpha \notin \mathbb{Z}$. The output is therefore an efficient evaluator for a non-scalar endomorphism of E_{pk}, as required. □

7.2 Simulation Hints and Min-Entropy

SQIsign2D–West proves honest-verifier zero knowledge using isogeny-sampling oracles such as UTO/FIDIO. For the non-interactive proof, we package the oracle answers into simulation-shaped hints. This does not assert an efficient public sampler for these hints; it records the exact auxiliary information needed by the simulator.

Definition 18 (SQIsign2D–West simulation hints). *For a public key E_{pk}, the distribution $\mathcal{H}^{2D\text{-}W}_{\mathsf{sim},E_{\mathsf{pk}}}$ samples the ordered list of oracle outputs used by the SQIsign2D–West honest-verifier simulator to generate one accepting transcript. The hint contains the sampled challenge, the represented response isogeny, the auxiliary isogenies and decomposition data used to format the two-dimensional response, and the torsion-image data required by the verifier. A simulator $\mathsf{Sim}^{2D\text{-}W}_{\mathsf{sim}}$ reads this list and executes the deterministic transcript-formatting algorithm of SQIsign2D–West.*

Lemma 3 (Hint-assisted HVZK for SQIsign2D–West). *Let s be the number of simulated signing transcripts. If the SQIsign2D–West oracle-assisted simulator has statistical distance $\varepsilon^{\mathsf{zk}}_{2D\text{-}W}(s)$ from the honest transcript distribution, then $\mathsf{Sim}^{2D\text{-}W}_{\mathsf{sim}}$ with hints from $\mathcal{H}^{2D\text{-}W}_{\mathsf{sim}}$ has the same statistical distance. In the parameter regime of SQIsign2D–West, the bound is negligible under the stated simulator conditions.*

Proof. By Definition 18, a hint is the complete ordered list of oracle outputs used by one run of the SQIsign2D–West simulator. The hint-assisted simulator replays the same deterministic computation and consumes those outputs at the same points at which the oracle-assisted simulator would receive them. For every fixed choice of simulator coins and oracle coins, the transcript produced by the replay is identical to the transcript produced by the oracle-assisted simulator. Therefore the two simulator distributions are identical. The distance from the honest transcript distribution is exactly the distance proved for the oracle-assisted simulator, namely $\varepsilon^{\mathsf{zk}}_{2D\text{-}W}(s)$. □

Lemma 4 (Commitment min-entropy for SQIsign2D–West). *Assume the SQIsign2D–West commitment distribution is at statistical distance at most δ_{com} from the uniform distribution on supersingular j-invariants over $\mathbb{F}_{p^2}$. Then*

$$\mathsf{MEnt}(\Sigma_{2D\text{-}W}) \le \delta_{\mathsf{com}} + \frac{12}{p-1}.$$

In particular, under the SQIsign2D–West condition $N_{\mathsf{com}} \ge 2^{4\lambda}$, this quantity is negligible for $p = \Theta(2^{2\lambda})$.

Proof. Let U be the uniform distribution on supersingular j-invariants over $\mathbb{F}_{p^2}$, and let D be the actual commitment distribution. For every commitment value c, statistical distance gives

$$\Pr[D = c] \le \Pr[U = c] + \delta_{\mathsf{com}}.$$

The number of supersingular j-invariants over $\mathbb{F}_{p^2}$ is at least $(p-1)/12$, so $\Pr[U = c] \le 12/(p-1)$. Taking the maximum over c gives the displayed bound. The final statement follows from the SQIsign2D–West commitment-mixing condition and from the chosen asymptotic size of p. □

7.3 The Exact Fiat–Shamir Theorem

The theorem below assembles the SQIsign2D–West proof ingredients in the order used by the generic hint framework. Accepted SQIsign2D–West responses are recoverable elliptic response maps by Proposition 7; two accepting transcripts yield a non-scalar elliptic endomorphism by Theorem 6; the UTO/FIDIO simulator is converted into $\mathcal{H}^{2D\text{-}W}_{\mathsf{sim}}$-hint simulation by Lemma 3; and the commitment distribution contributes the min-entropy loss of Lemma 4. The first security statement is deliberately phrased under simulation-shaped HintOneEnd, because those are the hints that reproduce the signing simulation exactly.

Theorem 7 (EUF-CMA security of SQIsign2D–West with simulation hints). *Assume that $\Sigma_{2D\text{-}W}$ satisfies the SQIsign2D–West correctness conditions, that*

$$e_{\mathsf{chl}} + e_{\mathsf{rsp}} \leq e,$$

that its challenge space has size $|\mathcal{C}_{2D\text{-}W}|$, and that $\mathsf{MEnt}(\Sigma_{2D\text{-}W}) \leq \mu_{2D\text{-}W}$. Let $\mathsf{SIG}[\Sigma_{2D\text{-}W}]$ be the uncompressed Fiat–Shamir signature obtained from the identification protocol. For every random-oracle EUF-CMA adversary $\mathcal{A}$ making at most Q_{RO} random-oracle queries and s signing queries, there is a polynomial-time algorithm $\mathcal{B}$ such that

$$\begin{aligned}\mathrm{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{SIG}[\Sigma_{2D\text{-}W}]}(\mathcal{A}) \leq (Q_{\mathsf{RO}}+1)\cdot\Big(&\mathrm{Adv}^{\mathsf{HintOneEnd}^{\mathcal{H}^{2D\text{-}W}_{\mathsf{sim}}}_{2D\text{-}W}}(\mathcal{B}) + \varepsilon^{\mathsf{zk}}_{2D\text{-}W}(s) + |\mathcal{C}_{2D\text{-}W}|^{-1}\Big)\\ &+ (Q_{\mathsf{RO}}+s+1)\cdot s\cdot\mu_{2D\text{-}W}.\end{aligned}$$

Proof. We verify the hypotheses of Theorem 4. By Lemma 3, the hint-assisted simulator using $\mathcal{H}^{2D\text{-}W}_{\mathsf{sim}}$ simulates s signing transcripts with error $\varepsilon^{\mathsf{zk}}_{2D\text{-}W}(s)$. By Theorem 6, two accepting transcripts with the same commitment and distinct challenges yield an efficient evaluator for a non-scalar endomorphism of E_{pk}. This is a witness for $R^{2D\text{-}W}_{\mathsf{OneEnd}}$, hence for the hinted relation. By assumption and Lemma 4, commitment hits are bounded by $\mu_{2D\text{-}W}$. Substituting these three ingredients into Theorem 4 gives the displayed inequality. □

7.4 Pushable-Hint and Commitment-Recovery Corollaries

Simulation-shaped hints are the right object for answering signing queries, but they are not necessarily the right object for EndRing-compatible transport. The corollary below changes the hint distribution to a pushable uniform family. This change is not a theorem about equality of distributions; it is a cryptographic replacement step measured by the explicit $2D$-West HintDist advantage. Commitment recovery is handled separately, because omitting the commitment from the signature is a compression transformation applied after the standard Fiat–Shamir-with-hints theorem has been instantiated.

Definition 19 (SQIsign2D–West pushable uniform hints). *A SQIsign2D–West pushable uniform hint family $\mathcal{H}^{2D\text{-}W}_{\mathsf{unif}}$ is a source-local distribution of represented response and auxiliary isogeny data satisfying two properties. First, it is pushable in the sense of Definition 16 through the self-reduction isogenies used to relate* EndRing *and* OneEnd*. Second, it supports the deterministic transcript-formatting algorithm of SQIsign2D–West after the same public transcript fields are chosen. The concrete distribution is required to be stated with its own $2D\text{-}W$-HintDist assumption when replacing $\mathcal{H}^{2D\text{-}W}_{\mathsf{sim}}$.*

Corollary 4 (EUF-CMA security with pushable SQIsign2D–West hints). *Let $\mathcal{H}^{2D\text{-}W}_{\mathsf{unif}}$ be a SQIsign2D–West pushable uniform hint family, and suppose that $\mathcal{H}^{2D\text{-}W}_{\mathsf{sim}}$ and $\mathcal{H}^{2D\text{-}W}_{\mathsf{unif}}$ are s-hint indistinguishable with advantage*

$$\mathrm{Adv}^{s\text{-}2D\text{-}W\text{-}\mathsf{HintDist}}(\mathcal{D}).$$

Under the hypotheses of Theorem 7, for every adversary $\mathcal{A}$ making at most Q_{RO} random-oracle queries and s signing queries, there are polynomial-time algorithms $\mathcal{B}$ and $\mathcal{D}$ such that

$$\begin{aligned}\mathrm{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{SIG}[\Sigma_{2D\text{-}W}]}(\mathcal{A}) \leq (Q_{\mathsf{RO}}+1)\cdot\Big(&\mathrm{Adv}^{\mathsf{HintOneEnd}^{\mathcal{H}^{2D\text{-}W}_{\mathsf{unif}}}_{2D\text{-}W}}(\mathcal{B}) + \mathrm{Adv}^{s\text{-}2D\text{-}W\text{-}\mathsf{HintDist}}(\mathcal{D})\\ &+ \varepsilon^{\mathsf{zk}}_{2D\text{-}W}(s) + |\mathcal{C}_{2D\text{-}W}|^{-1}\Big)\\ &+ (Q_{\mathsf{RO}}+s+1)\cdot s\cdot\mu_{2D\text{-}W}.\end{aligned}$$

If $\mathcal{H}^{2D\text{-}W}_{\mathsf{unif}}$ is pushable through the relevant EndRing*-to-*OneEnd *self-reductions, the resulting hardness statement is EndRing-compatible.*

Proof. Start from the simulated signing experiment in the proof of Theorem 7, where the signing simulations use s independently sampled $\mathcal{H}^{2D\text{-}W}_{\mathsf{sim}}$-hints. Replace these hints by s independently

sampled $\mathcal{H}_{\mathsf{unif}}^{2D\text{-}W}$-hints, leaving the adversary, random-oracle programming, and transcript-formatting algorithms unchanged. If the adversary's success probability changes by more than

$$\mathrm{Adv}^{s\text{-}2D\text{-}W\text{-}\mathsf{HintDist}}(\mathcal{D}),$$

then the standard hybrid that guesses the first changed signing transcript and gives the adversary's final success bit distinguishes the two s-sample hint distributions with the same advantage. Thus the replacement changes the success probability by at most the stated HintDist term.

After the replacement, the forking argument is identical to Theorem 7, except that the extracted relation uses $\mathcal{H}_{\mathsf{unif}}^{2D\text{-}W}$-hints. Applying Theorem 4 with the uniform hint family gives the same challenge-guessing and commitment-min-entropy terms. The conservative displayed bound charges the HintDist replacement inside the same $(Q_{\mathsf{RO}} + 1)$ factor as the hinted-relation and simulation terms. Pushability through the self-reduction class is exactly the hypothesis used by Proposition 6; hence the final hardness term is EndRing-compatible. □

Corollary 5 (Commitment-recoverable SQIsign2D–West signatures). *Assume the SQIsign2D–West commitment-recovery algorithm is correct and perfectly sound: every honestly generated transcript recovers the omitted commitment, and every recovered commitment together with the challenge and response is a valid transcript commitment. Then the EUF-CMA bounds of Theorem 7 and Corollary 4 also hold for the commitment-recoverable SQIsign2D–West signature format, with no additional loss beyond the commitment-recovery correctness and soundness errors if these are statistical rather than perfect.*

Proof. Given a forgery for the compressed format, the reduction runs the public commitment-recovery algorithm on the forged challenge and response. Perfect soundness ensures that the recovered commitment makes the tuple an accepting uncompressed Fiat–Shamir transcript whenever the compressed verification accepts. Thus every compressed forgery gives an uncompressed forgery for the same message and public key. If recovery correctness or soundness has a nonzero error probability, a union bound adds exactly that error term. Otherwise the loss is zero. □

8 The SQIPrime2D AIO Boundary

The second theorem package will repair the auxiliary layer of SQIPrime2D. Before defining that repair, this section isolates the obstruction in the published proof model. The issue is not elliptic extraction: accepted SQIPrime responses still recover elliptic isogenies. The issue is the auxiliary-isogeny oracle used by the two-dimensional variant. Published SQIPrime2D defines AIO through the distribution of an auxiliary isogeny produced by a protocol-specific Kani and endomorphism construction [DF24]. Such a distribution is not automatically curve-local or pushable.

8.1 Published SQIPrime2D Response Shape

A SQIPrime public key is a supersingular elliptic curve E_A, equipped with the public torsion data used to define a non-smooth challenge. In the clean coprime case considered here, a challenge determines a cyclic subgroup

$$L \subset E_A[q]$$

and the corresponding challenge isogeny

$$\phi : E_A \to E_2$$

of degree q. A response isogeny factors as

$$\kappa = \sigma \circ \phi : E_A \to E_1, \qquad \deg(\sigma) = d, \qquad q \nmid d,$$

so $\deg(\kappa) = q \cdot d$ and $L \subseteq \ker(\kappa)$. The two-dimensional variant also needs an auxiliary isogeny

$$\delta : E_A \to E_\delta, \qquad \deg(\delta) = 2^\alpha - q \cdot d,$$

so that the pair (κ, δ) satisfies the degree relation used by the Kani-style dimension-two verifier:

$$\deg(\kappa) + \deg(\delta) = 2^\alpha.$$

When d is odd and $q \nmid d$, the two degrees are coprime because

$$\gcd(q \cdot d, 2^\alpha - q \cdot d) = \gcd(q \cdot d, 2^\alpha) = 1.$$

The special-soundness object remains elliptic. Published SQIPrime proves that two accepting SQIPrime2D transcripts with the same commitment and distinct challenges yield a non-trivial endomorphism of the elliptic public-key curve E_A [DF24, Proposition 5]. The proof uses the response isogenies recovered by verification and forms the elliptic composition $\widehat{\kappa}_2 \circ \kappa_1$. Therefore the role of higher-dimensional data in SQIPrime2D is representational: it supplies an efficient verification format for elliptic maps, while the special-soundness witness is an elliptic OneEnd witness.

Proposition 8 (SQIPrime accepted responses are elliptic at extraction). *For accepted SQIPrime2D transcripts, the special-soundness extractor outputs an efficient evaluator for a non-scalar endomorphism of the public elliptic curve E_A. In the clean coprime case this follows from the q-primary kernel separation described above; in the full SQIPrime2D degree range it is exactly the scalar-exclusion theorem of SQIPrime Proposition 5.*

Proof. The statement is an application of SQIPrime's special-soundness theorem for SQIPrime2D. SQIPrime Proposition 5 states that, given two accepting SQIPrime2D transcripts

$$(E_1, \mathsf{chal}_1, T_1, U_1, V_1) \quad \text{and} \quad (E_1, \mathsf{chal}_2, T_2, U_2, V_2)$$

with the same commitment E_1 and distinct challenges $\mathsf{chal}_1 \neq \mathsf{chal}_2$, there is an efficient extractor that solves the one-endomorphism problem over E_A. In the notation of that proposition, verification recovers response isogenies

$$\kappa_i : E_A \longrightarrow E_1$$

and the extractor forms

$$\theta_A = \widehat{\kappa}_2 \circ \kappa_1 \in \mathrm{End}(E_A).$$

The proposition proves that θ_A is non-scalar. Its proof covers both the clean case in which the residual degrees are prime to q and the SQIPrime2D case in which one residual degree is divisible by q. In the latter case, the argument uses the constraints $q \cdot d_i < 2^\alpha$ and $q^3 > 2^\alpha$ to show that divisibility by q occurs to exactly the extent handled by the right q-primary factors; scalarity would force the two right q-primary factors to have the same kernel on $E_A[q]$, hence would force the two challenge lines to coincide, contradicting $\mathsf{chal}_1 \neq \mathsf{chal}_2$.

The present proposition uses exactly the same public objects and the same extracted endomorphism. The higher-dimensional or Kani-style data in an accepted transcript supplies an efficient representation of each elliptic response map κ_i, together with the evaluator and dual-evaluator interface needed to compute $\widehat{\kappa}_2 \circ \kappa_1$. Therefore the extractor described by SQIPrime Proposition 5 is directly applicable. Its output is an efficient evaluator for a non-scalar endomorphism of the elliptic public curve E_A, which is a valid OneEnd witness. No native higher-dimensional endomorphism is extracted. □

8.2 Why Published AIO Is Not Curve-Local

The auxiliary isogeny δ in published SQIPrime2D is not defined as a uniform curve-local sample from degree $2^\alpha - q \cdot d$ isogenies out of E_A. It is generated through a protocol-specific construction using a hidden context: the base curve, the secret path to E_A, ideals or orders representing that path, and randomized endomorphism/Kani-decomposition data. The AIO oracle is therefore best read as a sampler for "the distribution of the auxiliary isogeny produced by the response algorithm", not as a distribution determined only by $(E_A, 2^\alpha - q \cdot d)$.

This distinction is exactly what matters for hinted EndRing reductions. A pushable hint distribution must survive transport from one curve to another. If the target distribution is not defined by the target curve alone, then there is no well-defined distribution to push into without specifying how the hidden protocol context is also transported.

Definition 20 (Algorithm-context auxiliary distribution). *An SQIPrime2D auxiliary context is a tuple*

$$\mathsf{Ctx} = (E_0, E_A, \tau, \mathcal{O}_A, \mathsf{aux\text{-}coins})$$

containing the base curve, the public-key curve, the hidden path $\tau : E_0 \to E_A$, the order/ideal information used by the response algorithm, and the randomness of the auxiliary Kani and endomorphism routines. For an admissible residual degree d, write

$$\mathcal{D}^{\mathsf{Alg}}_{\mathsf{Ctx},d}$$

for the distribution of the auxiliary isogeny $\delta : E_A \to E_\delta$ of degree $2^\alpha - q \cdot d$ output by the published auxiliary-generation algorithm in that context.

Definition 21 (Curve-local AIO). *A curve-local AIO distribution for SQIPrime2D is a family*

$$\mathcal{D}^{\mathsf{AIO}}_{E,d}$$

depending only on the visible curve E and the residual degree d, such that for every admissible auxiliary context Ctx with public-key curve $E_A = E$, the distribution $\mathcal{D}^{\mathsf{Alg}}_{\mathsf{Ctx},d}$ is equal to, or computationally indistinguishable from, $\mathcal{D}^{\mathsf{AIO}}_{E,d}$.

8.3 A Formal Context-Independence Obstruction

The obstruction is already visible before considering any nontrivial transport. If an auxiliary distribution is claimed to be curve-local and pushable, then pushing through the identity isogeny on a visible curve E must leave the target distribution unchanged. Consequently all hidden Algorithm-contexts that realize the same visible input (E, d) must induce the same auxiliary distribution. This identity-transport argument is what the proposition formalizes.

Proposition 9 (AIO pushability requires context independence). *Suppose that the published SQIPrime2D auxiliary distribution admits an exact curve-local pushability theorem: for every admissible context Ctx over E, every transport isogeny $\rho : E \to F$ in the claimed pushability class, and every residual degree d, the pushforward of*

$$\delta \leftarrow \mathcal{D}^{\mathsf{Alg}}_{\mathsf{Ctx},d}$$

is distributed as a target distribution depending only on (F, d). Then the published auxiliary distribution satisfies context independence: for any two admissible contexts $\mathsf{Ctx}_1, \mathsf{Ctx}_2$ with the same visible curve E, the distributions

$$\mathcal{D}^{\mathsf{Alg}}_{\mathsf{Ctx}_1,d} \quad \text{and} \quad \mathcal{D}^{\mathsf{Alg}}_{\mathsf{Ctx}_2,d}$$

are equal, or computationally indistinguishable in the computational version.

Proof. Apply the claimed pushability theorem to the identity isogeny $\rho = \mathrm{id}_E$, which belongs to any reasonable pushability class and is prime to every admissible degree. Pushing through the identity does not change an auxiliary isogeny, so the output distribution is exactly $\mathcal{D}^{\mathsf{Alg}}_{\mathsf{Ctx},d}$. By the hypothesis, however, the pushed distribution must be the curve-local target distribution depending only on (E, d). Therefore every context Ctx over the same visible curve E induces the same distribution, namely that curve-local target distribution. This proves exact context independence. In the computational version, the same identity-push experiment gives a distinguisher reduction: any algorithm distinguishing the two context-conditioned distributions also distinguishes at least one of them from the asserted curve-local distribution. □

Corollary 6 (Published AIO needs an additional distributional claim). *The RUCO-DIO+AIO formulation of published SQIPrime2D does not by itself imply an EndRing-compatible pushable hint distribution for the auxiliary layer. Such a conclusion requires an additional statement asserting that the Algorithm-context auxiliary distributions are curve-local and pushable, or a replacement distribution that is curve-local by definition.*

Proof. AIO supplies the simulator with auxiliary isogenies having the distribution of the published auxiliary-generation algorithm. By Proposition 9, a curve-local pushability theorem for that distribution would imply context independence across hidden auxiliary contexts. The AIO model states the oracle distribution needed for simulation; it does not prove that all hidden contexts with the same visible curve induce the same auxiliary distribution, nor that this distribution is preserved by pushforward. Therefore a pushable EndRing-compatible hint theorem needs an extra distributional claim or a new curve-local auxiliary distribution. The next section follows the second route by replacing the Algorithm-context auxiliary sampler with UAux. □

9 SQIPrime2D-UAux: Context-Independent Auxiliary Hints

The preceding section isolates the auxiliary-isogeny oracle as the distributional component that prevents the published SQIPrime2D proof model from immediately yielding a pushable hint family. This section gives a constructive repair of that component. We replace the Algorithm-context AIO sampler by a curve-local uniform auxiliary distribution. The resulting variant, denoted $\mathsf{SQIPrime2D\text{-}UAux}$, keeps the SQIPrime2D challenge and response relation, and keeps the dimension-two Kani verification equation, but changes how the auxiliary isogeny is distributed.

The construction is a foundations variant rather than an efficiency claim. Sampling uniformly from the auxiliary distribution may require endomorphism-ring information and ideal-to-isogeny translation machinery. All security statements involving UAux are therefore interpreted either with a polynomial-time represented sampler for the stated distribution, or as ideal-hint experiments in which the represented sampler is part of the hinted hardness assumption. The point is distributional: the new auxiliary layer is defined only from the visible curve and the auxiliary degree. Consequently it is context-independent and exactly pushable.

9.1 UAux Sampling Models

The notation $\mathsf{UAux}(E, n)$ names a distribution, not automatically an efficient public algorithm. This distinction is important for the SQIPrime-side claims. Uniformly sampling represented cyclic degree-n auxiliary isogenies from a visible curve may require endomorphism-ring information or ideal-to-isogeny translation machinery. The paper therefore separates two models.

In the *ideal-hint model*, $\mathsf{UAux}(E, n)$ is sampled by the hinted hardness experiment. The sampled hint must be polynomial-size and verifier-facing, but the theorem does not claim that an external public sampler can produce it. This is the model used for the cleanest foundations statements.

In the *samplable-instantiation model*, a concrete represented sampler $\mathsf{SampUAux}$ is supplied. For every admissible (E, n), let

$$\Delta_{\mathsf{UAux}}(E, n)$$

be the statistical or computational distance between the output distribution of $\mathsf{SampUAux}(E, n)$ and $\mathsf{UAux}(E, n)$ in the representation format used by the protocol. We write

$$\varepsilon_{\mathsf{UAux}} = \max_{(E,n)} \Delta_{\mathsf{UAux}}(E, n)$$

over the admissible public instances and auxiliary degrees.

Proposition 10 (Sampling-loss accounting for UAux hints). *Any theorem in this paper stated in the ideal-hint model with at most s independently sampled UAux-containing hints remains valid with a concrete sampler $\mathsf{SampUAux}$, except that the success bound gains an additive term at most*

$$s \cdot \varepsilon_{\mathsf{UAux}}$$

in the underlying hint experiment. When the theorem is inserted under the conservative Fiat–Shamir-with-hints bound of Theorem 4, this loss may be charged as

$$(Q_{\mathsf{RO}} + 1) \cdot s \cdot \varepsilon_{\mathsf{UAux}}.$$

Proof. Consider the sequence of the at most s UAux-containing hints used by the hinted hardness or signing-simulation experiment. Construct hybrids $G_0, \ldots, G_s$, where G_i samples the first i UAux components using $\mathsf{SampUAux}$ and the remaining components using the ideal UAux distribution, with all non-UAux fields sampled identically in all games. For each adjacent pair (G_{i-1}, G_i), the only changed random variable is one represented UAux sample for some admissible pair (E, n). By definition of $\varepsilon_{\mathsf{UAux}}$, replacing this component changes the success probability of any distinguisher by at most $\varepsilon_{\mathsf{UAux}}$. Summing over the s hybrids gives a total change at most $s \cdot \varepsilon_{\mathsf{UAux}}$. If the corresponding EUF-CMA theorem places the hint-distribution replacement inside a $(Q_{\mathsf{RO}}+1)$ forking loss, the same conservative accounting multiplies the sampler error by $(Q_{\mathsf{RO}} + 1)$, yielding the second displayed term. □

9.2 Uniform Auxiliary Isogenies

Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve and let $n \geq 1$ be prime to p. Write

$$\mathsf{Cyc}_n(E)$$

for the finite set of cyclic finite subgroup schemes $K \subset E[n]$ of order n, over the same algebraic closure and field-of-definition convention used by the representation interface. Since $p \nmid n$, these subgroup schemes are etale, and each $K \in \mathsf{Cyc}_n(E)$ determines a separable cyclic quotient isogeny

$$\delta_K : E \longrightarrow E/K$$

of degree n. Two quotient maps with the same kernel differ only by post-composition with a target isomorphism. All distributions below identify such representatives, because the verifier-facing data in the paper's representation interface is invariant under this target-isomorphism ambiguity.

Definition 22 (Uniform auxiliary sampler). *For $p \nmid n$, the distribution*

$$\mathsf{UAux}(E, n)$$

is obtained by sampling $K \leftarrow \mathsf{Cyc}_n(E)$ uniformly and outputting an admissible efficient representation of the quotient isogeny

$$\delta_K : E \to E/K.$$

Equivalently, it is the uniform distribution on primitive cyclic degree-n isogenies from E, modulo target isomorphism, expressed in the chosen representation format. The definition is curve-local: its visible input is exactly (E, n).

Theorem 8 (Context independence of UAux). *Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve and let n be prime to p. The distribution $\mathsf{UAux}(E, n)$ depends only on the visible pair (E, n). In particular, if two SQIPrime-style hidden contexts contain different paths*

$$\tau_0 : E_0 \to E, \qquad \tau_1 : E_0 \to E$$

or different ideal and Kani-decomposition data leading to the same visible curve E, then the auxiliary distribution obtained by replacing the auxiliary sampler with $\mathsf{UAux}(E, n)$ is identical in the two contexts.

Proof. By Definition 22, the sample space is $\mathsf{Cyc}_n(E)$, with the uniform probability measure. The map from a sampled subgroup K to the represented quotient isogeny $\delta_K : E \to E/K$ is determined by E, n, and the fixed representation convention. None of these objects contains a path from E_0 to E, an ideal representing such a path, an endomorphism used to generate an auxiliary map, or randomness from KaniDoublePath or ExtKaniDoublePath. Therefore changing the hidden context while keeping (E, n) fixed leaves the sample space, the probability of every sample, and the quotient construction unchanged. Hence $\mathsf{UAux}(E, n)$ is a function only of (E, n). □

Theorem 9 (Pushability of uniform auxiliary isogenies). *Let $E, F/\mathbb{F}_{p^2}$ be supersingular elliptic curves, let $\rho : E \to F$ be a separable isogeny, and let n be prime to $p \cdot \deg(\rho)$. If*

$$\delta : E \to E'$$

is sampled from $\mathsf{UAux}(E, n)$, then the pushforward $\rho_*\delta : F \to F'$ is distributed exactly as $\mathsf{UAux}(F, n)$.

Proof. Let $m = \deg(\rho)$. Since $\gcd(m, n) = 1$, multiplication by m is an automorphism of $E[n]$. The dual isogeny satisfies

$$\widehat{\rho} \circ \rho = [m]_E.$$

Restricting to $E[n]$, if $P \in E[n]$ and $\rho(P) = 0$, then

$$[m]P = (\widehat{\rho} \circ \rho)(P) = 0.$$

Because $[m]$ is invertible on $E[n]$, this implies $P = 0$. Thus $\rho|_{E[n]}$ is injective. Since $p \nmid n$, both $E[n]$ and $F[n]$ have order n^2, so the injective map

$$\rho|_{E[n]} : E[n] \to F[n]$$

is an isomorphism. An explicit inverse is $[m]^{-1} \cdot \widehat{\rho}|_{F[n]}$, where $[m]^{-1}$ denotes the inverse of multiplication by m on $E[n]$.

Consequently,

$$K \longmapsto \rho(K)$$

is a bijection from $\mathsf{Cyc}_n(E)$ to $\mathsf{Cyc}_n(F)$. For $K \in \mathsf{Cyc}_n(E)$, the quotient map $\delta_K : E \to E/K$ pushes forward along ρ to the quotient map $F \to F/\rho(K)$, up to the canonical target isomorphism induced by the universal property of quotients. Equivalently, there is a commutative square

$$\begin{array}{ccc} E & \xrightarrow{\delta_K} & E/K \\ \rho\downarrow & & \downarrow \rho_K \\ F & \xrightarrow{\delta_{\rho(K)}} & F/\rho(K), \end{array}$$

and the lower horizontal map is the pushed auxiliary isogeny. Since uniform measure is preserved by a bijection, $\rho(K)$ is uniform in $\mathsf{Cyc}_n(F)$ when K is uniform in $\mathsf{Cyc}_n(E)$. Therefore $\rho_* \delta_K$ has exactly the distribution $\mathsf{UAux}(F, n)$. □

9.3 The SQIPrime2D-UAux Variant

Let q be the SQIPrime challenge degree. In the clean coprime setting used in this paper, a response map factors as

$$\kappa = \sigma \circ \phi : E_A \to E_1, \qquad \deg(\phi) = q, \qquad \deg(\sigma) = d,$$

where d is odd and $q \nmid d$. The auxiliary degree is

$$n_{\mathsf{aux}} = 2^\alpha - q \cdot d.$$

We require

$$q \cdot d < 2^\alpha, \qquad p \nmid n_{\mathsf{aux}}.$$

Since $q \cdot d$ is odd, n_{aux} is odd, and

$$\gcd(q \cdot d, n_{\mathsf{aux}}) = \gcd(q \cdot d, 2^\alpha) = 1.$$

Definition 23 (SQIPrime2D-UAux)**.** SQIPrime2D-UAux *is the SQIPrime2D identification protocol with the same public-key distribution, challenge family, constrained response relation, and dimension-two verifier relation as SQIPrime2D, except that the auxiliary isogeny accompanying the response map $\kappa : E_A \to E_1$ is sampled as*

$$\delta \leftarrow \mathsf{UAux}(E_A, 2^\alpha - q \cdot d)$$

instead of being sampled from the Algorithm-context auxiliary distribution. The response representation contains the evaluation data and Kani metadata required to verify the pair (κ, δ).

Theorem 10 (Preservation of the dimension-two Kani degree relation). *Let $\kappa : E_A \to E_1$ be a SQIPrime2D response isogeny of degree $q \cdot d$, where d is odd, $q \nmid d$, $q \cdot d < 2^\alpha$, and $p \nmid 2^\alpha - q \cdot d$. Let*

$$\delta : E_A \to E_\delta$$

be sampled from $\mathsf{UAux}(E_A, 2^\alpha - q \cdot d)$*. Then*

$$\deg(\kappa) + \deg(\delta) = 2^\alpha$$

and $\gcd(\deg(\kappa), \deg(\delta)) = 1$*. Thus (κ, δ) satisfies the degree and coprimality hypotheses of the dimension-two Kani verification relation. Acceptance additionally requires the supplied evaluation data and Kani metadata to certify the represented maps.*

Proof. The response degree is $\deg(\kappa) = q \cdot d$. By Definitions 22 and 23, the auxiliary degree is

$$\deg(\delta) = 2^\alpha - q \cdot d.$$

Adding the two identities gives $\deg(\kappa) + \deg(\delta) = 2^\alpha$. Since $q \cdot d$ is odd, it is coprime to 2^α, and therefore

$$\gcd(q \cdot d, 2^\alpha - q \cdot d) = \gcd(q \cdot d, 2^\alpha) = 1.$$

These are exactly the numerical hypotheses used to form the two-dimensional Kani relation in the clean SQIPrime2D response shape. The theorem does not assert that degree data alone is sufficient for verification; the verifier also checks the represented evaluations and metadata supplied with (κ, δ). The statement is that replacing the auxiliary sampler by UAux preserves the algebraic degree conditions of the original verifier relation. □

Proposition 11 (AIO obstruction and UAux replacement). *The published SQIPrime2D AIO distribution is not, by definition, a curve-local uniform auxiliary distribution. A pushable hint theorem for the published AIO requires an additional auxiliary-uniformity or auxiliary-pushability assumption. In* SQIPrime2D-UAux*, the auxiliary distribution is curve-local, context-independent, and exactly pushable by Theorems 8 and 9.*

Proof. Published AIO is specified by reference to the auxiliary isogeny produced by the SQIPrime2D response algorithm. That algorithm uses hidden context: the setup curve, the secret path to E_A, endomorphism or ideal data, and the Kani construction used to produce the auxiliary map. Therefore equality between the published AIO distribution at a curve E and the pushforward of an AIO distribution from an isogenous curve is not a consequence of the visible pair (E, n) alone. It requires a separate claim that the context-conditioned auxiliary distributions are indistinguishable under transport, or that they coincide with a curve-local uniform distribution.

By contrast, $\mathsf{UAux}(E, n)$ is defined directly as the uniform distribution on cyclic order-n kernels of $E[n]$. Theorem 8 proves that this distribution depends only on (E, n), and Theorem 9 proves exact preservation under pushforward through every isogeny of degree coprime to $p \cdot n$. Thus SQIPrime2D-UAux gives a context-independent and exactly pushable replacement for the auxiliary AIO layer. □

10 Fully Pushable Challenge-Carrying Hints for SQIPrime2D-UAux

The UAux replacement of Section 9 repairs the auxiliary layer, but it does not by itself make the whole SQIPrime2D hint distribution EndRing-compatible. A hint of the form (L, κ, δ), where $L \subset E[q]$ and $L \subseteq \ker(\kappa)$, is not fully pushable through the q-challenge quotient: the quotient map kills L, so its action on $E[q]$ is not invertible. The full repair is to carry the challenge isogeny itself in the pushable hint. The resulting distribution is stronger than the UAux-only hint, but it has the key property needed for hinted EndRing reductions: it is transported by bijections along smooth self-reduction isogenies.

10.1 Challenge-Carrying Uniform Hints

A levelled public curve is a tuple

$$\mathbf{E} = (E, P, Q, \mathcal{A}),$$

where $E/\mathbb{F}_{p^2}$ is supersingular, (P, Q) is a public basis of $E[q]$, and $\mathcal{A}$ is the challenge index set. For each $a \in \mathcal{A}$, the public rule defines a cyclic challenge line

$$L_a \subset E[q].$$

In the affine form one may take $L_a = \langle P + [a]Q \rangle$; a projective challenge space is handled identically by taking $a \in \mathbb{P}^1(\mathbb{Z}/q\mathbb{Z})$. Let

$$\phi_a : E \to F_a$$

be the degree-q quotient isogeny with kernel L_a.

Let $\mathcal{D}_{\mathsf{res}}$ be a residual-degree distribution supported on integers d satisfying

$$d \text{ odd}, \qquad q \nmid d, \qquad q \cdot d < 2^\alpha, \qquad p \nmid (2^\alpha - q \cdot d).$$

For each curve F and each such d, the admissible residual response sample space is either the full set of cyclic degree-d quotient isogenies from F, modulo target isomorphism, or an explicitly declared subset that is stable under pushforward by transport isogenies of degree prime to d. The pushability theorem below uses exactly this stability hypothesis; when the full cyclic quotient space is used, stability follows from the same torsion-kernel bijection as Theorem 9.

Definition 24 (Challenge-carrying uniform hints). *For a levelled public curve* $\mathbf{E} = (E, P, Q, \mathcal{A})$*, the distribution*

$$\mathcal{H}^{\mathsf{SQIP2D\text{-}UAux}}_{q\text{-}\mathsf{unif},\mathbf{E}}$$

samples a hint as follows.

(1) Sample $a \leftarrow \mathcal{A}$ *according to the honest challenge distribution and let* $\phi_a : E \to F_a$ *be the quotient by* L_a*.*
(2) Sample $d \leftarrow \mathcal{D}_{\mathsf{res}}$*.*
(3) Sample an admissibly represented residual response isogeny

$$\sigma : F_a \to E_1$$

uniformly from the degree-d residual response sample space at F_a*.*
(4) Sample

$$\delta \leftarrow \mathsf{UAux}(E, 2^\alpha - q \cdot d).$$

(5) Output

$$h = (a, \phi_a, \sigma, d, \delta, \mathsf{EvalMeta}, \mathsf{KaniMeta}),$$

where EvalMeta *and* KaniMeta *are the verifier-facing metadata needed to represent* $\kappa = \sigma \circ \phi_a$ *together with* δ*.*

This distribution is challenge-carrying because it contains an efficient representation of ϕ_a. It is therefore stronger than the line-constrained hint (L_a, κ, δ). The gain is that all random components are now sampled from source-local finite sets whose transport by smooth isogenies is bijective.

10.2 Transcript Reconstruction

A challenge-carrying hint is not itself a transcript. It stores a decomposed object: the challenge isogeny ϕ_a, the residual response σ, and the auxiliary isogeny δ. The transcript response is obtained by composing the first two maps,

$$\kappa = \sigma \circ \phi_a.$$

The theorem below verifies that this decomposition produces the same kind of accepted response checked by SQIPrime2D-UAux. This is the point at which the challenge-carrying representation is connected back to the verifier's line-constrained response relation.

Theorem 11 (Transcript reconstruction from challenge-carrying hints). *Let*

$$h = (a, \phi_a, \sigma, d, \delta, \mathsf{EvalMeta}, \mathsf{KaniMeta}) \leftarrow \mathcal{H}^{\mathsf{SQIP2D\text{-}UAux}}_{q\text{-}\mathsf{unif},\mathbf{E}}.$$

Set

$$\kappa = \sigma \circ \phi_a : E \to E_1.$$

Then κ has degree $q \cdot d$, its kernel contains the challenge line L_a, and δ has degree $2^\alpha - q \cdot d$. Consequently, under valid evaluation and Kani metadata, the hint reconstructs a valid SQIPrime2D-UAux *response for challenge a.*

Proof. By construction, ϕ_a is the quotient isogeny with kernel L_a, so $\deg(\phi_a) = q$. The residual response isogeny σ has degree d. Since $q \nmid d$, the degrees are coprime, and

$$\deg(\kappa) = \deg(\sigma \circ \phi_a) = q \cdot d.$$

Moreover $L_a \subseteq \ker(\phi_a) \subseteq \ker(\sigma \circ \phi_a) = \ker(\kappa)$. By the definition of the hint, δ is sampled from $\mathsf{UAux}(E, 2^\alpha - q \cdot d)$, hence has degree $2^\alpha - q \cdot d$. Therefore

$$\deg(\kappa) + \deg(\delta) = q \cdot d + (2^\alpha - q \cdot d) = 2^\alpha.$$

The coprimality of $q \cdot d$ and $2^\alpha - q \cdot d$ follows as in Theorem 10. Thus the pair (κ, δ) satisfies the numerical hypotheses of the SQIPrime2D Kani relation. If $\mathsf{EvalMeta}$ and $\mathsf{KaniMeta}$ are valid, the verifier can evaluate the represented maps and check the relation; hence the hint reconstructs an accepting response for challenge a. □

10.3 Full Pushability

The pushability theorem is stated for smooth self-reduction isogenies whose degrees are coprime to the challenge degree, the residual degree, and the auxiliary degree. This is the correct class for EndRing-compatible self-reductions; the theorem is not a claim that the hint can be pushed through the degree-q challenge quotient itself.

The proof decomposes the transport into three bijections. First, the smooth transport is prime to q, so it maps the public challenge line $\langle P + [a]Q \rangle$ bijectively to the transported line $\langle P' + [a]Q' \rangle$. Second, the induced transport on the codomain of the challenge isogeny is prime to the residual degree d, so it pushes residual response isogenies bijectively. Third, the transport is prime to $2^\alpha - q \cdot d$, so UAux pushability applies to the auxiliary isogeny. Since the hint samples these components conditionally independently, preserving each conditional distribution preserves the joint distribution.

This theorem asserts pushability through the smooth self-reduction isogenies used by the hinted EndRing reduction, under the stated coprimality and level-transport hypotheses. It does not assert pushability through the degree-q challenge quotient itself.

Definition 25 (Transport of levelled curves). *Let $\mathbf{E} = (E, P, Q, \mathcal{A})$ and $\mathbf{E}' = (E', P', Q', \mathcal{A})$ be levelled public curves. A transport isogeny*

$$\rho : \mathbf{E} \to \mathbf{E}'$$

is a separable isogeny $\rho : E \to E'$ with $q \nmid \deg(\rho)$ such that

$$P' = \rho(P), \qquad Q' = \rho(Q).$$

Then ρ maps each labelled challenge line $L_a \subset E[q]$ to the corresponding line $L'_a \subset E'[q]$.

Theorem 12 (Full pushability of challenge-carrying uniform hints). *Let $\rho : \mathbf{E} \to \mathbf{E}'$ be a transport isogeny as in Definition 25. Assume that for every d in the support of $\mathcal{D}_{\mathsf{res}}$,*

$$\gcd(\deg(\rho), q \cdot d \cdot (2^\alpha - q \cdot d)) = 1.$$

Then there is a polynomial-time push algorithm $\mathsf{PushHint}_q$ such that, for

$$h \leftarrow \mathcal{H}^{\mathsf{SQIP2D\text{-}UAux}}_{q\text{-}\mathsf{unif},\mathbf{E}},$$

the pushed hint $\mathsf{PushHint}_q(\mathbf{E}, h, \rho)$ is distributed exactly as

$$\mathcal{H}^{\mathsf{SQIP2D\text{-}UAux}}_{q\text{-}\mathsf{unif},\mathbf{E}'}.$$

Proof. Let

$$h = (a, \phi_a, \sigma, d, \delta, \mathsf{EvalMeta}, \mathsf{KaniMeta})$$

be sampled from $\mathcal{H}^{\mathsf{SQIP2D\text{-}UAux}}_{q\text{-}\mathsf{unif},\mathbf{E}}$. The push algorithm leaves a and d unchanged and transports the three isogeny components.

First consider the challenge isogeny. Since $q \nmid \deg(\rho)$, the restriction $\rho : E[q] \to E'[q]$ is an isomorphism. By Definition 25, it maps P to P', Q to Q', and therefore maps the labelled line L_a to L'_a. Form the pushout square

$$\begin{array}{ccc} E & \xrightarrow{\phi_a} & F_a \\ \rho\downarrow & & \downarrow\rho_a \\ E' & \xrightarrow{\phi'_a} & F'_a. \end{array}$$

The lower horizontal map ϕ'_a is the degree-q quotient by L'_a. Since a is sampled with the honest challenge distribution and the label a is preserved, the transported challenge component is distributed as a fresh challenge component at $\mathbf{E}'$.

Second, condition on the sampled pair (a, d). The residual response $\sigma : F_a \to E_1$ is uniform in the admissible degree-d response sample space at F_a. The isogeny $\rho_a : F_a \to F'_a$ has the same degree as ρ, hence is prime to d. As in Theorem 9, restriction to d-torsion is an isomorphism. Transport of kernels along ρ_a gives a bijection from degree-d residual response isogenies out of F_a to degree-d residual response isogenies out of F'_a, preserving admissibility by the representation-interface hypothesis. Thus the pushed residual map $\sigma' : F'_a \to E'_1$ is uniform in the target residual sample space.

Third, the auxiliary degree is $n_{\mathsf{aux}} = 2^\alpha - q \cdot d$. By hypothesis n_{aux} is prime to $p \cdot \deg(\rho)$, so Theorem 9 applies and gives

$$\rho_* \mathsf{UAux}(E, n_{\mathsf{aux}}) = \mathsf{UAux}(E', n_{\mathsf{aux}})$$

as exact distributions. Hence the pushed auxiliary δ' is a fresh UAux sample at E'.

The three transported components are obtained by bijections conditioned on the same preserved (a, d). The original hint samples σ and δ independently conditioned on (a, d), and componentwise bijections preserve independence. The metadata is recomputed deterministically from the pushed isogeny representations. Therefore the joint pushed tuple

$$(a, \phi'_a, \sigma', d, \delta', \mathsf{EvalMeta}', \mathsf{KaniMeta}')$$

has exactly the distribution specified by Definition 24 for $\mathbf{E}'$. This proves exact pushability. □

Corollary 7 (EndRing-compatible challenge-carrying hints). *The hint family $\mathcal{H}^{\mathsf{SQIP2D\text{-}UAux}}_{q\text{-}\mathsf{unif}}$ is EndRing-compatible in the sense of Proposition 6 for any* EndRing*-to-*OneEnd *self-reduction whose transport isogenies satisfy the coprimality condition of Theorem 12 and whose representation interface supports the pushout and metadata recomputation operations used in the proof.*

Proof. Theorem 12 proves exact pushability of the hint distribution through the stated transport isogenies. Therefore the hypothesis of Proposition 6 is satisfied for that self-reduction class. The reduction can push all available challenge-carrying hints from the source curve to the self-reduced curves, invoke the hinted OneEnd solver with correctly distributed hints, and transport the resulting non-scalar endomorphisms back exactly as in the unhinted EndRing-to-OneEnd reduction. No distributional loss is incurred beyond the stated success loss of the underlying algebraic reduction. □

10.4 Why Challenge-Carrying Hints Are Not Free

The challenge-carrying distribution is deliberately stronger than the UAux-only auxiliary repair. A hint sampled from $\mathcal{H}^{\mathsf{SQIP2D\text{-}UAux}}_{q\text{-}\mathsf{unif}}$ gives the adversary polynomially many represented triples

$$(\phi_a : E \to F_a,\ \sigma : F_a \to E_1,\ \delta : E \to E_\delta),$$

together with the challenge label a, the residual degree d, and the verifier-facing metadata. Thus the corresponding hardness assumption is not plain EndRing and not the published SQIPrime2D assumption. It is a hinted EndRing assumption for a stronger auxiliary-information distribution.

The assumption strength should be read in the following order:

$$\mathsf{PlainEndRing} \;<\; \mathsf{HintEndRing}^{\mathsf{UAux}} \;<\; \mathsf{HintEndRing}^{q\text{-}\mathsf{unif}}.$$

Here $<$ means "less auxiliary information is given to the adversary," and therefore the assumption on the left is weaker. The q-unif distribution is used because it is the distribution for which the full SQIPrime2D-UAux hint transport theorem is true: carrying ϕ_a avoids the non-bijective attempt to push a killed q-line. The paper's SQIPrime-side security theorem should therefore be read as a foundations result about this explicitly strengthened hinted hardness game.

10.5 Security Consequence Under Q-HintDist

The security consequence has three layers. The simulator naturally produces a simulation-shaped distribution $\mathcal{H}_{q\text{-}\mathsf{sim}}$, because that distribution is tied to the transcripts used to answer signing queries. The EndRing-compatible reduction uses the challenge-carrying uniform distribution $\mathcal{H}_{q\text{-}\mathsf{unif}}$, because Theorem 12 proves exact transport for that distribution. The Q-HintDist assumption is the explicit cryptographic bridge between the two. Thus the theorem is a hinted EndRing theorem for a stronger challenge-carrying hint family, not a plain EndRing theorem without auxiliary information.

The challenge-carrying uniform distribution is intended for hardness reductions, whereas transcript simulation uses the distribution generated by the simulator. As in the general hint framework, the two distributions are connected by an explicit indistinguishability assumption.

Definition 26 (SQIPrime2D-UAux Q-HintDist). *Let $\mathcal{H}_{q\text{-}\mathsf{sim}}^{\mathsf{SQIP2D\text{-}UAux}}$ be the simulation-shaped challenge-carrying hint distribution: it contains the same fields as Definition 24, but samples them according to the SQIPrime2D-UAux simulator's constrained response distribution. The s-sample* SQIP2D-UAux-QHintDist *advantage is the hint-distinguishing advantage between*

$$\mathcal{H}_{q\text{-}\mathsf{sim}}^{\mathsf{SQIP2D\text{-}UAux}} \quad \textit{and} \quad \mathcal{H}_{q\text{-}\mathsf{unif}}^{\mathsf{SQIP2D\text{-}UAux}}$$

with s independently sampled hints, conditioned on the same public levelled curve distribution.

Theorem 13 (EUF-CMA theorem with challenge-carrying EndRing hints). *Let $\Sigma_{\mathsf{SQIP2D\text{-}UAux}}$ be the* SQIPrime2D-UAux *identification protocol and let $\mathsf{SIG}[\Sigma_{\mathsf{SQIP2D\text{-}UAux}}]$ be its Fiat–Shamir transform in the ROM. Suppose that:*

(i) the protocol is correct except with probability $\varepsilon_{\mathsf{corr}}$;
(ii) two accepting transcripts with the same commitment and distinct challenges extract an elliptic OneEnd *witness;*
(iii) the simulator using $\mathcal{H}_{q\text{-}\mathsf{sim}}^{\mathsf{SQIP2D\text{-}UAux}}$ is hint-assisted honest-verifier zero knowledge with error $\varepsilon_{\mathsf{HVZK}}(s)$ for s simulated signing transcripts;
(iv) the commitment min-entropy parameter is $\mu_{\mathsf{SQIP2D\text{-}UAux}}$;
(v) the challenge space has size $|\mathcal{C}|$;
(vi) $\mathcal{H}_{q\text{-}\mathsf{unif}}^{\mathsf{SQIP2D\text{-}UAux}}$ is used for the hinted EndRing *hardness experiment and is related to $\mathcal{H}_{q\text{-}\mathsf{sim}}^{\mathsf{SQIP2D\text{-}UAux}}$ by s-sample* QHintDist*;*
(vii) the UAux component is sampled in the ideal-hint model, or else by a sampler whose distance from ideal UAux is at most $\varepsilon_{\mathsf{UAux}}$ as in Proposition 10.

Write

$$\mathcal{H}_{qU} = \mathcal{H}_{q\text{-}\mathsf{unif}}^{\mathsf{SQIP2D\text{-}UAux}}.$$

Then, for every ROM EUF-CMA adversary $\mathcal{A}$ making at most Q_{RO} random-oracle queries and at most s signing queries, there are algorithms $\mathcal{B}$ and $\mathcal{D}$ such that

$$\begin{aligned}\mathrm{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{SIG}[\Sigma_{\mathsf{SQIP2D\text{-}UAux}}]}(\mathcal{A}) \le (Q_{\mathsf{RO}}+1)\Big(&\mathrm{Adv}^{\mathsf{HintEndRing}^{\mathcal{H}_{qU}}}(\mathcal{B}) + \mathrm{Adv}^{s\text{-}\mathsf{QHintDist}}_{\mathsf{SQIP2D\text{-}UAux}}(\mathcal{D}) \\ &+ \varepsilon_{\mathsf{HVZK}}(s) + |\mathcal{C}|^{-1} + s\cdot\varepsilon_{\mathsf{UAux}}\Big) \\ &+ (Q_{\mathsf{RO}}+s+1)\cdot s\cdot\mu_{\mathsf{SQIP2D\text{-}UAux}} + \varepsilon_{\mathsf{corr}},\end{aligned}$$

up to the algebraic loss of the hinted EndRing*-to-*OneEnd *reduction used in Corollary 7.*

Proof. Apply the Fiat–Shamir-with-hints theorem of Theorem 4 to the simulation-shaped challenge-carrying hint distribution $\mathcal{H}_{q\text{-sim}}^{\mathsf{SQIP2D\text{-}UAux}}$. Assumption (ii) gives special soundness for the elliptic OneEnd relation, assumption (iii) gives hint-assisted simulation for the s signing queries, assumption (iv) gives the commitment-hit term, and assumption (v) gives the challenge-guessing term. This yields a bound with a HintOneEnd hardness term using $\mathcal{H}_{q\text{-sim}}$ hints.

Next replace the s simulation-shaped hints by s challenge-carrying uniform hints. By Lemma 2 and Definition 26, the success probability changes by at most the displayed QHintDist advantage. If the theorem is instantiated with an approximate represented UAux sampler instead of the ideal-hint experiment, Proposition 10 adds at most $s \cdot \varepsilon_{\mathsf{UAux}}$ to the hint-experiment loss, charged inside the same conservative $(Q_{\mathsf{RO}} + 1)$ factor. The random-oracle programming, forking, challenge-guessing, and commitment-min-entropy terms are unchanged by these replacements.

Finally, Corollary 7 lifts the resulting $\mathsf{HintOneEnd}^{\mathcal{H}_{q\text{-unif}}}$ hardness term to the corresponding $\mathsf{HintEndRing}^{\mathcal{H}_{q\text{-unif}}}$ term for the stated self-reduction class. Substituting this lifted term gives the displayed inequality. The correctness error is added by a union bound over the event that an honestly generated transcript fails the ideal verification relation. □

The theorem is the strongest SQIPrime-side statement in the paper. It does not claim plain EndRing security without hints. It says that, after replacing the Algorithm-context auxiliary sampler by UAux and after using a challenge-carrying uniform hint distribution, SQIPrime2D-UAux has a fully pushable EndRing-compatible hint layer, with the remaining gap isolated as QHintDist and hinted EndRing hardness for the stronger challenge-carrying distribution.

11 Assumption Map, Related Work, and Applicability Boundary

The preceding sections establish three separate layers. First, Sections 4 and 5 show that recoverable higher-dimensional responses extract elliptic OneEnd witnesses. Second, Sections 6 and 7 give an exact hint-based Fiat–Shamir theorem for SQIsign2D–West, with the usual distinction between simulation-shaped hints and pushable hints. Third, Sections 8–10 isolate the SQIPrime2D auxiliary-isogeny obstruction, replace it by a curve-local UAux sampler, and then give a challenge-carrying uniform hint distribution that is fully pushable for the stated smooth transports. This section records the resulting assumption map and places the two theorem packages relative to prior work.

11.1 Layered Foundation Profiles

The profile notation below is a compact way to compare proof layers rather than a taxonomy of implementations. It records the object extracted by special soundness, the distribution needed for simulation, the distribution that is pushable through self-reductions, and the resulting EndRing-compatible hardness statement. Keeping these fields separate is essential: the same protocol can have elliptic extraction, simulation-shaped hints, and no pushable EndRing-compatible hint layer unless an additional distributional statement is supplied.

Definition 27 (Layered SQI foundation profile). *A SQI-style protocol has a layered foundation profile*

$$(\mathsf{Ext}, \mathsf{Sim}, \mathsf{Push}, \mathsf{Dist})$$

where Ext *is the object extracted from two accepting transcripts,* Sim *is the auxiliary information used by the simulator,* Push *records whether the corresponding hint distribution is transportable through the self-reduction isogenies used for EndRing compatibility, and* Dist *is the distributional assumption needed to replace simulation-shaped hints by a pushable distribution. The profile is* elliptic at extraction *when* Ext *is an efficiently evaluable non-scalar element of* $\mathrm{End}(E)$*, or an equivalent* OneEnd *witness, for a supersingular elliptic public-key curve* E*.*

The definition deliberately separates extraction from simulation. Higher-dimensional representation affects how a transcript is checked and how a simulator samples valid-looking responses. It does not by itself change the extracted hardness object once accepted responses publicly recover elliptic Hom elements.

Table 2. Foundation profiles for the schemes and variants used in the paper. The entries identify which layer carries the additional assumption; they are not performance comparisons.

Family	Extracted object	Simulation/hint layer	EndRing-compatible layer
SQIsign	Elliptic OneEnd	$\mathcal{H}_{\mathsf{sim}}$ and $\mathcal{H}_{\mathsf{unif}}$ hints	Pushable $\mathcal{H}_{\mathsf{unif}}$ plus HintDist
SQIsignHD	Elliptic Hom/OneEnd	HD response representation hints	Covered when an elliptic component is recovered
SQIsign2D–West	Elliptic OneEnd	Packaged UTO/FIDIO simulation hints	Exact theorem under 2D-West HintDist
Published SQIPrime2D	Elliptic OneEnd	RUCODIO+AIO shaped simulation	AIO is not curve-local as published
SQIPrime2D–UAux foundations variant	Elliptic OneEnd	Curve-local UAux auxiliary hints	Auxiliary layer exactly pushable
SQIPrime2D–UAux foundations variant with q-uniform hints	Elliptic OneEnd	Challenge-carrying (ϕ, σ, δ) hints	Fully pushable for smooth prime-to-degree transports

Proposition 12 (Separation of extraction and hint assumptions). *Let Π be a dual-closed SQI-style transcript system satisfying the recoverability interface of Definition 10. If Π satisfies the challenge-separation hypothesis of Theorem 2, then its extraction profile is elliptic* OneEnd*. If its simulator uses stateless polynomial-size representation oracles, then its simulation profile can be expressed as a simulation-shaped hint distribution. Replacing that distribution by an EndRing-compatible one requires an explicit pushability theorem and, unless the two distributions are identical, a hint-distinguishing assumption.*

Proof. The first claim is Theorem 2: two accepting transcripts give recovered elliptic maps whose composition with the challenge maps is an efficiently evaluable non-scalar endomorphism. The second claim is Theorem 3: the simulator's oracle outputs can be sampled in advance and packaged as a non-interactive hint without changing the transcript distribution. The final claim follows from Definition 16, Proposition 6, and Lemma 2. Pushability is the property that permits hints to be transported across the self-reduction curves, and hint indistinguishability is exactly the hybrid loss incurred when the simulator's hint distribution is replaced by the pushable distribution. □

11.2 Assumption Strength and Auxiliary Information

The SQIPrime-side results use stronger auxiliary information than the published scheme gives to an adversary. The correct ordering is

$$\mathsf{PlainEndRing} \;<\; \mathsf{HintEndRing}^{\mathsf{UAux}} \;<\; \mathsf{HintEndRing}^{q\text{-}\mathsf{unif}}.$$

The symbol $<$ denotes increasing auxiliary information given to the adversary, not a formal reduction direction. Plain EndRing gives no hints. $\mathsf{HintEndRing}^{\mathsf{UAux}}$ additionally gives curve-local uniform auxiliary isogeny hints. $\mathsf{HintEndRing}^{q\text{-}\mathsf{unif}}$ additionally gives represented random q-challenge isogenies and residual response isogenies. Thus the q-unif hardness assumption is the strongest SQIPrime-side assumption in the paper. It is used only because it is the distribution for which Theorem 12 proves full pushability.

This hierarchy is also the reason why the paper states separate SQIPrime-side conclusions. UAux gives a context-independent and exactly pushable replacement for the auxiliary AIO layer. Challenge-carrying q-unif hints give a fully pushable EndRing-compatible hinted layer for the UAux foundations variant. Neither statement is a proof of plain EndRing security for the published SQIPrime2D auxiliary sampler.

11.3 Map of the Main Schemes and Hint Layers

Table 2 records the main distinction. SQIsign2D–West is the cleanest direct instantiation of the general hint theorem: the simulator's higher-dimensional sampling oracles become $\mathcal{H}_{\mathsf{sim}}$ hints,

and a pushable distribution gives the EndRing-compatible formulation under an explicit HintDist assumption. Published SQIPrime2D has the same elliptic extraction target, but its two-dimensional auxiliary isogeny is generated by an Algorithm-context distribution. Section 9 removes this auxiliary obstruction by defining $\mathsf{UAux}(E, n)$, and Section 10 gives the stronger challenge-carrying distribution needed for full hint pushability.

The fully pushable SQIPrime2D–UAux statement should be read with its qualifier: it is a theorem about the challenge-carrying uniform hint layer, not about the UAux auxiliary component alone. UAux makes the auxiliary AIO layer context-independent and exactly pushable. The q-uniform hint distribution additionally carries the challenge isogeny ϕ_a and the residual response isogeny σ, so that the whole hint can be transported by bijections along smooth self-reduction isogenies satisfying the coprimality hypotheses of Theorem 12.

11.4 Relation to Existing Work

SQIsign and complete hint-based security. SQIsign introduced compact isogeny signatures based on quaternionic knowledge and the hardness of endomorphism-ring problems for supersingular elliptic curves [De +20]. The complete SQIsign security proof gives the Fiat–Shamir-with-hints framework used as the proof template in this paper: simulation-shaped hints support transcript simulation, pushable uniform hints support EndRing-compatible reductions, and hint indistinguishability bridges the two distributions [Ård+25]. The present paper uses this separation as a foundation and extends its role to higher-dimensional response representations and to the SQIPrime2D non-smooth challenge setting.

Higher-dimensional SQIsign variants. SQIsignHD introduced higher-dimensional response representations for SQI-style signatures [Dar+24]. SQIsign2D–West gives a fast two-dimensional variant with a security analysis based on higher-dimensional sampling oracles, while SQIsign2D–East explores a different two-dimensional route [Bas+24; NO24]. The recoverability theorem of Section 4 explains why these representation changes do not automatically create native higher-dimensional extraction assumptions: once the accepted data recovers an elliptic Hom element, special soundness extracts elliptic OneEnd.

SQIPrime and the auxiliary-isogeny boundary. SQIPrime introduces non-smooth challenge isogenies and, in the two-dimensional variant, an auxiliary isogeny that enables Kani-style verification of a response of degree $q \cdot d$ through a complementary degree $2^\alpha - q \cdot d$ map [DF24]. The published AIO model captures the distribution produced by the response algorithm, rather than a curve-local uniform auxiliary distribution. Section 8 formalizes this context dependence, Section 9 gives the UAux replacement, and Section 10 shows how carrying the challenge isogeny in the hint gives a fully pushable EndRing-compatible hint layer for the UAux variant.

Level structures and constrained Hom spaces. SQInstructor formulates SQIsign-style protocols using level structures and constrained Hom spaces and also uses hint distributions and hint-distinguishing assumptions in its security analysis [Bor+26]. This is conceptually close to the constrained-hint language of Section 6. The present SQIPrime2D–UAux construction addresses a different constraint: the non-smooth q-challenge line is made transportable by carrying the corresponding q-isogeny ϕ_a inside the uniform hint.

Ideal-to-isogeny translation. Two-dimensional ideal-to-isogeny algorithms and Qlapoti improve the algorithmic translation from quaternion ideals to isogeny representations [ON24; Bor+25]. These algorithms are representation-layer tools. They may be used to instantiate uniform auxiliary sampling or to produce efficient higher-dimensional representations, but their cryptographic role depends on the distribution in which their outputs are exposed to the simulator or adversary.

11.5 Applicability Boundary

The theorems in this paper apply to protocols whose accepted data either directly exposes an elliptic response map or exposes enough public higher-dimensional structure to recover one. The

decisive condition is not that the verifier performs a higher-dimensional computation. The decisive condition is the existence of public maps and evaluators certifying

$$\sigma_D = \pi \circ \Phi_D \circ \iota \in \mathrm{Hom}(E, F).$$

When this interface is present, the extraction theorem is elliptic. When it is absent, the present results do not classify the extraction assumption.

This boundary excludes native higher-dimensional relations in which the accepted object is a Hom element, endomorphism, or path relation for a non-split principally polarized abelian surface or higher-dimensional abelian variety with no public elliptic component. Such a protocol could require a genuinely higher-dimensional hardness assumption. The paper does not rule out that possibility; it proves that the covered SQI-family variants do not obtain native higher-dimensional extraction merely by using higher-dimensional response representations.

The pushability results have their own boundary. UAux pushability requires transport isogenies of degree coprime to the auxiliary degree. Challenge-carrying SQIPrime2D–UAux hints require smooth levelled transports whose degrees are coprime to the challenge degree q, the residual degree d, and the auxiliary degree $2^\alpha - q \cdot d$. These conditions are exactly what make the torsion and kernel maps bijective. Without them, the theorem gives no pushability claim.

Finally, the EndRing-compatible security consequences remain hinted assumptions. They do not assert plain EndRing security without auxiliary information. The role of pushability is to make the auxiliary information stable under the self-reductions used by the hinted EndRing-to-OneEnd argument. The role of HintDist or Q-HintDist is to justify replacing the simulator's distribution by the pushable one.

12 Conclusion

Higher-dimensional SQI-style signatures use higher-dimensional computations for representation, verification, and simulation. This paper shows that these roles should be separated. For the recoverable class, an accepted higher-dimensional response induces an elliptic Hom element

$$\sigma_D = \pi \circ \Phi_D \circ \iota,$$

and dual-closed special soundness extracts an elliptic non-scalar endomorphism. Thus the extraction layer is elliptic OneEnd; EndRing enters through the corresponding hinted reduction, not through the dimension of the response representation.

The simulation layer is the source of the additional assumptions. Oracle-assisted higher-dimensional simulators can be expressed as non-interactive simulation hints. A stronger EndRing-compatible statement requires a pushable hint distribution and a distributional bridge from the simulator's hints to that pushable distribution. This gives a uniform interpretation of the complete SQIsign proof strategy and of the higher-dimensional variants considered here.

The two concrete theorem packages instantiate this separation. For SQIsign2D–West, the paper gives an exact hint-based ROM EUF-CMA theorem from elliptic special soundness, simulation-shaped two-dimensional hints, commitment min-entropy, and Fiat–Shamir with hints. The pushable-hint corollary is obtained under the corresponding 2D-West HintDist assumption. For SQIPrime2D, the paper identifies the published AIO layer as Algorithm-context shaped and then defines a foundations variant, SQIPrime2D-UAux, which replaces that auxiliary layer by the context-independent uniform auxiliary distribution UAux. UAux is a proof-model distribution unless an efficient represented sampler is separately instantiated; approximate instantiations incur the explicit sampler loss stated in the theorem. The challenge-carrying uniform hints for this foundations variant carry the non-smooth challenge isogeny, the residual response, and the uniform auxiliary map; they are fully pushable under the stated smooth transport conditions and yield an EndRing-compatible hinted security theorem under Q-HintDist. This is not a claim that the published SQIPrime2D AIO sampler itself is pushable.

Consequently, representation dimension and hardness dimension are distinct invariants. Higher-dimensional SQI variants may be higher-dimensional in their algorithms and verifier-facing encodings while remaining elliptic at extraction. Their security foundations are governed by the hint distributions needed for simulation and by whether those distributions can be made pushable for the EndRing-compatible reductions.

Use of LLMs. Large language models (LLMs) were used exclusively to enhance the presentation quality of the paper, including spelling correction and English proofreading. None of the results, proofs, constructions, assumptions, or claims in the paper were generated by an LLM.

References

[Ård+25] Marius Andre Årdal, Andrea Basso, Luca De Feo, Sikhar Patranabis, and Benjamin Wesolowski. *A Complete Security Proof of SQIsign*. In: Advances in Cryptology – CRYPTO 2025, Part VI. 2025. IACR: 2025/379.

[Bas+24] Andrea Basso, Pierrick Dartois, Luca De Feo, Antonin Leroux, Luciano Maino, Giacomo Pope, Damien Robert, and Benjamin Wesolowski. *SQIsign2D–West: The Fast, the Small, and the Safer*. In: Advances in Cryptology – ASIACRYPT 2024, Part III. 2024. IACR: 2024/760.

[Bor+25] Giacomo Borin, Maria Corte-Real Santos, Jonathan Komada Eriksen, Riccardo Invernizzi, Marzio Mula, Sina Schaeffler, and Frederik Vercauteren. *Qlapoti: Simple and Efficient Translation of Quaternion Ideals to Isogenies*. In: Advances in Cryptology – ASIACRYPT 2025, Part IV. 2025. IACR: 2025/1604.

[Bor+26] Giacomo Borin, Luca De Feo, Guido Maria Lido, and Sina Schaeffler. *The SQInstructor: A Guide to SQIsign and the Deuring Correspondence with Level Structures*. Also available as arXiv:2603.09899. 2026. IACR: 2026/493.

[Dar+24] Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski. *SQIsignHD: New Dimensions in Cryptography*. In: Advances in Cryptology – EUROCRYPT 2024, Part I. 2024. IACR: 2023/436.

[De +20] Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski. *SQISign: Compact Post-quantum Signatures from Quaternions and Isogenies*. In: Advances in Cryptology – ASIACRYPT 2020, Part I. 2020. IACR: 2020/1240.

[DF24] Max Duparc and Tako Boris Fouotsa. *SQIPrime: A Dimension 2 Variant of SQISignHD with Non-smooth Challenge Isogenies*. In: Advances in Cryptology – ASIACRYPT 2024, Part III. 2024. IACR: 2024/773.

[NO24] Kohei Nakagawa and Hiroshi Onuki. *SQIsign2D–East: A New Signature Scheme Using 2-dimensional Isogenies*. In: Advances in Cryptology – ASIACRYPT 2024, Part III. 2024. IACR: 2024/771.

[ON24] Hiroshi Onuki and Kohei Nakagawa. *Ideal-to-Isogeny Algorithm Using 2-Dimensional Isogenies and Its Application to SQIsign*. In: Advances in Cryptology – ASIACRYPT 2024, Part III. 2024. IACR: 2024/778.

A Proof-Location Index

This appendix records where the complete proof details for the main formal claims are located. The main text states the theorem packages and explains their proof dependencies; the appendices provide the algebraic and distributional details needed to verify the claims without relying on informal proof summaries.

B Proof Details for Recoverable Higher-Dimensional Responses

This appendix gives the algebraic details underlying the recoverability statements in Section 4. The arguments are purely structural: they use only the public embedding-projection data and the correctness of the representation evaluator.

B.1 HD recoverability gives an elliptic Hom element

Let an accepted response datum certify a higher-dimensional morphism

$$\Phi_D : A \longrightarrow B$$

Table 3. Location of complete proof details for the main theorem packages.

Claim or theorem package	Complete proof details
HD recoverability and elliptic Hom extraction	Section B
Elliptic special-soundness transfer and scalar-exclusion interface	Section C
SQIsign2D–West recoverability, non-scalarity, Fiat–Shamir instantiation, and commitment min-entropy	Section D
SQIPrime2D-UAux context independence, UAux pushability, sampling-loss accounting, and Kani degree preservation	Section E and Section 9.1
Challenge-carrying q-unif pushability, joint distribution preservation, and EndRing compatibility	Section F and Section 10.3

together with public morphisms

$$\iota : E \longrightarrow A, \qquad \pi : B \longrightarrow F.$$

Assume that the representation interface evaluates Φ_D correctly on the image of ι, and that the public interface evaluates ι and π correctly. Define

$$\sigma_D = \pi \circ \Phi_D \circ \iota : E \longrightarrow F.$$

For every geometric point $P \in E$, the evaluator computes

$$P \longmapsto \iota(P) \longmapsto \Phi_D(\iota(P)) \longmapsto \pi(\Phi_D(\iota(P))).$$

This is exactly the evaluation of σ_D. Since a composition of homomorphisms of abelian varieties is a homomorphism, $\sigma_D \in \mathrm{Hom}(E, F)$. If $\sigma_D \neq 0$, then σ_D is a nonzero homomorphism between elliptic curves, hence has finite kernel and is a separable isogeny after removing any inseparable component. In the prime-to-p settings considered in the paper, the represented maps are separable, so σ_D is a separable isogeny. This proves the theorem stated in Theorem 1.

B.2 Product-component recovery

Suppose that

$$A = E_1 \times \cdots \times E_m, \qquad B = F_1 \times \cdots \times F_n,$$

and let

$$\Phi : A \longrightarrow B$$

be represented by its matrix of components

$$\Phi_{j,i} \in \mathrm{Hom}(E_i, F_j).$$

Let $\iota_i : E_i \to A$ be the inclusion into the i-th factor and let $\pi_j : B \to F_j$ be the projection from the j-th factor. For every $P \in E_i$,

$$\pi_j\big(\Phi(\iota_i(P))\big) = \Phi_{j,i}(P).$$

Therefore

$$\pi_j \circ \Phi \circ \iota_i = \Phi_{j,i}.$$

If the accepted transcript certifies the evaluation of Φ, then every component $\Phi_{j,i}$ is publicly recoverable by composing the certified evaluator with ι_i and π_j. This proves Proposition 2.

B.3 Kani matrix component recovery

Let a dimension-two Kani relation be represented by a pair of isogenies with common source

$$u : E \longrightarrow E_1, \qquad v : E \longrightarrow E_2,$$

with coprime degrees satisfying the Kani compatibility conditions for the displayed relation. The associated product morphism has the form

$$\Phi : EimesE \longrightarrow E_1 \times E_2, \qquad (P, Q) \longmapsto (u(P) + u'(Q), v(P) + v'(Q))$$

for the component maps certified by the response representation. Taking the inclusion into the first factor and the projection onto either output factor gives

$$\pi_1 \circ \Phi \circ \iota_1 = u, \qquad \pi_2 \circ \Phi \circ \iota_1 = v.$$

Thus the elliptic components needed by the extraction argument are not hidden in the product representation: they are public components of the certified higher-dimensional morphism. Target isomorphism ambiguity does not affect this conclusion. If the represented target is replaced by an isomorphic target, all recovered components are post-composed by a public isomorphism. Such post-composition preserves kernel containment, degree, and non-scalarity of the eventual endomorphism after dual composition. This gives the component-recovery statement used in Corollary 2 and Proposition 3.

C Proof Details for Elliptic Extraction

This appendix expands the special-soundness transfer of Section 5. The purpose is to record the exact algebraic object extracted from two accepting recoverable transcripts.

C.1 Composition extracted from two transcripts

Let two accepting transcripts have the same commitment object and distinct challenges $c_1 \neq c_2$. By recoverability, the two responses determine elliptic maps

$$\sigma_1 : E_c \longrightarrow F_1, \qquad \sigma_2 : E_c \longrightarrow F_2$$

or the protocol-specific dual-orientation equivalents. Dual closure gives public duals or public equivalent reverse maps, so the extractor forms an endomorphism of the public-key curve by composing the response components with the public challenge and commitment maps. In the common forward orientation, this has the form

$$\alpha = \beta_2 \circ \widehat{\sigma_2} \circ \sigma_1 \circ \beta_1 \in \mathrm{End}(E_{\mathsf{pk}}),$$

where β_1 and β_2 are the public maps dictated by the commitment and challenge data. In the dual orientation, the extractor replaces each response component by its certified dual and obtains an endomorphism of the same public-key curve. All maps in the displayed composition are efficiently evaluable by the transcript interface, so α is an efficiently evaluable endomorphism.

C.2 Nonzeroness

The composition is nonzero whenever the certified response and public challenge maps are nonzero and their degrees are finite. Indeed, nonzero homomorphisms between elliptic curves are isogenies onto their image, and compositions of nonzero isogenies are nonzero. The only way for the displayed composition to be zero would be for one of the certified maps to be zero, which is excluded by the acceptance relation, since the verifier checks degree, kernel, or evaluation constraints that a zero morphism cannot satisfy.

C.3 Scalar exclusion

The abstract extraction theorem isolates scalar exclusion as a separate scheme-specific property. The transfer statement is then exact: once a protocol proves that the composed endomorphism satisfies

$$\alpha \notin \mathbb{Z},$$

the extracted object is a valid OneEnd witness. This paper uses the following two scalar-exclusion forms.

Degree-based exclusion. Suppose the extracted endomorphism has certified degree D and the verified degree data imply that D is not a square integer. If $\alpha = [m]$ were scalar, then

$$\deg(\alpha) = \deg([m]) = m^2,$$

which is a square in $\mathbb{Z}$. This contradicts the certified value of D. Hence α is non-scalar.

Scheme-proved kernel-line exclusion. Some SQI-style protocols prove scalar exclusion from exact q-primary kernel conditions. The abstract condition used in Theorem 2 is the following: for any two accepting transcripts with the same commitment and distinct challenges, the protocol-specific composition

$$\widehat{\kappa}_2 \circ \kappa_1 \in \mathrm{End}(E)$$

is non-scalar. The proof of this condition must be supplied by the concrete scheme, because it depends on how the verifier certifies exact q-primary kernel structure and on whether residual degrees may contain extra factors of q. For SQIPrime2D, this condition is exactly the statement of the published SQIPrime special-soundness theorem [DF24, Proposition 5]. In the notation used in Proposition 8, that theorem constructs

$$\theta_A = \widehat{\kappa}_2 \circ \kappa_1$$

from two accepted transcripts and proves $\theta_A \notin \mathbb{Z}$ by a degree-divisibility argument. If the residual degrees are prime to q, scalarity forces equality of the two certified challenge lines. If a residual degree is divisible by q, the SQIPrime2D parameter bound forces exact single divisibility, and scalarity forces the exact q^2-parts of the two response isogenies to coincide; this again implies equality of the challenge lines. In all cases distinct challenges contradict scalarity.

Combining nonzeroness, efficient evaluability, and the relevant scalar-exclusion theorem gives the OneEnd witness claimed in Theorem 2.

D Proof Details for SQIsign2D–West

This appendix records the concrete algebraic steps used in the SQIsign2D–West theorem package of Section 7. The proof details are organized around the two nontrivial inputs to the security theorem: the non-scalarity of the extracted endomorphism and the Fiat–Shamir-with-hints instantiation, including commitment min-entropy and hint replacement.

D.1 Recoverability of accepted responses

An accepted SQIsign2D–West response contains higher-dimensional data certifying a response isogeny in dimension two together with the torsion evaluations required by the verifier. The verifier's acceptance condition ensures that this data represents an elliptic response component

$$\widetilde{\sigma} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$$

of the prescribed response degree bound. This is exactly an instance of HD recoverability: the component is recovered by composing the certified product morphism with the public inclusion and projection maps. The response is therefore an elliptic Hom element, not a native higher-dimensional witness.

D.2 Concrete non-scalarity argument

Let two accepting responses share the same commitment curve E_{com} and have distinct challenge isogenies

$$\phi_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}, \qquad \phi'_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E'_{\mathsf{chl}}.$$

Let the recovered response maps be

$$\sigma : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}, \qquad \sigma' : E_{\mathsf{com}} \longrightarrow E'_{\mathsf{chl}},$$

each of degree at most $2^{e_{\mathsf{rsp}}}$. The extractor forms

$$\alpha = \widehat{\phi'_{\mathsf{chl}}} \circ \sigma' \circ \widehat{\sigma} \circ \phi_{\mathsf{chl}} \in \mathrm{End}(E_{\mathsf{pk}}).$$

Assume for contradiction that $\alpha = [m]$ for some integer m. Compose the equality

$$\widehat{\phi'_{\mathsf{chl}}} \circ \sigma' \circ \widehat{\sigma} \circ \phi_{\mathsf{chl}} = [m]$$

on the left by ϕ'_{chl} and on the right by $\widehat{\phi}_{\mathsf{chl}}$. Using

$$\phi'_{\mathsf{chl}} \circ \widehat{\phi'_{\mathsf{chl}}} = [\deg(\phi'_{\mathsf{chl}})] \quad \text{and} \quad \phi_{\mathsf{chl}} \circ \widehat{\phi}_{\mathsf{chl}} = [\deg(\phi_{\mathsf{chl}})],$$

and the commutativity of scalar multiplication with homomorphisms, we obtain the exact equality

$$[m] \circ \phi'_{\mathsf{chl}} \circ \widehat{\phi}_{\mathsf{chl}} = [\deg(\phi_{\mathsf{chl}}) \cdot \deg(\phi'_{\mathsf{chl}})] \circ \sigma' \circ \widehat{\sigma}. \tag{A.1}$$

For a nonzero homomorphism f, write $\mathrm{scal}(f)$ for the largest positive integer u such that $f = [u] \circ f_0$ for some homomorphism f_0. The quotient f_0 is then the cyclic part of f: its kernel has no nontrivial scalar multiplication factor left. This decomposition is unique up to target isomorphism, because the largest scalar subgroup contained in $\ker(f)$ is determined by $\ker(f)$.

Decompose the two homomorphisms appearing in (A.1) as

$$\phi'_{\mathsf{chl}} \circ \widehat{\phi}_{\mathsf{chl}} = [2^a] \circ \psi, \qquad \sigma' \circ \widehat{\sigma} = [d] \circ \nu,$$

where ψ and ν have cyclic kernels. Equality (A.1) then becomes

$$[2^a m] \circ \psi = [d \cdot \deg(\phi_{\mathsf{chl}}) \cdot \deg(\phi'_{\mathsf{chl}})] \circ \nu. \tag{A.2}$$

The integer multiplying the cyclic part on the left of (A.2) is the largest scalar divisor of the left-hand side, and the integer multiplying the cyclic part on the right is the largest scalar divisor of the right-hand side. Since the homomorphisms in (A.2) are equal, their largest scalar divisors are equal and their cyclic parts are equal up to target isomorphism. Hence

$$\psi = \nu$$

up to the canonical target isomorphism.

The response side gives an upper bound for this common cyclic component. Since ν is the cyclic part of $\sigma' \circ \widehat{\sigma}$,

$$\deg(\nu) \le \deg(\sigma' \circ \widehat{\sigma}) = \deg(\sigma') \cdot \deg(\sigma) \le 2^{2e_{\mathsf{rsp}}}. \tag{A.3}$$

The challenge side gives a lower bound. SQIsign2D–West's distinct-challenge lemma states that the largest integer divisor of

$$\phi'_{\mathsf{chl}} \circ \widehat{\phi}_{\mathsf{chl}}$$

is smaller than $2^{e_{\mathsf{chl}}}$. Since each challenge isogeny has degree 2^e, the degree of the cyclic part ψ is

$$\deg(\psi) = \frac{\deg(\phi'_{\mathsf{chl}} \circ \widehat{\phi}_{\mathsf{chl}})}{2^{2a}} = \frac{2^{2e}}{2^{2a}} > 2^{2(e-e_{\mathsf{chl}})}. \tag{A.4}$$

The parameter condition $e_{\mathsf{chl}} + e_{\mathsf{rsp}} \le e$ gives

$$2^{2(e-e_{\mathsf{chl}})} \ge 2^{2e_{\mathsf{rsp}}}.$$

Combining (A.3) and (A.4) contradicts $\psi = \nu$. Therefore α cannot be scalar. The extracted endomorphism is non-scalar and hence is a valid elliptic OneEnd witness.

D.3 Hint-assisted HVZK and Fiat–Shamir instantiation

Let s be the number of signing queries. The simulation-shaped hint distribution for SQIsign2D–West contains exactly the non-interactive data returned by the UTO/FIDIO-style simulation oracles for the s simulated transcripts. Once these values are placed in the hint, the simulator performs only deterministic transcript assembly and random-oracle programming. Therefore the hint-assisted simulated transcript distribution is identical to the oracle-assisted simulated transcript distribution, up to the statistical error already present in the SQIsign2D–West simulator.

The Fiat–Shamir theorem with hints is then instantiated with: the special-soundness extractor above; the s-hint HVZK simulator; the commitment min-entropy bound for the commitment curve distribution; and challenge space of size $2^{e_{\mathsf{chl}}}$. The adversary's EUF-CMA advantage is bounded by the corresponding $\mathcal{H}_{\mathsf{sim}}$-HintOneEnd advantage, the HVZK error, the challenge-guessing term, and the commitment min-entropy term. Replacing $\mathcal{H}_{\mathsf{sim}}$ by a pushable uniform hint family is a separate distribution-replacement hybrid. The difference between the two games is exactly the SQIsign2D–West $\mathsf{HintDist}$ advantage, which yields the corollary under pushable hints.

E Proof Details for SQIPrime2D-UAux

This appendix gives the proof details for the context-independent auxiliary layer introduced in Section 9.

E.1 SQIPrime2D scalar-exclusion theorem used by the boundary section

Proposition 8 uses the special-soundness theorem proved for published SQIPrime2D. For completeness of the dependency, we record the exact imported statement in the notation of this paper. Given two accepting SQIPrime2D transcripts with the same commitment E_1 and distinct challenges, verification recovers elliptic response isogenies

$$\kappa_1,\, \kappa_2 : E_A \longrightarrow E_1.$$

SQIPrime Proposition 5 proves that

$$\theta_A = \widehat{\kappa}_2 \circ \kappa_1$$

is a non-scalar endomorphism of E_A. The proof in SQIPrime first handles the case in which both residual degrees are prime to the challenge prime q by the same kernel-line separation used for the four-dimensional variant. If $q \mid d_1$, it writes $d_1 = d_1' \cdot q$; the constraints $q \cdot d_1 < 2^\alpha$ and $q^3 > 2^\alpha$ imply $q \nmid d_1'$. If θ_A were scalar, the equality of degrees forces $q \mid d_2$, so $d_2 = d_2' \cdot q$ with $q \nmid d_2'$. The two response maps then factor through right factors κ_1' and κ_2' of degree q^2, followed by outer factors of degrees d_1' and d_2' prime to q. SQIPrime Proposition 5 proves that scalarity of $\widehat{\kappa}_2 \circ \kappa_1$ forces these right q-primary factors to have the same kernel on $E_A[q]$. Hence the two challenge kernels coincide, contradicting the distinctness of the challenges. This proves non-scalarity.

The paper uses this imported theorem only for the extraction profile of published SQIPrime2D. The UAux and challenge-carrying theorems proved below are separate distributional statements about auxiliary and hint layers.

E.2 Sampling model and sampler-loss accounting

The main text treats UAux first as an ideal hinted distribution. If a concrete sampler is supplied, Proposition 10 gives the exact hybrid argument: replacing at most s ideal samples by approximate represented samples changes any experiment by at most $s \cdot \varepsilon_{\mathsf{UAux}}$, and by $(Q_{\mathsf{RO}} + 1) \cdot s \cdot \varepsilon_{\mathsf{UAux}}$ when charged inside the conservative Fiat–Shamir-with-hints loss. This statement is purely distributional and does not assume a particular ideal-to-isogeny algorithm.

E.3 Uniform auxiliary sample space

Fix a supersingular elliptic curve $E/\mathbb{F}_{p^2}$ and an integer $n > 0$ with $p \nmid n$. The distribution $\mathsf{UAux}(E, n)$ is uniform over the finite set of cyclic subgroup schemes

$$K \subset E[n], \qquad |K| = n,$$

modulo the usual target isomorphism of the quotient isogeny

$$\delta_K : E \longrightarrow E/K.$$

Equivalently, when an endomorphism-ring representation is used, the same distribution is the uniform distribution over the corresponding primitive norm-n ideals, modulo the equivalence identifying target-isomorphic quotient isogenies. The paper uses one fixed interpretation throughout the theorem statement; the pushability proof only needs the kernel-set interpretation.

E.4 Context independence

Let two signing contexts produce the same visible curve E. These contexts may contain different paths

$$\tau_1 : E_0 \longrightarrow E, \qquad \tau_2 : E_0 \longrightarrow E,$$

different ideal representatives, and different KaniDoublePath randomness. The sample space defining $\mathsf{UAux}(E, n)$ is nevertheless the same set of cyclic subgroup schemes of $E[n]$. Since the sampling distribution is uniform over that set and contains no reference to τ_i, the distributions induced by the two contexts are identical. This proves context independence.

E.5 Auxiliary pushability

Let

$$\rho : E \longrightarrow F$$

be a separable isogeny such that

$$\gcd(\deg \rho, n) = 1.$$

The dual isogeny satisfies

$$\widehat{\rho} \circ \rho = [\deg \rho].$$

Since multiplication by $\deg \rho$ is invertible on $E[n]$, the restriction

$$\rho|_{E[n]} : E[n] \longrightarrow F[n]$$

is an isomorphism with inverse

$$[\deg \rho]^{-1} \cdot \widehat{\rho}|_{F[n]}.$$

Consequently

$$K \longmapsto \rho(K)$$

is a bijection between cyclic order-n subgroup schemes of $E[n]$ and cyclic order-n subgroup schemes of $F[n]$. The quotient by K pushes out to the quotient by $\rho(K)$, so the image of a uniform $\delta_K : E \to E/K$ is a uniform $\delta_{\rho(K)} : F \to F/\rho(K)$. Therefore

$$\rho_* \mathsf{UAux}(E, n) = \mathsf{UAux}(F, n).$$

E.6 Preservation of the dimension-two Kani degree relation

In SQIPrime2D-UAux, the response map has degree

$$\deg(\kappa) = q \cdot d$$

with $q \nmid d$, and the auxiliary isogeny is sampled with degree

$$n_{\mathsf{aux}} = 2^\alpha - q \cdot d.$$

The admissibility conditions include

$$q \cdot d < 2^\alpha, \qquad p \nmid n_{\mathsf{aux}}.$$

The degree sum is

$$\deg(\kappa) + \deg(\delta) = q \cdot d + (2^\alpha - q \cdot d) = 2^\alpha.$$

Moreover,

$$\gcd(q \cdot d, 2^\alpha - q \cdot d) = \gcd(q \cdot d, 2^\alpha) = 1$$

because $q \cdot d$ is odd. Thus the degree and coprimality hypotheses of the dimension-two Kani relation hold. The verifier still requires valid evaluation and Kani metadata for (κ, δ); UAux changes the distribution of δ, not the algebraic verification relation.

F Proof Details for Challenge-Carrying SQIPrime2D-UAux Hints

This appendix expands the fully pushable challenge-carrying hint theorem of Section 10. The proof is decomposed into the challenge component, the residual response component, the UAux component, and the final joint-distribution argument, matching the structure of Theorem 12.

F.1 Why line-constrained hints are not enough

A line-constrained hint has the form

$$(L, \kappa, \delta), \qquad L \subset E[q], \qquad L \subseteq \ker(\kappa).$$

If the transport map is the challenge quotient $\phi_L : E \to E/L$, then ϕ_L kills L. Hence $\phi_L|_{E[q]}$ is not an isomorphism, and there is no bijective pushforward of the full line-constrained sample space through ϕ_L. This is the obstruction that motivates challenge-carrying hints.

F.2 Challenge-carrying distribution

Let the public instance be a levelled curve

$$\mathbf{E} = (E, P, Q),$$

where (P, Q) is the public q-torsion basis. A challenge-carrying hint samples

$$a \leftarrow \mathbb{Z}/q\mathbb{Z}, \qquad L_a = \langle P + [a]Q \rangle,$$

then samples the q-isogeny

$$\phi_a : E \longrightarrow F_a$$

with kernel L_a. It samples a residual response isogeny

$$\sigma : F_a \longrightarrow E_1$$

of degree d, with $q \nmid d$, and an auxiliary isogeny

$$\delta \leftarrow \mathsf{UAux}(E, 2^\alpha - q \cdot d).$$

The represented response is

$$\kappa = \sigma \circ \phi_a.$$

By construction, $L_a \subseteq \ker(\kappa)$ and $\deg(\kappa) = q \cdot d$. Together with δ, the degree equality and coprimality conditions of the dimension-two Kani relation are satisfied as in Section E.

F.3 Pushout of the challenge component

Let

$$\rho : (E, P, Q) \longrightarrow (E', P', Q')$$

be a smooth transport isogeny with

$$P' = \rho(P), \qquad Q' = \rho(Q), \qquad \gcd(\deg \rho, q) = 1.$$

Then $\rho|_{E[q]}$ is an isomorphism and

$$\rho(L_a) = \langle P' + [a]Q' \rangle.$$

The pushout square

$$\begin{array}{ccc} E & \xrightarrow{\phi_a} & F_a \\ \downarrow \rho & & \downarrow \rho_F \\ E' & \xrightarrow{\phi'_a} & F'_a \end{array}$$

exists in the category of elliptic curves and separable isogenies. The lower horizontal map ϕ'_a is the q-isogeny with kernel $\rho(L_a)$. Since a is uniform and the map $L_a \mapsto \rho(L_a)$ is a bijection on the challenge-line family, ϕ'_a is distributed exactly as a fresh challenge isogeny from (E', P', Q').

F.4 Pushout of the residual and auxiliary components

Conditioned on a, the right vertical map $\rho_F : F_a \to F'_a$ has degree $\deg \rho$. If

$$\gcd(\deg \rho, d) = 1,$$

then the same kernel-bijection proof gives a bijection between admissible degree-d residual response isogenies from F_a and admissible degree-d residual response isogenies from F'_a. Thus a uniform σ pushes to a uniform σ'.

If

$$\gcd(\deg \rho, 2^\alpha - q \cdot d) = 1,$$

then Theorem 9 gives

$$\rho_* \mathsf{UAux}(E, 2^\alpha - q \cdot d) = \mathsf{UAux}(E', 2^\alpha - q \cdot d).$$

Thus δ pushes to a fresh auxiliary sample δ' from the target distribution.

F.5 Joint distribution and EndRing compatibility

The challenge scalar a, residual degree d, residual response isogeny σ, and auxiliary isogeny δ are sampled with the conditional independence stated in Definition 24. Each conditional sample space is transported by a bijection under the coprimality condition

$$\gcd(\deg \rho, q \cdot d \cdot (2^\alpha - q \cdot d)) = 1.$$

Therefore the pushed tuple

$$(a, \phi'_a, \sigma', \delta')$$

has exactly the distribution of a fresh challenge-carrying hint on (E', P', Q'). This proves full pushability.

The EndRing-compatible statement follows from this exact transport property. The hinted OneEnd-to-EndRing reduction transports hard instances along smooth self-reduction isogenies. The reduction uses disjoint independent batches of source hints for distinct calls to the hinted OneEnd solver and pushes each batch through the corresponding transport isogeny. Since the challenge-carrying hint family is preserved by the same transports, each pushed batch has the distribution of fresh independent target hints. Thus a solver for OneEnd with these hints can be used in the corresponding hinted $\mathsf{EndRing}$ reduction without changing the hint distribution. The only distributional replacement needed for signatures is the explicitly stated Q-HintDist assumption between simulation-shaped hints and challenge-carrying uniform hints.