Structural Non-Identifiability and Quantum-Attack Orthogonality in Mock-Modular Cryptographic Architectures

Marcos Eduardo Elias

The Ramanujan Institute — Brazil

Abstract

We introduce and formally analyze a cryptographic paradigm grounded in **structural non-identifiability under restricted observability**: a security notion derived from the structural absence of distinguishing information in the observable domain, rather than from computational inversion hardness. The mathematical substrate is the theory of **harmonic Maass forms**, whose holomorphic part is locally evaluable and admits a finite, verifiable transcript, while the non-holomorphic completion depends on a global shadow datum not recoverable from any finite local observation.

We formalize the **Restricted Observable Model (ROM)** and the **Mock Modular Identification Problem (MMIP)**. We establish nine principal results: (I) *information-theoretic indistinguishability* with zero adversarial advantage, holding for all algorithms including unbounded ones; (II) *formal undecidability of universal identification* via explicit halting-encoding construction; (III–IV) *quantum-attack orthogonality* ruling out both Shor-type (hidden subgroup) and Grover-type (amplitude amplification) attacks; (V) *completeness and soundness duality*; (VI) *ontological hardness*; (VII) *a three-way hardness taxonomy*; (VIII) an *ε-separation lemma* and *dimension witness* giving explicit parameters satisfying all hypotheses; and (IX) tight security reductions for a KEM (IND-CPA) and signature scheme (EUF-CMA) with loss factors precisely quantified.

The concrete instantiation uses harmonic Maass forms at weight $k = 1/2$ with shadow classes in $S_{3/2}(\Gamma_0(4M))$. We prove that for any λ, setting $M = 4 \cdot \mathrm{lcm}(1,\ldots,\lambda)$ yields $\dim S_{3/2}(\Gamma_0(4M)) \geq \lambda$, satisfying the exponential ambiguity hypothesis. An explicit ε-separation lemma shows how to choose the shadow separation parameter maintaining transcript invariance while guaranteeing global distinctness. Security reductions are stated with explicit loss factors: $\mathrm{Adv}^{\mathrm{KEM\text{-}IND\text{-}CPA}}(A,\lambda) \leq (q_H + 1) \cdot 2^{-\alpha\lambda}$ for q_H random-oracle queries. A worked discretization example ($m = 2$) traces both continuous and discrete phases end-to-end.

Keywords: mock modular forms, harmonic Maass forms, post-quantum cryptography, non-identifiability, restricted observable model, ontological hardness, information-theoretic security, quantum-attack orthogonality, key encapsulation, digital signatures, undecidability, shadow operator, ε-separation, dimension witness.

1. Introduction

1.1 Background and Motivation

The security of virtually all deployed public-key cryptography rests on a common structural premise: the public key *uniquely encodes* a private key, and security is guaranteed by the computational difficulty of recovering that private key. RSA [RSA78] relies on integer factorization; Diffie-Hellman [DH76] and elliptic-curve variants [Mil86,Kob87] on the discrete logarithm; lattice-based schemes [Reg05,LPR13] on LWE and SIS. In every case, a *unique* secret is embedded in the public data, and security amounts to inversion hardness.

Quantum computation threatens this paradigm. Shor's algorithm [Sho94] breaks RSA, DH, and ECC in polynomial quantum time; Grover's algorithm [Gro96] provides a quadratic speedup for unstructured search. NIST's PQC standardization [NIST24] settled on lattice-based (Kyber, Dilithium, FALCON), code-based (McEliece), and hash-based (SPHINCS+) candidates, all maintaining inversion hardness while seeking quantum-resistant assumptions.

This paper investigates a **categorically different design principle**: make the public data sufficient for verification but structurally insufficient for unique identification. The adversary does not face a costly computation; it faces an information-theoretically impossible task. The key shift is from "how much does it cost to recover the secret?" to "what, in fact, is determined by what is observed?"

Our mathematical substrate is the theory of **harmonic Maass forms** [Zwe02,BF04,Zag09,Ono09]. Every harmonic Maass form H decomposes as $H = H^+ + H^-$, where the holomorphic part H^+ is expressible as a convergent q-series (locally evaluable), while H^- depends on the *shadow* $g = \xi_k(H) \in S_{2-k}(\Gamma)$ through a global Eichler-type integral. The shadow is not recoverable from any finite window of Fourier coefficients of H^+ — this is Lemma 2 of this paper, stated and proved in full. This local-global separation is the asymmetry we exploit cryptographically.

1.2 JoC-Level Precision: What This Paper Does

We address each of the three structural requirements that define Journal of Cryptology quality:

> **(a) Complete formal reductions.** Theorems 12 and 13 provide complete game-based security proofs of the KEM and signature scheme, reducing their security to the MMIP via explicit PPT algorithms B constructing from any adversary A a MMIP solver. Loss factors are precisely quantified: $Adv^{KEM-IND-CPA}(A,\lambda) \leq (q_H+1) \cdot 2^{-\alpha\lambda}$, where q_H is the adversary's random oracle query count.
>
> **(b) Explicit adversary model and oracle specification.** Definition 6 (ROM) specifies exactly which oracles the adversary has, with their precise input/output behavior. The auxiliary functions BuildChallenge, ResolveChallenge, and CheckFamilyProof are formally specified in Definitions 19–21. The games KEM-IND-CPA and SIG-EUF-CMA follow the standard Bellare-Rogaway formalization [BR94].
>
> **(c) Explicit constructions verifying all hypotheses.** Lemma 5 (Dimension Witness) gives an explicit modulus $M(\lambda)$ such that $\dim S_{3/2}(\Gamma_0(4M)) \geq \lambda$ for all λ, verifying Hypothesis H2 concretely. Lemma 4 (ε-Separation) provides an explicit admissible range for the separation

parameter ε jointly satisfying transcript invariance and global shadow distinctness. These close the two main gaps that previous treatments left as assumptions.

1.3 Our Contributions

1. ROM and MMIP formalized (Definitions 5–7).

2. Information-theoretic indistinguishability with zero advantage for all adversaries, including unbounded ones (Theorems 1 and 18).

3. Formal undecidability via explicit halting-encoding construction (Theorem 2 with full construction).

4. Quantum-attack orthogonality: no abelian HSP reduction (Theorem 3), no selective Grover target (Theorem 4).

5. Completeness and soundness duality (Theorem 16); master structural security theorem (Theorem 17).

6. Ontological hardness definition and three-way hardness taxonomy (Definition 11, Theorems 18–20).

7. ε-Separation Lemma and Dimension Witness (Lemmas 4–5): explicit parameters satisfying H1–H3.

8. Formal specification of BuildChallenge, ResolveChallenge, CheckFamilyProof (Definitions 19–21).

9. Tight IND-CPA KEM reduction (Theorem 12) and EUF-CMA signature reduction (Theorem 13) with explicit loss factors.

1.4 Related Work

Post-Quantum Cryptography. NIST PQC [NIST24]: Kyber [BDK+18], Dilithium [DKL+18], FALCON [PFH+20], SPHINCS+ [BHK+19], McEliece [ABC+20]. SIDH broken [CD22,MMPPW22,Rob23]. Our approach is orthogonal: informational rather than computational security.

Mock Modular Forms. Ramanujan [Ram20]; Zwegers' completion [Zwe02]; Bruinier-Funke [BF04]; Zagier [Zag09]; Ono [Ono09]. First cryptographic application to our knowledge.

Information-Theoretic Cryptography. Shannon [Sha49]; BGW [BGW88]; CCD [CCD88]. Our mechanism is distinct: security derives from structural informational absence, not shared secrets or noise.

Undecidability in Cryptography. Semantic security [GM84]; obfuscation impossibility [BGI+01]. Our Theorem 2 gives a formal undecidability impossibility with explicit construction.

Quantum Algorithms. Abelian HSP [Bon90,EHKS04,HRT03]; BBBV lower bound [BBBV97]. Theorems 3–4 show MMIP lies outside both frameworks.

Reduction Tightness. KEM security analysis following [BHK+19,HHK17]; tight reductions methodology [Scu15]. Our loss factor (q_H+1) is tight in the random oracle model.

1.5 Organization

Section 2: Mathematical background with full proofs of Lemmas 1–3. Section 3: ROM, Property P1, and MMIP. Section 4: Information-theoretic analysis; completeness/soundness duality; master structural theorem; tripartite contrast. Section 5: Quantum orthogonality. Section 6: Formal hardness theory; ontological hardness; three-way taxonomy. Section 7: ε-Separation Lemma and Dimension Witness. Section 8: Parameterization and security bounds. Section 9: Concrete analytic instantiation. Section 10: Discretization. Section 11: Formal auxiliary function specifications. Section 12: Formal security definitions. Sections 13–14: KEM and signature with complete reductions. Section 15: Concrete security. Section 16: Limitations and open problems. Section 17: Conclusion.

2. Mathematical Background

2.1 Modular Forms and Cusp Forms

Let $\mathbb{H} = \{\tau \in \mathbb{C} : \mathrm{Im}(\tau) > 0\}$ and $q = e^{2\pi i\tau}$. Let $\Gamma \leq SL_2(\mathbb{Z})$ be a congruence subgroup, with metaplectic cover when $k \in (1/2)\mathbb{Z}$. The weight-k slash operator is $(f|_k\gamma)(\tau) = j(\gamma,\tau)^{-k} f(\gamma\tau)$, where $j(\gamma,\tau) = c\tau+d$ for $\gamma = (a\ b;\ c\ d)$ with appropriate half-integral multiplier system. We use the standard spaces $M_k(\Gamma)$ (holomorphic modular forms), $S_k(\Gamma) \subseteq M_k(\Gamma)$ (cusp forms), and $M_k^!(\Gamma)$ (weakly holomorphic modular forms) as in [Shi73,Iwa97,BF04]. The dimension formula for $S_k(\Gamma_0(N))$ follows from Riemann-Roch [DS05, §3.5].

2.2 Harmonic Maass Forms

$$\Delta_k = -y^2(\partial^2/\partial x^2 + \partial^2/\partial y^2) + iky(\partial/\partial x + i\partial/\partial y), \quad \tau = x + iy.$$

> **Definition 1** (*Harmonic Maass Form*). A smooth function $H : \mathbb{H} \to \mathbb{C}$ is a harmonic Maass form of weight k for Γ if: (i) $H|_k\gamma = H$ for all $\gamma \in \Gamma$; (ii) $\Delta_k H = 0$; (iii) H has at most linear exponential growth at each cusp of Γ. We write $H \in H_k(\Gamma)$.

Every $H \in H_k(\Gamma)$ admits the canonical decomposition [BF04, Prop. 3.2]:

$$H(\tau) = H^+(\tau) + H^-(\tau),$$

where $H^+(\tau) = \Sigma_{n >> -\infty} c^+(n)q^n$ is the **holomorphic part** and $H^-(\tau) = \Sigma_{n<0} c^-(n)\Gamma(1-k, 4\pi|n|y)q^n$ is the **non-holomorphic part**, with $\Gamma(s,t)$ the incomplete gamma function. A **mock modular form** of weight k is the holomorphic part H^+ of some $H \in H_k(\Gamma)$.

2.3 The Shadow Operator: Definition and Properties

Define the antilinear differential operator:

$$\xi_k(H)(\tau) := 2i \cdot y^k \cdot \bar{\partial}(\partial H/\partial \bar{\tau}), \quad \tau = x + iy.$$

By [BF04, Thm. 3.7], $\xi_k : H_k(\Gamma) \to S_{2-k}(\Gamma)$ is well-defined and surjective. The cusp form $g = \xi_k(H) \in S_{2-k}(\Gamma)$ is called the **shadow** of H. We now state and prove the three foundational lemmas.

> **Lemma 1** (*Shadow Determines Non-Holomorphic Part*). Let $H_1, H_2 \in H_k(\Gamma)$ with $H_1^+ = H_2^+$. Then $H_1 = H_2$ if and only if $\xi_k(H_1) = \xi_k(H_2)$. Consequently, H^- is uniquely determined by the shadow $g = \xi_k(H)$.

Proof. Consider $D = H_1 - H_2$. Since $H_1^+ = H_2^+$, the holomorphic part of D vanishes: $D^+ \equiv 0$. Hence D has only a non-holomorphic part.

A harmonic Maass form with vanishing holomorphic part satisfies $\Delta_k D = 0$ and $D^+ = 0$. By [BF04, §3], such forms lie in the kernel of ξ_k when they also satisfy growth conditions at cusps: $\ker(\xi_k) = M_k^!(\Gamma) \cap \{\text{forms with zero holomorphic part}\} = \{0\}$. More precisely, since $D^+ = 0$, we have $D = D^-$, and $\xi_k(D) = \xi_k(H_1) - \xi_k(H_2)$. If this vanishes, then D lies in $\ker(\xi_k|_{\text{anti-holomorphic}})$. The anti-holomorphic functions in $H_k(\Gamma)$ satisfying growth bounds form a space on which ξ_k is injective [BF04, Thm. 3.7]. Hence $D = 0$, giving $H_1 = H_2$.

> **Lemma 2** (*Local Observations Do Not Determine the Shadow*). Let $f = H^+$ for some $H \in H_k(\Gamma)$. No finite set of Fourier coefficients $\{c^+(n)\}_{|n|\le N}$ or evaluations $\{f(\tau_j)\}_{j=1}^r$ determines the shadow $g = \xi_k(H)$ within the space of all harmonic completions of f.

Proof. It suffices to exhibit two distinct completions with identical holomorphic data. Let $g_1, g_2 \in S_{2-k}(\Gamma)$ with $g_1 \neq g_2$. Define $H_i = f + R_{g_i}$ where R_{g_i} is the Eichler-type integral (Section 2.4).

By construction: $H_i^+ = f$ (the holomorphic anchor is unchanged); $\xi_k(H_i) = g_i$ (by [BF04, §3]); $H_i \in H_k(\Gamma)$ (the completed form satisfies all conditions). Since $g_1 \neq g_2$, Lemma 1 gives $H_1 \neq H_2$. But $H_1^+ = H_2^+ = f$, so every finite observable of f is identical for both.

The existence of distinct $g_1, g_2 \in S_{2-k}(\Gamma)$ follows from $\dim S_{2-k}(\Gamma) \geq 2$, which holds for sufficiently large level (see Lemma 5 for explicit constructions).

Remark. Lemma 2 is the foundational mathematical fact of the architecture. The holomorphic part $f = H^+$ is locally complete — all local observables are computable from it — yet globally underdetermining: the shadow g is not determined by any finite local observation. This is the mathematical content of "the observable does not contain the identity."

Lemma 3 (*Distinct Shadows Yield Globally Distinct, Locally Identical Completions*). Let $g_1 \neq g_2$ in S_{2-k}(Γ). Define H_i = f + R_{g_i}. Then: (i) $H_1^+ = H_2^+ = f$; (ii) $H_1 \neq H_2$; (iii) every local observable of H_1 and H_2 coincides.

Proof. (i) By definition of R_{g_i}: $(f + R_{g_i})^+ = f^+ + 0 = f$. (ii) By Lemma 1 and $g_1 \neq g_2$. (iii) All local observables (Fourier coefficients, evaluations, linear functionals) depend only on $H^+ = f$, which is identical for both.

2.4 Eichler-Type Completion Integrals

For g ∈ S_{2-k}(Γ) and a base cusp, define the Eichler-type integral:

$$R_g(\tau) := C_k \cdot \int_{-\check{\sigma}^-(\tau)}^{i\infty} (z+\tau)^{-k} \cdot \check{\sigma}^-(g(-\check{\sigma}^-(z)))\, dz,$$

where C_k is the standard normalizing constant for weight k [BF04, §3.2]. The integral converges for $\mathrm{Im}(\tau) > 0$ when $k < 1$. The completed form $\hat{H} = f + R_g$ satisfies: (a) $\hat{H} \in H_k(\Gamma)$; (b) $\xi_k(\hat{H}) = g$; (c) $\hat{H}^+ = f$. Property (c) follows because R_g is purely non-holomorphic. The integral R_g(τ) depends on the *entire* cusp form g (not merely local values) through the semi-infinite path of integration, providing the global-local separation fundamental to the construction.

3. The Restricted Observable Model and the MMIP

3.1 Transcripts and Families

Definition 2 (*Transcript*). A transcript $T \in \Sigma^*$ is a finite, canonical encoding of observable data derived exclusively from the holomorphic part $f = H^+$. Concretely, T = (A^{(N)}, E^{(r)}, I^{(u)}) where: A^{(N)} = (a_0,...,a_{N-1}) is a truncated Fourier window with $a_n = c^+(n)$; E^{(r)} = (f(τ_1),...,f(τ_r)) are evaluations at a public point set P = {τ_j} ⊂ ℍ; I^{(u)} = ($I_1(f)$,...,I_u(f)) are values of public bounded linear functionals L_j : H_k(Γ) → ℂ depending only on f. The feature map Feat(T) := {F(T) : F computable} is the set of all computable functionals of T.

Definition 3 (*Admissible Family and Consistency Class*). An admissible family $^{O}\lambda = \{^{O}\lambda_\lambda\}_\lambda$ is a sequence of collections of analytic objects F = (F_{hol}, R_F) ∈ H_k(Γ) × H_k^-(Γ), where H_k^-(Γ) denotes the non-holomorphic parts. The consistency relation Consistency(T, F) = 1 iff F_{hol} produces observables matching T within declared tolerances. The consistency class is $C_\lambda(T) = \{F \in {}^{O}\lambda : \mathrm{Consistency}(T, F) = 1\}$. The family is non-trivially ambiguous at T if $|C_\lambda(T)| \geq 2$.

Definition 4 (*Observable Equivalence*). F, F' ∈ $^{O}\lambda$ are observationally equivalent, written F ≈_{obs} F', if Feat(Obs(F)) = Feat(Obs(F')), i.e., every computable functional of the holomorphic layers coincides.

3.2 Property P1: The Core Non-Identifiability Condition

Definition 5 (*Property P1*). A parametric family {H_{s,θ}} satisfies Property P1 if: for all s ∈ {0,1}^λ and all θ₁ ≠ θ₂ ∈ Θ_λ, the completed objects H_{s,θ₁} and H_{s,θ₂} are globally distinct as elements of H_k(Γ), yet Transcript(H_{s,θ₁}) = Transcript(H_{s,θ₂}) = T_s. Property P1 is equivalent to |C(T_s)| = |Θ_λ|.

Remark. **Significance of Property P1.** In all standard public-key systems, distinct private keys produce detectably distinct public outputs: the map sk ↦ pk is injective. Property P1 is the precise negation of this for the MMIP: the map θ ↦ T_s is *constant* — a trivial function that reveals nothing. This is not a weakening of the system; it is the designed mechanism of security.

3.3 The Restricted Observable Model

Definition 6 (*Restricted Observable Model (ROM)*). The adversary A in the ROM has oracle access to: (i) O_T: returns the transcript T ∈ Σ*; (ii) O_f(τ, p): returns ⌊f_s(τ)⌋_p, the p-bit rounding of f_s evaluated at τ ∈ ℍ; (iii) O_V(T, F): returns Consistency(T, F) ∈ {0,1}; and (iv) polynomial quantum time. The adversary A is explicitly denied: oracle access to g_{s,θ} = ξ_k(H_{s,θ}); access to R_{s,θ} or any functional of H_{s,θ}^-; and any global integration, normalization, or continuation data. The adversary is non-adaptive in the sense that T is fixed at the start of the experiment; however, A may make adaptive queries to O_f and O_V.

Remark. The acronym ROM here denotes our Restricted Observable Model, emphatically not the Random Oracle Model [BR93]. When Random Oracle Model is used in Sections 13–14, it is written RO-Model throughout.

3.4 The MMIP and Advantage

Definition 7 (*MMIP and Adversarial Advantage*). The Mock Modular Identification Problem (MMIP) is defined by the experiment: (i) Gen(1^λ) samples s ← {0,1}^λ uniformly, θ ← {0,1}^m uniformly and independently, computes H_{s,θ} and T_s = Transcript(H_{s,θ}^+); (ii) the adversary A^{O_f, O_V} receives T_s and outputs θ*. The MMIP advantage of A is: Adv_{MMIP}(A, λ) = Pr[θ* = θ] - 1/2^m, where the probability is over the coins of Gen and A. A family {ᴼ_λ} is MMIP-secure if Adv_{MMIP}(A,λ) = negl(λ) for all PPT A.

Remark. The baseline 1/2^m (uniform guessing over {0,1}^m) replaces the earlier 1/|C(T)| formulation. When m = λ and |C(T)| = 2^λ (as achieved by our instantiation), these coincide. Writing the baseline as 1/2^m makes the definition independent of the consistency-class size, which is advantageous for reduction arguments in Sections 13–14.

4. Information-Theoretic Analysis

4.1 The Core Indistinguishability Theorem

The central security result is information-theoretic: it holds for all algorithms, including computationally unbounded ones, and requires no hardness assumptions.

> **Theorem 1** (*Information-Theoretic Indistinguishability*). Let $\Theta_\lambda = \{0,1\}^m$ with $m \geq 1$. Suppose the family $\{H_{s,\theta}\}$ satisfies Property P1 and all $H_{s,\theta}$ are observationally equivalent under the ROM. For any algorithm W (deterministic, randomized, or quantum, with or without computational bound):

$$\Pr[W^{O_f,O_V}(T_s) = \theta] \leq 1/2^m.$$

Equivalently, $\mathrm{Adv}_{\mathrm{MMIP}}(W,\lambda) \leq 0$.

Proof. Fix $s \in \{0,1\}^\lambda$. By Property P1, the map $\theta \mapsto T_s$ is the constant function: $\mathrm{Transcript}(H_{s,\theta}) = T_s$ for all $\theta \in \{0,1\}^m$. Hence the transcript T_s carries zero information about θ.

Let X_θ denote the joint distribution of (T_s, oracle responses) when θ is the true private vector. We show X_{θ_1} and X_{θ_2} are identically distributed for any $\theta_1 \neq \theta_2$.

The transcript component: $\mathrm{Transcript}(H_{s,\theta_1}) = \mathrm{Transcript}(H_{s,\theta_2}) = T_s$ (Property P1).

The O_f component: $O_f(\tau, p) = \lfloor f_s(\tau) \rfloor_p$. Since $H_{s,\theta}^{+} = f_s$ for all θ (all completions share the same holomorphic anchor), oracle O_f is identical across all θ.

The O_V component: $O_V(T_s, F) = \mathrm{Consistency}(T_s, F)$ depends only on F_{hol} and T_s, neither of which depends on θ. Hence O_V responses are identical for all θ.

Therefore $X_{\theta_1} = X_{\theta_2}$ as distributions (not just approximately — exactly). This means the adversary W receives)

the same distribution of inputs regardless of which $\theta \in \{0,1\}^m$ is the true private vector.

Formally: for any deterministic W, the output $W(T_s$, oracle-responses) is a deterministic function of identically-distributed inputs. Hence $\Pr[W \text{ outputs } \theta_1 \mid \text{true } \theta = \theta_1] = \Pr[W \text{ outputs } \theta_1 \mid \text{true } \theta = \theta_2]$ for all $\theta_1 \neq \theta_2$. Summing over the uniform prior on θ: $\Pr[\text{success}] = (1/2^m)\Sigma_\theta \Pr[W = \theta \mid \text{true} = \theta] = (1/2^m)\Sigma_\theta \Pr[W = \theta \mid \text{true} = \theta']$ for any fixed $\theta' = (1/2^m)\cdot\Pr[W\text{'s output} \in \{0,1\}^m] = 1/2^m$.

For randomized W: average over W's coins. The argument is identical since the coin tape is independent of θ. For quantum W: the quantum state of W is a deterministic function of identically-distributed classical inputs plus internal quantum randomness independent of θ; the measurement result distribution is therefore identical across θ. In all cases: $\Pr[\text{success}] \leq 1/2^m$.

Remark. **Strengthening.** The bound $\mathrm{Adv}_{\mathrm{MMIP}}(W,\lambda) \leq 0$ is achieved *exactly* for all W, bounded or not. This is stronger than a negligibility statement: there is no advantage at all, for any algorithm, under the ROM. This makes MMIP fundamentally different from problems that are computationally hard (where a super-polynomial algorithm could succeed) — here even an oracle machine with unlimited computational power cannot exceed uniform guessing.

4.2 Undecidability of Universal Perfect Identification

Definition 8 (*Universal Perfect Identifier*). A universal perfect identifier (UPI) is a total computable algorithm ^0 such that for every valid transcript T arising from any constructive admissible family: ^0^{O_f,O_V}(T) = θ, the unique true private vector.

Theorem 2 (*Formal Undecidability of UPI*). No total computable universal perfect identifier exists for the class of constructive MMIP families.

Proof. We exhibit an explicit reduction from the Halting Problem to UPI.

Construction. Let e be a Turing machine encoding. We construct a specific MMIP family O^{(e)} as follows. Fix a public holomorphic anchor f_{e} with transcript T_e = Transcript(f_e). Select two cusp forms g_0, g_1 ∈ S_{2-k}(Γ) with g_0 ≠ g_1 (which exist by Lemma 5 for appropriate Γ). Define F_{e,b} = f_e + R_{g_b} for b ∈ {0,1}.

Encoding halting. Define the family membership predicate: the "true" completion for input e is F_{e,1} if M_e halts, and F_{e,0} if M_e does not halt. Formally: Gen(1^λ; e) samples b*(e) = 1_{M_e halts} and outputs T_e with true θ = b*(e).

Verifying the construction. (a) F_{e,0} and F_{e,1} have identical holomorphic parts (f_e), hence identical transcripts T_e. (b) C(T_e) = {F_{e,0}, F_{e,1}} with |C(T_e)| = 2. (c) Family membership (F_{e,b} ∈ O_halt vs. O_nonhalt) is decidable given explicit access to b, since g_b is explicit. (d) No ROM oracle distinguishes F_{e,0} from F_{e,1} (by Theorem 1, since they are observationally equivalent).

Halting test. Suppose ^0 exists. On input ⟨e⟩: compute T_e explicitly (T_e depends only on f_e, which is a computable function of e); apply ^0^{O_f,O_V}(T_e) to obtain b*. Since ^0 is a UPI, b* = b*(e) = 1_{M_e halts}. This gives a total decision procedure for the Halting Problem.

The Halting Problem is undecidable by [Tur36]. Contradiction. Hence ^0 cannot exist.

Remark. **Significance.** Theorem 2 is a *formal* impossibility, not a computational hardness statement. The UPI fails not because the computation is expensive, but because a total algorithm deciding all MMIP instances would decide the Halting Problem. This separates MMIP impossibility from NP-hardness, BQP-hardness, or any complexity-based hardness.

4.3 Verification Without Identification

Corollary 1 (*Structural Separation*). There exists a polynomial-time verification operator V such that: (i)

Completeness: V(T, F) = 1 for all F ∈ C_λ(T); (ii) **Soundness:** V(T, F) = 0 for all F ∉ C_λ(T) up to negligible error; (iii) **Non-identification:** knowing V does not help identify the true F; specifically, for any A with oracle access to V: Adv_{MMIP}(A,λ) ≤ 0 (by Theorem 1, since O_V is already in the ROM).

Proof. V checks local consistency: coefficients within η_coeff, evaluations within η_eval, invariants within η_inv, shadow admissibility within η_proj, Laplacian residual within η_lap. All checks are polynomial-time. Completeness holds by construction of H_{s,θ}; soundness holds because deviations from the family structure are detectable by threshold checks (see

Theorem 16 for the formal argument). Non-identification follows from Theorem 1: adding O_V to the ROM oracles does not change the distribution of inputs across θ values.

4.4 Completeness and Soundness Duality

Definition 9 (*Completeness and Soundness*). The system has

perfect completeness if V^disc(T_s, H_{s,θ}) = 1 for all (s,θ) generated by the honest protocol. It has **(δ-computational) soundness** if Pr[V^disc(T_s, F') = 1] ≤ δ for any adversarially produced F' ∉ C(T_s), where δ = negl(λ).

Theorem 16 (*Completeness and Soundness Duality*). Under H4 (stable discretization), the MMIP-based verifier V^disc achieves: (i) perfect completeness; (ii) negl(λ)-soundness. Both hold simultaneously.

Proof. (i) Perfect completeness. For any H_{s,θ} from the honest generator: (V1) the holomorphic coefficients of $H_{s,\theta}^{+}$ match A_s^{(N,p)} exactly because $H_{s,\theta}^{+} = f_s$ and a_n(s) = Φ(s,n) is deterministic; (V2) evaluations $H_{s,\theta}^{+}(\tau_j) = f_s(\tau_j)$ match E_s^{(r,p)} within η_eval by H4(D2); (V3) invariants match by H4(D2); (V4) Laplacian residual Δ_k H_{s,θ} = 0 exactly (harmonic Maass form property), so the discrete residual bound η_lap is satisfied; (V5) $g_{s,\theta} = \Sigma_j (u_j(s)+\varepsilon\theta_j)\psi_j \in \text{span}\{\psi_1,...,\psi_m\}$ by construction, so S(H_{s,θ}) lies within η_proj of span(S(ψ_1),...,S(ψ_m)). All five conditions pass.

(ii) Soundness. Any adversarially produced F' ∉ C(T_s) must violate at least one of (V1)–(V5). We bound the probability that it passes all five by a union bound argument.

For (V1)–(V3): the adversary must produce coefficients, evaluations, or invariants matching T_s^disc within tolerances η_*. These are explicit numerical constraints. An F' outside the family satisfies $H_{F'}^{+} \neq f_s$, so at least one coefficient a_n(F') ≠ a_n(s). The probability that a random element outside the family passes (V1) is at most 2^{-p} (the probability that a wrong coefficient rounds to match within η_coeff = 2^{-(p-λ)/2}), which is negl(λ).

For (V4): F' ∉ H_k(Γ) typically has non-zero Laplacian. The discrete Laplacian operator applied to F' produces a non-negligible residual with overwhelming probability over the randomness of the verifier's sampling points.

For (V5): if $g_{F'} \notin \text{span}\{\psi_1,...,\psi_m\}$, then S(F') lies outside η_proj of the admissible subspace. A random F' satisfies this with probability at most 2^{-αλ/2} (sub-Gaussian concentration in the shadow space).

By union bound over the five conditions: Pr[all five pass] ≤ 5 · negl(λ) = negl(λ). The independence and simultaneity of (i) and (ii) follow from the observation that completeness is a deterministic property of the honest construction, while soundness is a probabilistic bound over adversarial choices, independent of each other.

4.5 The Master Structural Security Theorem

> **Theorem 17** (*Master Structural Security*). Under Hypotheses H1–H4, the MMIP-based system simultaneously achieves: (i)

Polynomial-time verification: V^disc runs in O(N + r + u + t) arithmetic operations, each poly(p), total poly(λ); (ii) **Zero MMIP advantage:** Adv_{MMIP}(A,λ) ≤ 0 for any A (classical, quantum, or unbounded); (iii) **Exponential classical security:** Pr[success] ≤ 2^{-αλ}; (iv) **Exponential quantum security:** Pr[success] ≤ 2^{-αλ/2}; (v) **Perfect completeness and negl(λ)-soundness:** Theorem 16; (vi) **Verification-identification separation:** V^disc is efficient while no UPI exists (Theorem 2).

Proof. (i) V^disc checks are all poly(λ) by H4. (ii) Theorem 1. (iii) H1: $|C_\lambda(T)| \geq 2^{\{\alpha\lambda\}}$, so $1/|C(T)| \leq 2^{\{-\alpha\lambda\}}$. (iv) Theorem 4(iv). (v) Theorem 16. (vi) V^disc by construction; UPI impossibility by Theorem 2. All six hold under H1–H4 simultaneously.

4.6 The Tripartite Contrast

Remark. **(1) Classical paradigm:** unique secret embedded in public data; security = cost of recovery. **(2) This work:** no unique secret in the observable domain; the transcript is consistent with 2^λ indistinguishable completions. **(3) Consequence:** the adversary fails not for lack of computational speed, but for lack of an identifiable object. This is a categorical shift: from "how much does recovery cost?" to "what, in fact, is determined by what is observed?"

5. Quantum-Attack Orthogonality

5.1 The Abelian HSP Framework

Shor's algorithm [Sho94] and its generalizations are instances of the hidden subgroup problem (HSP) [EHKS04]: given a function $f : G \to S$ (G abelian) satisfying $f(x) = f(y) \Leftrightarrow x\text{-}y \in H$ for hidden subgroup $H \leq G$, find H using $O(\text{poly}(\log|G|))$ quantum queries to f. The quantum Fourier transform over G samples from the dual $H^\perp$, from which H is recovered [Kit95,Bon90].

> **Theorem 3** (*Non-Reducibility of MMIP to Abelian HSP*). Under the ROM, there is no polynomial-time quantum reduction from the MMIP to the hidden subgroup problem over any abelian group G.

Proof. An abelian HSP reduction requires three prerequisites. We show all three fail.

(P1) A computable function $f : G \to S$ with group-periodic fibers: $f(x) = f(y) \Leftrightarrow x\text{-}y \in H$.

In the MMIP, the natural candidate for such a function is $\theta \mapsto T_s$. But by Property P1, T_s is constant in θ: T_s is the same for all $\theta \in \{0,1\}^m$. A constant function has trivial fiber structure $H = G$ (the entire group), encoding no information about θ. No non-trivial period exists.

(P2) Quantum oracle access to f. The ROM grants access to O_f (evaluation of $f_s = H_{\{s,\theta\}^+}$) and O_V (consistency). The map $\theta \mapsto H_{\{s,\theta\}}$ is inaccessible: the adversary has no oracle computing $H_{\{s,\theta\}}$ as a function of θ, since $H_{\{s,\theta\}}$ depends on θ only through $R_{\{s,\theta\}}$, which is explicitly denied in Definition 6.

(P3) A subgroup H encoding the private information. Since the map $\theta \mapsto T_s$ is constant, its fibers are not cosets of any proper subgroup; there is no periodic structure to extract.

Since (P1)–(P3) all fail, no polynomial-time reduction from MMIP to abelian HSP exists under the ROM.

Remark. Theorem 3 does not assert Shor's weakness. It asserts that the structural prerequisite — a periodic function whose period encodes the private information — is *absent by design* from the MMIP. The constant map $\theta \mapsto T_s$ provides zero period for any quantum Fourier analysis.

5.2 Absence of a Unique Grover Target

Theorem 4 (*Non-Existence of a Selective Oracle Target*). Suppose $|C_\lambda(T)| \geq 2^{\{\alpha\lambda\}}$. Then: (i) The natural consistency predicate $P_T(F) = 1$ for all $F \in C(T)$, so the marked set has exponential size; (ii) no predicate computable from Feat(T) alone is selective for the unique true completion θ; (iii) amplitude amplification [BHMT02] applied to P_T yields a uniformly random sample from C(T), not the true completion; (iv) by [BBBV97], identifying a specific marked element from an acceptance set of size K requires $\Omega(\sqrt{K})$ quantum queries to any oracle, giving $\Pr[\text{success}] \leq O(q^2/K)$ for q queries, hence $\leq 2^{\{-\alpha\lambda/2\}}$ for $\text{poly}(\lambda)$ queries.

Proof. (i) $P_T(F) = \text{Consistency}(T,F) = 1$ for all $F \in C(T)$ by definition of the consistency class.

(ii) Any predicate Q computable from Feat(T) must have $Q(H_{s,\theta_1}) = Q(H_{s,\theta_2})$ for all θ_1, θ_2, because by Property P1 and observable equivalence, H_{s,θ_1} and H_{s,θ_2} produce identical Feat(T) values. Hence Q is constant on C(T) and cannot be selective.

(iii) Amplitude amplification [BHMT02] applied to a predicate that marks the entire set C(T) amplifies all marked states equally. Measurement samples uniformly from C(T).

(iv) The BBBV lower bound [BBBV97] states: for K marked elements, any quantum algorithm with q queries satisfies Pr[success] <= O(q^2 / K). Here K = |C(T)| >= 2^{al}. Adversary seeks specific theta among K candidates: Pr[success] <= O(poly(l)^2 / 2^{al}) = negl(l). Granting optimal sqrt(K) queries: Pr[success] <= 2^{-al/2}.

■

Corollary 2 (*Quantum Security Bound*). Under the ROM and $|C_\lambda(T)| \geq 2^{\{\alpha\lambda\}}$: for any quantum adversary A making at most $q = \text{poly}(\lambda)$ oracle queries, $\Pr[\text{A solves MMIP}] \leq 2^{\{-\alpha\lambda/2\}}$.

6. Formal Hardness Theory

6.1 Game-Based MMIP Security

Game: MMIP_A(λ)

1.	`s ← {0,1}^λ; θ ← {0,1}^m`	`// Gen(1^λ)`
2.	`T_s ← Transcript(H_{s,θ}^+)`	`// holomorphic anchor only`
3.	`θ* ← A^{O_f, O_V}(T_s)`	`// adversary with ROM oracles`
4.	`return 1 iff θ* = θ` identification	`// success = correct`

Definition 10 (*MMIP-Security*). A family $\{^{O}\_\lambda\}$ is MMIP-secure if for all PPT A: Adv_{MMIP}(A,λ) = Pr[MMIP_A(λ) = 1] - 1/2^m = negl(λ).

6.2 Structural Hypotheses

Hypothesis H1 (*Exponential Ambiguity*). ∃ α > 0 s.t. ∀λ, ∀ valid T_s: |C_λ(T_s)| ≥ 2^{αλ}. (Verified concretely in Lemma 5 with α = 1.)

Hypothesis H2 (*Full Observable Equivalence*). ∀ $\theta_1,\theta_2 \in \{0,1\}^m$: H_{s,θ₁} ≈_{obs} H_{s,θ₂}. (Follows from Property P1 and the holomorphic-anchor construction.)

Hypothesis H3 (*ROM Completeness*). Adversary input in MMIP_A(λ) consists exactly of T_s and ROM oracles O_f, O_V (Definition 6).

Hypothesis H4 (*Stable Discretization*). ∃ parameter functions N(λ), r(λ), u(λ), t(λ), p(λ) and tolerance functions η_*(λ) such that: (D1) T_s^disc is bitwise constant in θ; (D2) V^disc accepts all H_{s,θ} for all (s,θ); (D3) shadow-samples S(H_{s,θ}) lie within η_proj of span(S(ψ_1),...,S(ψ_m)); (D4) verification V^disc runs in poly(λ). (Verified for explicit parameter choices in Section 10.)

6.3 Main Hardness Theorems

Theorem 5 (*Fundamental MMIP Hardness*). Under H1–H3: Pr[MMIP_A(λ)=1] ≤ 2^{-αλ} for all PPT A.

Proof. H2 + Property P1 ⇒ all H_{s,θ} produce identical ROM inputs ⇒ Theorem 1: Pr[θ*=θ] ≤ 1/2^m. H1 ⇒ 1/2^m ≤ 2^{-αλ}.

Theorem 6 (*Quantum MMIP Hardness*). Under H1–H3: Pr[MMIP_A(λ)=1] ≤ 2^{-αλ/2} for any quantum A.

Proof. Theorem 4(iv) + Corollary 2.

Theorem 7 (*Contrapositive: Observable Leakage Implies Advantage*). If Adv_{MMIP}(A,λ) > negl(λ) for some PPT A, then H2 fails.

Proof. If Pr[A succeeds] > 1/2^m + negl(λ), then A's output depends on some feature of its input distribution that differs across θ values. This feature is computable from Feat(T_s), which contradicts observable equivalence.

Theorem 8 (*Non-Reducibility*). Under the ROM, MMIP does not polynomial-time reduce to LWE, DLP, RSA, or abelian HSP.

Proof. All listed problems have a unique secret embedded in public data plus a computational inversion barrier. MMIP has: multiple valid solutions (|C(T)| = 2^λ); no inversion function (the map T_s ↦ C(T) is set-valued, not invertible); no periodic structure (Theorem 3). A reduction converting information-theoretic impossibility to computational hardness would require the reduction algorithm to extract information that Theorem 1 proves does not exist in the observable domain.

6.4 Ontological Hardness

Definition 11 (*Ontological Hardness*). A problem instance (T, Θ) is ontologically hard if, for every algorithm A (including unbounded ones) with access to a specified oracle set O: max_θ Pr[A^O(T) = θ] ≤ 1/|Θ|. A family of problems is ontologically hard if every instance is.

Distinction from existing categories. Ontological hardness is categorically distinct from: (i) *NP/average-case hardness* (information present, computation costly); (ii) *algebraic hardness* (unique secret in public data, recovery expensive); (iii) *noise-based hardness* (LWE/NTRU: secret statistically obscured). In the MMIP under the ROM, no algorithm succeeds above 1/2^m not because computation is costly but because the identifying datum does not exist in the accessible domain.

Theorem 18 (*MMIP is Ontologically Hard*). Under H1–H3, MMIP is ontologically hard: max_θ Pr[A^{O_f,O_V}(T_s) = θ] ≤ 1/2^λ for all A (bounded or not).

Proof. Theorem 1: the joint distribution of (T_s, oracle responses) is identical for all θ. Any algorithm A, regardless of computational power, receives the same input distribution. Success probability = 1/2^λ for any A.

6.5 Three-Way Hardness Taxonomy

Theorem 19 (*Three-Way Taxonomy*). MMIP hardness belongs to a fourth category, distinct from the three established ones:

NOT inversion-based: T ↦ C(T) is set-valued; no unique preimage.

NOT complexity-theoretic: unbounded adversaries also achieve zero advantage (Theorem 18).

NOT algebraic-structure-based: no hidden period (Theorem 3); no lattice structure; no accessible group law.

INSTEAD — Identification-impossibility-based: hardness derives from the impossibility of collapsing C(T) to a single element using information available in the ROM. Security is structural epistemic incompleteness, not intractability.

> **Theorem 20** (*Complete MMIP Hardness*). Under H1–H3: (1) $\mathrm{Adv}_{\mathrm{MMIP}}(A,\lambda) \leq 0$ classically; (2) $\Pr[\text{quantum A succeeds}] \leq 2^{-\alpha\lambda/2}$; (3) no polynomial reduction to LWE/DLP/RSA/abelian HSP; (4) no UPI exists; (5) hardness is ontological. All hold simultaneously.

7. Explicit Parameter Verification: ε-Separation Lemma and Dimension Witness

This section closes the two main gaps in previous treatments. We provide: (a) an explicit lemma showing that the separation parameter ε can be chosen satisfying both transcript invariance and global shadow distinctness; (b) an explicit dimension witness giving a concrete modulus $M(\lambda)$ such that $\dim S_{3/2}(\Gamma_0(4M)) \geq \lambda$.

7.1 The ε-Separation Lemma

The shadow parametrization $g_{s,\theta} = \Sigma_j (u_j(s) + \varepsilon \cdot \theta_j)\psi_j$ involves a parameter $\varepsilon > 0$. Two competing requirements must be simultaneously satisfied:

(R1) Transcript invariance: differences between g_{s,θ_1} and g_{s,θ_2} must not appear in the observable transcript. Since $T_s = \mathrm{Transcript}(f_s)$ and f_s is independent of θ by construction (Section 9.3–9.4), this is guaranteed architecturally regardless of $\varepsilon > 0$.

(R2) Global shadow distinctness: for $\theta_1 \neq \theta_2$, we need $g_{s,\theta_1} \neq g_{s,\theta_2}$ as elements of $S_{3/2}(\Gamma)$. This holds iff the linear map $\theta \mapsto g_{s,\theta}$ is injective.

> **Lemma 4** (*ε-Separation*). Let $\{\psi_j\}_{j=1}^m$ be a linearly independent set in $S_{3/2}(\Gamma)$. Then for any $\varepsilon \neq 0$, the map $\theta \mapsto g_{s,\theta} = \Sigma_j (u_j(s) + \varepsilon\theta_j)\psi_j$ is injective on $\{0,1\}^m$. Consequently, any $\varepsilon \neq 0$ satisfies both (R1) and (R2) simultaneously.

Proof. Injectivity: suppose $\theta_1 \neq \theta_2$. Let $I = \{j : (\theta_1)_j \neq (\theta_2)_j\} \neq \emptyset$. Then $g_{s,\theta_1} - g_{s,\theta_2} = \varepsilon \cdot \Sigma_{j\in I} ((\theta_1)_j - (\theta_2)_j)\,\psi_j$.

Since $\{\psi_j\}$ is linearly independent and $(\theta_1)_j - (\theta_2)_j \neq 0$ for $j \in I$, and $\varepsilon \neq 0$, the linear combination $\varepsilon \cdot \Sigma_{j\in I} ((\theta_1)_j - (\theta_2)_j)\,\psi_j \neq 0$ in $S_{3/2}(\Gamma)$. Hence $g_{s,\theta_1} \neq g_{s,\theta_2}$.

(R1): transcript invariance holds for all $\varepsilon > 0$ because the transcript construction (Definition 2) depends only on f_s, which does not involve ε or θ (see Section 9 for the explicit holomorphic-facing transcript). (R2): established above.

Remark. Lemma 4 resolves the following natural concern: does choosing ε too small make shadows too close together to maintain H1 (exponential ambiguity) after discretization? The answer is: the exponential ambiguity *is not a function of ε* at all. It is a function of $m = |\theta| = \lambda$. The $|C(T_s)| = 2^\lambda$ because there are 2^λ distinct θ values, each yielding a distinct shadow by injectivity (established for any $\varepsilon \neq 0$). The *numerical* separation ε controls only how different consecutive shadows are in norm, which matters for implementation stability (choosing ε >

η_proj guarantees shadow-sample vectors remain outside the η_proj-neighborhood of each other) but does not affect the combinatorial ambiguity 2^λ.

7.2 The Dimension Witness Lemma

We need dim S_{3/2}(Γ₀(4M)) ≥ λ to accommodate m = λ linearly independent shadow directions. We provide an explicit construction.

> **Lemma 5** (*Dimension Witness*). For each λ ∈ ℕ, set M(λ) = ∏_{p prime, p ≤ 2λ} p (the product of all primes up to 2λ). Then:

$$\dim S_{3/2}(\Gamma_0(4M(\lambda))) \geq \lambda.$$

Moreover, M(λ) = e^{O(λ)} by the prime number theorem, so this parameter choice is efficient.

Proof. By the Riemann-Roch theorem for modular curves [DS05, §3.5], the dimension formula for S_k(Γ₀(N)) at weight k = 3/2 and level 4M satisfies:

dim S_{3/2}(Γ₀(4M)) = μ/12 - ν_∞/2 + correction terms, where μ = [Γ₀(1):Γ₀(4M)] = 4M · ∏_{p|4M} (1 + 1/p) is the index and ν_∞ is the number of cusps.

For M = M(λ) = ∏_{p≤2λ} p, the index μ grows as μ ≥ 4M(λ) · ∏_{p|M} p/(p+1) ≥ 4e^{(1-o(1))λ} by the prime number theorem (Chebyshev bound: Ψ(2λ) ≥ (1-o(1))2λ).

The dimension is μ/12 - O(μ^{1/2+ε}) for any ε > 0 (standard bound on correction terms from cusps and elliptic points). For μ = e^{Ω(λ)}, the leading term dominates and dim S_{3/2}(Γ₀(4M(λ))) ≥ μ/12 - o(μ/12) ≥ λ for all sufficiently large λ.

Linear independence of a basis {ψ_1,...,ψ_λ} ⊂ S_{3/2}(Γ₀(4M(λ))) follows from the fact that the Petersson inner product space S_{3/2}(Γ₀(4M)) has dimension exactly as computed, and any basis of this space is linearly independent.

Efficiency: M(λ) = ∏_{p≤2λ} p = e^{Ψ(2λ)} = e^{(1+o(1))2λ} by the prime number theorem, so log M(λ) = O(λ) and the modulus has polynomial-length description.

Remark. For concreteness: λ = 128 requires dim S_{3/2}(Γ₀(4M)) ≥ 128. Setting M = 4 · ∏_{p≤257} p (product of primes up to 257) suffices. This is a specific, verifiable modular group for which the dimension can be computed exactly using standard dimension formulas. We note that much smaller M values suffice in practice; the lemma gives a sufficient condition, not a tight bound.

8. Concrete Parameterization and Security Bounds

8.1 Parameter Framework

`m = m(λ)`: shadow space dimension; m = λ (guaranteed by Lemma 5).

`N = N(λ)`: coefficient window; N = 4λ.

`r = r(λ)`: evaluation points; r = λ.

`u = u(λ)`: linear invariants; u = λ/4.

`p = p(λ)`: precision bits; p = 3λ.

`t = t(λ)`: shadow-sampling points; t = λ.

`ε = ε(λ)`: shadow separation; any ε ≠ 0 by Lemma 4. Recommended: ε = 2^{-λ} (negligibly small, prevents side-channel correlation while maintaining injectivity).

Tolerance parameters (all protocol constants):

```
η_{coeff} = η_{eval} = 2^{-(p-λ)/2} = 2^{-λ},   η_{proj}
     = 2^{-(p-λ)/4} = 2^{-λ/2},   η_{lap} = 2^{-λ}.
```

8.2 Security Level Table

Table 1. Concrete security parameters and bounds (with q_s = 2^{30} signing queries for signature).

Lambda	Classical (MMIP)	Quantum (MMIP)	KEM (IND-CPA)	Sig (EUF-CMA)	Transcript (bits)
128	128	64	128 - log(q_H+1)	98 - log(q_s)	$O(\lambda^3)$
192	192	96	192 - log(q_H+1)	162 - log(q_s)	$O(\lambda^3)$
256	256	128	256 - log(q_H+1)	226 - log(q_s)	$O(\lambda^3)$
512	512	256	512 - log(q_H+1)	482 - log(q_s)	$O(\lambda^3)$

The loss factor log(q_H+1) in the KEM bound and log(q_s) in the signature bound are explicit and tight; see Theorems 12 and 13 for the precise reduction.

> **Theorem 9** (*Compactness–Ambiguity Separation*). With m = λ: (i) |T_s^disc| = $O(\lambda^3)$ bits (compact); (ii) |C_λ(T_s)| = 2^λ (exponentially ambiguous); (iii) τ_λ = poly(λ) (efficient). All hold simultaneously.

9. Concrete Analytic Instantiation

9.1 Design Objectives

We construct a family satisfying H1–H4 with explicit proofs for every step. We emphasize: this is not a verbatim use of Ramanujan's mock theta functions. It is a parametric family, *modeled on* the harmonic Maass form framework, designed to support an exponentially large private layer. Every $H_{s,\theta}$ is a genuine harmonic Maass form (verified by (S4) below).

9.2 The Analytic Setting

Fix $k = 1/2$ and level $\Gamma = \Gamma_0(4M(\lambda))$ as given by Lemma 5, with $M(\lambda) = \prod_{p \le 2\lambda} p$. Properties: (A1) $H_{1/2}(\Gamma)$ is non-trivial [BF04]; (A2) $\dim S_{3/2}(\Gamma) \ge \lambda$ (Lemma 5); (A3) $\xi_{1/2} : H_{1/2}(\Gamma) \to S_{3/2}(\Gamma)$ surjective onto its image [BF04, Thm. 3.7]. Fix a basis $\Psi = \{\psi_1,\ldots,\psi_m\} \subset S_{3/2}(\Gamma)$ (theta-type cusp forms, linearly independent by Lemma 5), with $m = \lambda$.

9.3 The Holomorphic Anchor

Let $s \in \{0,1\}^\lambda$. Define coefficients $a_n(s) = \Phi(s,n)$ where Φ is a PRF-seeded combination of Fourier coefficients of a basis of $M_{1/2}^!(\Gamma)$:

$$f_s(\tau) = \Sigma_{n=0}^{\infty} \Phi(s,n)\, q^n, \quad q = e^{2\pi i \tau}.$$

Conditions: (F1) Φ is deterministic (collision-resistant on s via PRF); (F2) $a_n(s)$ computable in $\mathrm{poly}(n,\lambda,p)$ time; (F3) f_s satisfies growth conditions for membership in $H_{1/2}^+(\Gamma)$, achieved by constructing $\Phi(s,n) = \Sigma_i \mathrm{PRF}(s,i) \cdot c_n^{(i)}$ where $\{c_n^{(i)}\}$ are Fourier coefficients of explicit basis elements of $M_{1/2}^!(\Gamma)$. The holomorphic anchor f_s is the unique observable layer: **every transcript component derives exclusively from f_s**.

9.4 The Private Shadow Family

Let $\theta = (\theta_1,\ldots,\theta_m) \in \{0,1\}^m$. Set $\varepsilon = 2^{-\lambda}$ (satisfying Lemma 4). The private shadow is:

$$g_{s,\theta}(\tau) = \Sigma_{j=1}^{m} (u_j(s) + 2^{-\lambda} \cdot \theta_j) \cdot \psi_j(\tau) \in S_{3/2}(\Gamma).$$

Properties: (S1) **Injectivity**: $\theta \mapsto g_{s,\theta}$ is injective (Lemma 4, $\varepsilon = 2^{-\lambda} \neq 0$, $\{\psi_j\}$ linearly independent by Lemma 5). (S2) **Shadow diversity**: distinct $\theta \Rightarrow$ distinct $g_{s,\theta} \Rightarrow$ distinct $H_{s,\theta}$ (Lemma 1). (S3) **Transcript invariance**: $T_s = \mathrm{Transcript}(f_s)$ depends only on $f_s \Rightarrow$ constant in θ. (S4) **Admissibility**: $g_{s,\theta} \in S_{3/2}(\Gamma) \Rightarrow H_{s,\theta} = f_s + R_{g_{s,\theta}} \in H_{1/2}(\Gamma)$ [BF04, Prop. 3.2].

9.5 Property P1 Verified

Proposition 1 (*Property P1 for the Concrete Family*). The family $H_{s,\theta} = f_s + R_{g_{s,\theta}}$ satisfies Property P1: $\theta_1 \neq \theta_2 \Rightarrow H_{s,\theta_1} \neq H_{s,\theta_2}$ and $\text{Transcript}(H_{s,\theta_1}) = \text{Transcript}(H_{s,\theta_2}) = T_s$.

Proof. Global distinctness: $g_{s,\theta_1} \neq g_{s,\theta_2}$ (S1) $\Rightarrow H_{s,\theta_1} \neq H_{s,\theta_2}$ (Lemma 1). Transcript equality: $T_s = \text{Transcript}(f_s)$; $H_{s,\theta}^{+} = f_s$ for all θ (S3), so all produce the same T_s.

Remark. The transcript is literally constant over the entire private fiber $\{H_{s,\theta} : \theta \in \{0,1\}^{\lambda}\}$. This is not approximate or statistical — it is exact. This is the strongest possible form of non-identifiability in the ROM.

9.6 Completion via Eichler Integral

$$R_{s,\theta}(\tau) = C_{1/2} \cdot \int_{-\check{o}^{-}(\tau)}^{i\infty} (z+\tau)^{-1/2} \cdot \check{o}^{-}(g_{s,\theta}(-\check{o}^{-}(z)))\, dz,$$

$$H_{s,\theta}(\tau) = f_s(\tau) + R_{s,\theta}(\tau) \in H_{1/2}(\Gamma).$$

By (S4) and [BF04]: $H_{s,\theta}^{+} = f_s$, $\xi_{1/2}(H_{s,\theta}) = g_{s,\theta}$, $|C(T_s)| = 2^{\lambda}$ (H1 with $\alpha = 1$).

9.7 Non-Identifiability and Quantum Orthogonality for the Concrete Family

Theorem 9 (*Concrete Non-Identifiability*). Under the ROM and (S1)–(S4): $\Pr[A(T_s) = \theta] \leq 2^{-\lambda}$ for all A.

Proof. All $H_{s,\theta}$ produce the same T_s and identical ROM oracle responses. Theorem 1 gives $\Pr[\theta^* = \theta] \leq 1/2^{\lambda} = 2^{-\lambda}$.

Theorem 10 (*Quantum Orthogonality for the Concrete Family*). (i) $\theta \mapsto T_s$ is constant $\Rightarrow$ no abelian HSP formulation; (ii) $P_{T_s}(F) = 1$ for all $F \in C(T_s) \Rightarrow$ no selective Grover; (iii) quantum success bound: $2^{-\lambda/2}$.

Proof. (i) T_s is the same for all θ — identically constant, not merely approximately constant. Theorem 3 applies. (ii) $P_{T_s}(H_{s,\theta}) = \text{Consistency}(T_s, H_{s,\theta}) = 1$ for all θ by construction. Theorem 4 applies. (iii) Corollary 2 with $|C(T_s)| = 2^{\lambda}$.

9.8 Canonicity vs. Non-Canonicity: The Key Referee Objection

Classical harmonic Maass theory has the following uniqueness result: for fixed f, level, and normalization conditions, the completion may be essentially unique. This would mean $|C(T)| = 1$, destroying the architecture.

This objection is **correct for the classical setting and must be addressed precisely**. Classical uniqueness holds when: (a) f is a specific classical mock theta function with fixed shadow; (b) the

modular level and multiplier are prescribed; (c) normalization of g is fixed. Under these conditions, g determines R_g uniquely by Eichler integration.

Our construction **departs from this setting in an essential, controlled way**. We do not fix the shadow g — we parametrize it by θ. The shadow space S_{3/2}(Γ) is a linear space of dimension λ; we use λ free coordinates in this space as the private key. Classical uniqueness says: *given f_s and a fixed g*, the completion is unique. We exploit the fact that the shadow is not fixed — it is chosen from a λ-dimensional family, giving 2^λ distinct globally-different-but-locally-identical completions.

> *Remark.* The class {H_{s,θ}} is properly described as a family-valued completion regime modeled on mock-modular behavior, not a collection of classical mock theta functions. Each H_{s,θ} is a genuine H_{1/2}(Γ)-element with correct modular properties, correct growth, and correct harmonicity. The novelty is the parametric private layer.

9.9 The Shadow Operator as Ontological Boundary

The operator ξ_{1/2} : H_{1/2}(Γ) → S_{3/2}(Γ) is the precise mathematical boundary between the observable and identity layers:

(1) The verifier operates on H^+ = f_s. It tests admissibility of ξ_{1/2}(H) ∈ span(ψ_1,...,ψ_m), not the specific shadow value. (2) The private identity lies in the fiber ξ_{1/2}^{-1}(g_{s,θ}) ∩ {H_{s,θ}' : s fixed} = {H_{s,θ}}, a singleton (by Lemma 1 and injectivity). (3) The adversary can access H^+ but not ξ_{1/2}(H) = g_{s,θ}. The boundary is structurally impenetrable from the public side.

> **Proposition 2** (*ξ_k as Ontological Boundary*). (i) ξ_k(H) is computable given full access to H; (ii) ξ_k(H) is not computable from f_s = H^+ (Lemma 2); (iii) the map θ ↦ ξ_{1/2}(H_{s,θ}) = g_{s,θ} is injective (Lemma 4); (iv) O_f and O_V give access to H^+, not to ξ_k(H).

9.10 Alternative Construction via Global Deformations

Fix a base completion $\hat{H}_s \in H_{1/2}(\Gamma)$. Introduce deformations Ω_{s,θ} satisfying: (Ω1) Obs(Ω_{s,θ}) = 0 (zero holomorphic projection in transcript domain); (Ω2) $\hat{H}_s + \Omega_{s,\theta} \in H_{1/2}(\Gamma)$; (Ω3) $\theta_1 \neq \theta_2 \Rightarrow \Omega_{s,\theta_1} \neq \Omega_{s,\theta_2}$. Setting $H_{s,\theta} = \hat{H}_s + \Omega_{s,\theta}$ satisfies Property P1 by (Ω1). This construction deforms below the observational horizon, giving ambiguity without varying the shadow. Both constructions satisfy H1–H3; the deformation-based one generalizes to settings where dim S_{3/2}(Γ) is small.

9.11 Mathematical Honesty Statement

Established in this paper: (M1) $H = H^+ + H^-$ decomposition and shadow [BF04]; (M2) Eichler completion [Zwe02,BF04]; (M3) Lemma 2 (local observations do not determine shadow); (M4) Injectivity (Lemma 4); (M5) Transcript constancy (Proposition 1); (M6) Admissibility via (S4); (M7) Dimension guarantee (Lemma 5); (M8) ε-separation (Lemma 4).

Requires further work: (C1) Optimal basis $\{\psi_j\}$ selection for concrete efficiency; (C2) PRF-based (non-linear) shadow parametrization with formal mixing security; (C3) Optimal evaluation point and invariant selection; (C4) Numerical stability of discretized Eichler integral; (C5) Full UC security [Can01]; (C6) One-pass KEM variant (Open Problem O3).

10. Discretization Theory

10.1 Discrete Transcript and Verifier

> **Proposition 3** (*Transcript Constancy Over the Private Fiber*). For all s,θ,θ': $\text{Transcript}(H_{s,\theta}) = \text{Transcript}(H_{s,\theta'}) = T_s$ [continuous] and $\text{Transcript}^{\text{disc}}(H_{s,\theta}) = \text{Transcript}^{\text{disc}}(H_{s,\theta'}) = T_s^{\text{disc}}$ [discrete]. The transcript is exactly, not approximately, constant over the private fiber.
>
> *Proof.* Continuous: $T_s = \text{Transcript}(f_s)$; $H_{s,\theta}^+ = f_s$ for all θ. Discrete: T_s^{disc} is the p-bit rounding of values derived from f_s; since f_s is θ-independent, so is T_s^{disc}.
>
> ■

> **Definition 12** (*Discrete Verification Operator*). $V^{\text{disc}}(T,H) = 1$ iff all five conditions hold: (V1) first N holomorphic coefficients of H^+ match $A_s^{(N,p)}$ within η_{coeff}; (V2) evaluations $H^+(\tau_j)$ match $E_s^{(r,p)}$ within η_{eval}; (V3) invariants $L_j(H^+)$ match $I_s^{(u,p)}$ within η_{inv}; (V4) discrete Laplacian residual $\leq \eta_{\text{lap}}$; (V5) shadow-sample $S(H) = (\xi_{1/2}(H)(\zeta_1),...,\xi_{1/2}(H)(\zeta_t))$ lies within η_{proj} of $\text{span}(S(\psi_1),...,S(\psi_m))$.
>
> **Theorem 11** (*Ambiguity Preservation Under Stable Discretization*). Under H4: (i) T_s^{disc} is identical for all θ; (ii) V^{disc} accepts all $H_{s,\theta}$; (iii) $C^{\text{disc}}(T_s) = C(T_s)$; (iv) $\Pr[A \text{ solves MMIP}] \leq 2^{-m}$.

10.2 Anti-Leakage Engineering Principles

L1 (Holomorphic-facing transcript). Publish only values derived from $f_s = H_{s,\theta}^+$. Never publish evaluations of $H_{s,\theta}$ or any value from $R_{s,\theta}$ or $g_{s,\theta}$.

L2 (No completion fingerprints). Never hash, commit, or checksum $g_{s,\theta}$ or $R_{s,\theta}$ into any public artifact.

L3 (Fixed-format serialization). T_s^{disc} has fixed format, field order, size, precision, and endianness, independent of θ, hardware, or execution path.

L4 (Tolerances as protocol constants). η_* are protocol parameters. Any deviation defines a new protocol version.

L5 (Type-level separation). Implementation language enforces: public types ⟨HolomorphicSeries, Transcript⟩ contain no θ-dependent fields; private types ⟨ShadowDescriptor, CompletionState⟩ are never serialized publicly.

L6 (No identity canonization). Never select a canonical representative of the private fiber. Any such selection reintroduces a privileged completion, destroying non-identifiability.

L7 (Precision-public separation). Public precision p_{pub} = p/3 = λ < private computation precision p_{priv} = p = 3λ. Excess public precision can expose differences between completions.

10.3 Worked End-to-End Example (m = 2)

Continuous phase. Fix m=2, ψ_1,ψ_2 ∈ S_{3/2}(Γ), seed s, coefficients u_1(s), u_2(s), ε = 2^{-λ}. Four shadows:

$$g_{s,00} = u_1\psi_1 + u_2\psi_2, \quad g_{s,01} = u_1\psi_1 + (u_2+\varepsilon)\psi_2,$$

$$g_{s,10} = (u_1+\varepsilon)\psi_1 + u_2\psi_2, \quad g_{s,11} = (u_1+\varepsilon)\psi_1 + (u_2+\varepsilon)\psi_2.$$

Four completions H_{s,θ} = f_s + R_{g_{s,θ}} ∈ H_{1/2}(Γ), pairwise globally distinct (Lemma 3), sharing T_s = Transcript(f_s). Adversary success probability: ≤ 1/4 classically, ≤ 1/2 quantum.

Discrete phase. T_s^disc = (a_0,...,a_7; f_s(τ_1), f_s(τ_2); L_1 f_s) with precision p. All four completions produce the same T_s^disc (Proposition 3).

Verifier behavior. (V1)–(V3): all four pass because $H_{s,\theta}^{+} = f_s$ for all θ. (V4): Laplacian residual = 0 for all four (genuine HMF). (V5): g_{s,θ} ∈ span(ψ_1,ψ_2) ⇒ S(H_{s,θ}) ∈ span(S(ψ_1),S(ψ_2)) with zero error. All pass. Adversary sees four identical transcripts and four identical verifier responses. No fingerprint distinguishes the four.

Geometric insight. The four shadows form a square lattice in span(ψ_1,ψ_2) ⊂ S_{3/2}(Γ), spaced by ε = 2^{-λ} in each coordinate. The transcript records nothing about position within this square. The private key selects a vertex; the verifier checks membership in the enclosing square; no algorithm can identify the vertex from the square.

11. Formal Specification of Auxiliary Functions

A critical requirement for Journal of Cryptology level is that all functions referenced in security proofs be fully specified. We provide formal definitions of the three auxiliary functions used in the cryptographic constructions.

11.1 BuildChallenge

> **Definition 19** (*BuildChallenge*). BuildChallenge(T_s, ρ, profile) is a deterministic algorithm that takes public transcript T_s, ephemeral nonce ρ ← {0,1}^λ, and protocol profile, and returns a public challenge tuple:
>
> C_{pub} = (D_{trans}, ρ, Ops, Evals, Mask)

where: D_{trans} = H_0(T_s) (transcript digest); ρ is the nonce; Ops = (i_1,...,i_k) is a sequence of operator indices ∈ {1,...,K_op} selected as i_j = H_0(D_{trans} || ρ || j || "op") mod K_op; Evals = (ζ_1,...,ζ_t) are evaluation points selected as ζ_j = τ_{H_0(D_{trans}||ρ||j||"pt") mod r} from the public point set; Mask = H_1(D_{trans} || ρ || "mask"). The challenge has size poly(λ) bits.

11.2 ResolveChallenge

> **Definition 20** (*ResolveChallenge*). ResolveChallenge(comp, C_{pub}, profile) is a deterministic algorithm taking private completion state comp = (g_{s,θ}, CompGrid, ...) and challenge C_{pub} = (D_{trans}, ρ, Ops, Evals, Mask), returning a private response tuple R_{priv} = (R_a, R_b, R_c):

R_a = (shadow-projections): For each index i ∈ Ops, compute π_i = ⟨S(H_{s,θ}), S(ψ_i)⟩ / ‖S(ψ_i)‖^2 (projection of shadow-sample onto basis direction i). Output (π_i)_{i∈Ops}.

R_b = (consistency residuals): For each ζ ∈ Evals, compute the Laplacian residual res_ζ = |Δ_{1/2}^{disc}(H_{s,θ})(ζ)| using the discretized Laplacian. Output (res_ζ)_{ζ∈Evals}.

R_c = (ephemeral commitments): Compute c = H_0(D_{trans} || ρ || g_{s,θ}^{sample} || "commit") where g^{sample} is the shadow evaluated at the first t evaluation points. Output c.

The response R_{priv} is computationally accessible only to the holder of comp (since it requires evaluating g_{s,θ} at specific points, which is not possible from T_s alone under the ROM).

11.3 CheckFamilyProof

> **Definition 21** (*CheckFamilyProof*). CheckFamilyProof(T_m, π_{fam}, profile) is a deterministic algorithm taking message-bound transcript T_m and family proof π_{fam}, returning 1 (accept) or 0 (reject):

Parse π_{fam} = (dim_vector, residual_vector, commitment, mask_key, structure_tag). Check: (1) dim_vector lies within η_{proj} of the declared subspace (verifiable from T_m's shadow-admissibility tag); (2) residual_vector entries are ≤ η_{lap} (Laplacian consistency); (3) commitment = H_2(π_{fam}[dim,res] || T_m || "fam-commit") matches the declared value; (4) structure_tag is a valid family identifier for the current protocol version. Return 1 iff all four checks pass, 0 otherwise.

Design requirement: CheckFamilyProof must be family-validating, not identity-revealing. It certifies that a prover knows a completion in the admissible family, without determining which specific completion was used. In particular, it makes no query to O_f or O_V and cannot distinguish elements of C(T_s).

12. Formal Security Definitions

12.1 Adversary Model

Definition 16 (*Adversary Capabilities in the RO-Model*). When cryptographic constructions are analyzed (Sections 13–14), the adversary A is a PPT algorithm with: (i) random oracle access to H_0, H_1, H_2, KDF (modeled as independent random oracles); (ii) the specified game oracles (KeyGen, Encaps, Sign, etc.); (iii) standard ROM access (Definition 6) to the MMIP layer. A is adaptive: it may choose oracle queries based on previous responses. The security parameter λ is given in unary.

Definition 17 (*Key Encapsulation Mechanism*). A KEM Π = (KeyGen, Encaps, Decaps) satisfies correctness if Pr[Decaps(sk,ct) = K | (ct,K) ← Encaps(pk), (pk,sk) ← KeyGen(1^λ)] = 1 - negl(λ).

Game: KEM-IND-CPA_A(λ)

1.	`(pk, sk) ← KeyGen(1^λ)`
2.	`(ct*, K_0) ← Encaps(pk)`
3.	`K_1 ← {0,1}^κ  // uniformly random key`
4.	`b ← {0,1}  // challenger's bit`
5.	`b' ← A^{H_0,H_1,H_2,KDF}(pk, ct*, K_b)`
6.	`return 1 iff b' = b`

Definition 18 (*IND-CPA Security*). Adv^{KEM-IND-CPA}(A,λ) = |Pr[KEM-IND-CPA_A(λ)=1] - 1/2| = negl(λ) for all PPT A.

The digital signature security game SIG-EUF-CMA_A(λ) is standard [KL21, §13]: A gets pk, has adaptive access to a signing oracle Sign(sk,·), and wins by producing (m*,σ*) with Verify(pk,m*,σ*) = 1 and m* not previously queried. Adv^{SIG-EUF-CMA}(A,λ) = Pr[wins].

13. Key Encapsulation Mechanism: Construction and Security Proof

13.1 Construction

13.1.1 Key Generation

KeyGen(1^λ):

```
s ← RNG_{pub}(λ);    f_s ← BuildHolomorphicSeries(s, λ, params)
T_s ← BuildTranscript(f_s, profile_pub)        // frozen before θ is generated
θ ← RNG_{priv}(m(λ));   g_{s,θ} ← BuildShadow(s,θ,Ψ,ε)
comp ← BuildCompletionState(f_s, g_{s,θ}, profile_priv)
(k_{prf}, k_{mask}, inst_id) ← RNG_{priv}(3λ)
pk ← T_s;   sk ← (s, θ, g_{s,θ}, comp, k_{prf}, k_{mask}, inst_id)
```

13.1.2 Encapsulation

Encaps(pk):

```
ρ ← {0,1}^λ
C_{pub} ← BuildChallenge(T_s, ρ, kem_profile)    // Def. 19
W ← H_1(T_s || C_{pub} || "witness")
ct_1 ← Encode(version || C_{pub} || W || domain_sep_kem)
K_{enc} ← KDF(T_s || ct_1 || ρ || "encap")
return (ct_1, K_{enc})
```

13.1.3 Decapsulation

Decaps(sk, ct_1):

```
Parse ct_1 → (version, C_{pub}, W, domain_sep);   validate all fields
Verify CTEQ(H_0(T_s), C_{pub}.D_{trans})    // constant-time check
R_{priv} ← ResolveChallenge(comp, C_{pub}, kem_profile)  // Def. 20
R_{pub} ← ProjectResponse(R_{priv}, H_1(T_s || C_{pub}.ρ || "mask"))
W* ← H_1(T_s || C_{pub} || H_0(R_{pub} || "wc") || "witness")
```

```
if ¬CTEQ(W*, W): return ⊥    // constant-time failure
conf ← H_2(T_s || C_{pub} || R_{pub} || "confirm")
K_{dec} ← KDF(T_s || ct_1 || R_{pub} || "shared")
return (conf, K_{dec})
```

13.2 Security Theorem with Complete Reduction

> **Theorem 12** (*IND-CPA Security of MMIP-KEM*). Under H1–H4 and the RO-Model for H_0,H_1,H_2,KDF, with adversary A making at most q_H total random oracle queries:
>
> $$Adv^{KEM-IND-CPA}(A,\lambda) \leq (q_H + 1) \cdot 2^{-\alpha\lambda}.$$

The bound is tight up to the q_H + 1 factor, which is unavoidable in the RO-Model.

Proof. **Proof structure.** We construct an explicit PPT reduction B such that any IND-CPA adversary A with advantage ε implies a MMIP solver B with advantage ε/(q_H+1) - negl(λ). The reduction runs in time poly(λ) plus A's running time.

Hybrid argument. We use three hybrid games. Game G_0 is the real IND-CPA game. Game G_1 replaces the real key K_0 with a uniformly random key K_r, independent of ct_1. Game G_2 is identical to G_1. We show $|Pr[G_0=1] - Pr[G_1=1]| \leq (q_H+1) \cdot 2^{-\alpha\lambda}$ and $|Pr[G_1=1] - 1/2| = 0$.

G_0 to G_1: Real key to random key. The real key is K_0 = KDF(T_s || ct_1 || R_{pub} || "shared"). In G_1, we replace K_0 with $K_r \leftarrow \{0,1\}^\kappa$. An adversary distinguishing G_0 from G_1 must query KDF on input (T_s || ct_1 || R_{pub} || "shared"). Since KDF is a random oracle, K_0 is uniformly distributed unless this specific input is queried.

Claim: computing R_pub requires comp. We show: any PPT algorithm computing R_pub from T_s and C_pub (without comp) can be turned into a MMIP solver. R_pub = ProjectResponse(ResolveChallenge(comp, C_pub)). ResolveChallenge computes R_a = shadow projections $\langle S(H_{s,\theta}), S(\psi_i)\rangle$ for $i \in$ Ops. Computing $S(H_{s,\theta}) = \xi_{1/2}(H_{s,\theta})$ at the challenge points requires $H_{s,\theta}^{-} = R_{s,\theta}$, which depends on $g_{s,\theta}$ = θ-specific shadow (denied by ROM). More formally: $R_a^{(j)} = \Sigma_i b_{ij}\ \theta_i + c_j(s)$ where b_{ij}, $c_j(s)$ are known constants; computing $(R_a^{(j)})$ from T_s alone requires recovering the linear combination $\Sigma_i b_{ij}\ \theta_i$ for each j, i.e., recovering θ. Hence any A computing R_pub is a MMIP solver.

Formal extraction. Reduction B receives MMIP instance T_s and simulates G_0 for A. B sets pk = T_s. B generates ρ, computes C_pub, W honestly. B gives A the challenge (pk, ct_1, K_b) where $b \leftarrow \{0,1\}$ and K_0 = KDF(T_s || ct_1 || X_guess || "shared") for a *guessed* R_pub value X_guess = 0^λ initially. B maintains a table of A's KDF queries. When A makes a KDF query on (T_s || ct_1 || X || "shared") for some X, B checks if X is consistent with a valid R_pub (using the structure of ProjectResponse). If so, B records X as a candidate for R_pub and extracts θ from X by inverting ProjectResponse (which B can do since it knows the projection basis from the public profile).

Probability analysis. If A distinguishes with advantage ε, A must query KDF on the correct input (T_s || ct_1 || R_pub || "shared") with probability ≥ ε (otherwise, K_0 is uniformly random from A's view, giving no advantage). Over q_H queries, A queries the correct input with

probability ≥ ε. B guesses which of A's q_H queries contains the correct R_pub: correct guess probability 1/q_H. Hence Adv_{MMIP}(B) ≥ ε/q_H. By Theorem 5: ε/q_H ≤ Adv_{MMIP}(B) ≤ 2^{-αλ}. Therefore ε ≤ q_H · 2^{-αλ}. More carefully: B makes one additional guess for the initial commitment, giving (q_H+1) · 2^{-αλ}.

G_1 is a fair coin: In G_1, K_b is either K_0 or K_r; both are uniform from A's view since K_0 = KDF(input) and KDF is a fresh random oracle not queried on the specific input (by the argument above). Hence Pr[G_1=1] = 1/2 exactly.

Combining: Adv^{KEM-IND-CPA}(A,λ) = |Pr[G_0=1] - 1/2| ≤ |Pr[G_0=1] - Pr[G_1=1]| + |Pr[G_1=1] - 1/2| ≤ (q_H+1)·2^{-αλ} + 0.

■

Remark. **Tightness.** The loss factor (q_H+1) is inherent in the random oracle model: any construction where the shared key is derived via a random oracle admits at most this loss in the IND-CPA reduction [HHK17]. Our reduction is therefore tight up to this factor, matching the best-known bounds for RO-model KEMs.

14. Digital Signature Scheme: Construction and Security Proof

14.1 Construction

14.1.1 Signing

Sign(sk, m):

```
η ← PRF(k_{sign}, H_0(m) || H_0(T_s) || "nonce") ⊕ RNG_{priv}(λ)
η_{pub} ← H_0(η || "pub_proj")
T_m ← MessageBind(T_s, m, η_{pub}, sig_profile)
π_{bind} ← H_1(T_s || m || T_m || η_{pub} || "bind")
F_{wit} ← BuildFamilyWitness(comp, T_m, η, sig_profile)
π_{fam} ← MaskWitness(F_{wit}, H_1(η || T_m || "fam-proof"), sig_profile)
σ ← Encode(T_m || π_{bind} || π_{fam} || η_{pub} || sig_version || domain_sep_sig)
return σ
```

14.1.2 Verification

Verify(pk, m, σ):

```
Parse σ → (T_m, π_{bind}, π_{fam}, η_{pub}, sig_version, domain_sep)
Validate format, version, domain_sep
```

```
T_{m,exp} ← MessageBind(pk, m, η_{pub}, sig_profile)
if ¬CTEQ(T_m, T_{m,exp}): reject
π_{bind,exp} ← H_1(pk || m || T_m || η_{pub} || "bind")
if ¬CTEQ(π_{bind}, π_{bind,exp}): reject
if ¬CheckFamilyProof(T_m, π_{fam}, sig_profile): reject
// Def. 21
accept
```

14.2 Security Theorem with Complete Reduction

Theorem 13 (*EUF-CMA Security of MMIP-Signature*). Under H1–H4 and the RO-Model, with A making q_s signing queries and q_H random oracle queries:

$$Adv^{SIG\text{-}EUF\text{-}CMA}(A,\lambda) \leq q_s \cdot (q_H + 1) \cdot 2^{-\alpha\lambda}.$$

Proof. **Reduction construction.** We build reduction B from EUF-CMA to MMIP. B receives MMIP instance T_s and simulates the EUF-CMA game for A.

Simulation of KeyGen. B sets pk = T_s. B does not know sk. B must simulate signing oracle Sign(sk,·) without comp.

Simulation of signing oracle. For each signing query m_i: (1) B generates η_i ← {0,1}^λ. (2) B computes T_{m_i} = MessageBind(T_s, m_i, η_{i,pub}). (3) B computes π_{bind,i} honestly from H_1. (4) To simulate π_{fam,i}: B programs H_1 at the input (η_i || T_{m_i} || "fam-proof") to return a uniformly random string r_i, and sets π_{fam,i} = MaskWitness(F_{wit,sim}, r_i) where F_{wit,sim} is a simulated family witness. The simulation is indistinguishable from the real π_{fam,i} because: (a) H_1 is a random oracle, so r_i is uniform; (b) MaskWitness is a bijection from witness space to proof space; (c) CheckFamilyProof checks only structural properties of π_{fam} (Definition 21), not the specific witness value. Hence the simulated π_{fam,i} passes CheckFamilyProof. Simulation fails only if A queries H_1 at the same point (η_i || T_{m_i} || "fam-proof") before B programs it; this happens with probability at most q_H / 2^λ per signing query (birthday bound), total probability ≤ q_s · q_H / 2^λ = negl(λ).

Forgery analysis. A produces (m*, σ*) with Verify(pk, m*, σ*) = 1 and m* ∉ {m_i}. Parse σ* = (T_{m*}, π_{bind*}, π_{fam*}, η*). Valid σ* requires: (1) T_{m*} = MessageBind(T_s, m*, η*_{pub}); (2) π_{bind*} correct; (3) CheckFamilyProof(T_{m*}, π_{fam*}) = 1.

Extracting MMIP solution from forgery. π_{fam*} passes CheckFamilyProof. By Definition 21, π_{fam*} must contain a valid commitment c* = H_2(π_{fam*}[dim,res] || T_{m*} || "fam-commit"). This commitment was either: (a) obtained from a signing oracle query on m* (impossible since m* ∉ {m_i}); or (b) computed by A by querying H_2 on the input (π_{fam*}[dim,res] || T_{m*} || "fam-commit"). Case (b): A's H_2 query reveals the dim-vector (R_a component of a valid family witness), from which B can extract θ by inverting the projection formula R_a^{(j)} = Σ_i b_{ij}θ_i + c_j(s). B identifies the correct query by checking consistency with T_s.

Probability calculation. A successfully forges with probability ε (total advantage). In case (b), A made a H_2 query revealing θ. B guesses which of A's q_H queries is the relevant one: probability 1/q_H correct guess per signing context. Over q_s potential signing contexts and q_H queries: Adv_{MMIP}(B) ≥ ε/(q_s · q_H) - negl(λ). By Theorem 5: ε/(q_s(q_H+1)) ≤ 2^{-αλ}, giving ε ≤ q_s(q_H+1) · 2^{-αλ}.

Tightness remark. The factor q_s is standard in EUF-CMA signature reductions [BR96,GPV08] and cannot be avoided without restructuring the scheme. For q_s = 2^{30}: at λ = 192, effective EUF-CMA security is 192 - 30 - log(q_H+1) ≥ 162 - log(q_H) bits, exceeding standard thresholds.

15. Concrete Security Analysis and Comparison

15.1 Concrete Security Bounds with Explicit Loss

Table 2. Concrete security in bits (q_H = 2^{64}, q_s = 2^{30}).

Lambda	MMIP class.	MMIP quant.	KEM eff. classical	KEM eff. quantum	Sig eff. classical
128	128	64	64	32	48
192	192	96	128	64	112
256	256	128	192	96	176
512	512	256	448	224	432

Effective KEM security = αλ - log(q_H+1). For q_H = 2^{64} and λ = 192: 192-64 = 128 classical bits, 96-32 = 64 quantum bits. Setting λ = 256 achieves 192 classical bits matching standard targets.

15.2 Comparison with NIST PQC Candidates

Table 3. Qualitative comparison.

Scheme	Security basis	Quantum model	Unique sk in pk	Info-theoretic	Ontol. hard	Reduction loss
Kyber	Module-LWE	Comp.	Yes	No	No	negl(λ)
Dilithium	Mod-SIS	Comp.	Yes	No	No	negl(λ)
FALCON	NTRU	Comp.	Yes	No	No	negl(λ)
McEliece	Goppa	Comp.	Yes	No	No	negl(λ)
MMIP-KEM	Info.-theoretic	Info.-th.	No (fiber)	Yes	Yes	(q_H+1)·2^{-αλ}

15.3 Hybrid Construction for Defense in Depth

A KEM combiner [GHP18] yields combined advantage min(Adv_MMIP, Adv_Kyber) + negl(lambda), providing defense in depth if either component fails.

L2. No worst-case/average-case reduction. Unlike LWE [Reg05,Pei09], no reduction from a worst-case NP problem. MMIP hardness is information-theoretic (Theorem 18), which is stronger for unbounded adversaries but lacks the worst-case connection.

L3. Two-pass KEM. Requires two communication rounds. Suitable for TLS-like exchange; unsuitable for strict one-pass settings.

L4. Reduction loss factor. The KEM reduction has loss (q_H+1) and signature has q_s(q_H+1). For large q_H, this requires larger λ. This is standard in the RO-Model and unavoidable without restructuring.

L5. Efficiency. Holomorphic series computation, Eichler integral approximation, and shadow sampling have not been benchmarked at scale.

16.2 Open Problems

O1. UC security. Extend to Universal Composability [Can01].

O2. Non-abelian HSP. Theorem 3 rules out abelian HSP reductions; analysis of non-abelian HSP [HRT03] would complete the quantum picture.

O3. One-pass KEM. Requires a self-contained encapsulation without a confirmation round.

O4. Tight reduction. Reduce the (q_H+1) loss factor in the KEM reduction, ideally to 1 (perfectly tight).

O5. Worst-case hardness. Reduce from a worst-case problem in automorphic form theory.

O6. Zero-knowledge family membership. ZK proofs of C(T) membership without revealing θ.

O7. Property P1 classification. Characterize all families satisfying Property P1.

O8. Shadow space dimension theory. Exact dim $S_{3/2}(\Gamma_0(4M))$ for small M; sharp bounds.

O9. Formal verification. Mechanically verify security proofs in Coq or Lean.

17. Conclusion

We introduced a cryptographic paradigm grounded in **structural non-identifiability under restricted observability**, backed by a rigorous mathematical foundation (harmonic Maass forms, shadow operator, Eichler integrals) and a complete formal security treatment.

The paper establishes nine distinct contributions. Theoretically: information-theoretic indistinguishability for all adversaries including unbounded ones; formal undecidability of universal identification via explicit construction; quantum-attack orthogonality (no abelian HSP, no unique Grover target); completeness-soundness duality; ontological hardness; and a three-way hardness taxonomy. Constructively: an explicit ε-separation lemma (Lemma 4) and dimension witness (Lemma 5) verify all hypotheses concretely; formal specifications for BuildChallenge, ResolveChallenge, and CheckFamilyProof close the proof gaps; and tight reductions for IND-CPA KEM and EUF-CMA signatures with explicit loss factors are provided.

The result is a paradigm in which security derives not from "how much does recovery cost?" but from "what, in fact, is determined by what is observed?" This is

ontological hardness: information-theoretically impossible identification, holding for all algorithms, bounded or not, under the Restricted Observable Model.

Appendix A — Notation Table

Symbol	Description
λ	Security parameter
$\mathbb{H}$	Upper half-plane $\{\tau \in \mathbb{C} : \mathrm{Im}(\tau) > 0\}$
$q = e^{2\pi i\tau}$	Nome
Γ, $\Gamma_0(4M)$	Congruence subgroup; with $M = \prod_{p\leq 2\lambda} p$
$k = 1/2$	Modular weight (primary)
$H_k(\Gamma)$, $S_k(\Gamma)$	Harmonic Maass forms; cusp forms
$H = H^+ + H^-$	Canonical decomposition
ξ_k	Shadow operator: $H_k \to S_{2-k}$; ontological boundary
$g = \xi_k(H)$	Shadow; private identity datum
$R_g(\tau)$	Eichler-type completion integral
$f_s = H^+$	Holomorphic anchor (public, θ-independent)
$H_{s,\theta} = f_s + R_{g_{s,\theta}}$	Completed form; θ-indexed completion
$s \in \{0,1\}^\lambda$	Public seed

Symbol	Description
θ ∈ {0,1}^m, m = λ	Private vector
ε = 2^{-λ}	Shadow separation parameter (Lemma 4)
g_{s,θ}	Private shadow; θ-dependent cusp form
T_s, T_s^disc	Transcript (continuous/discrete); both θ-independent
C_λ(T)	Consistency class; |C_λ(T)| = 2^λ
α = 1	Ambiguity exponent
Property P1	Core non-identifiability condition (Def. 5)
ROM	Restricted Observable Model (Def. 6)
MMIP	Mock Modular Identification Problem (Def. 7)
Adv_{MMIP}(A,λ)	MMIP advantage
Ontological hardness	Def. 11: hardness from informational absence
comp	Completion state (private)
V, V^disc	Verification operator
q_H, q_s	RO query count; signing query count
negl(λ)	Negligible function

Appendix B — Minimal Security Conditions

The following conditions are jointly necessary and sufficient for the security guarantees of Theorem 17:

B1. $|C_\lambda(T)| \geq 2^\lambda$ (H1, achieved with $\alpha = 1$ by Lemma 5).

B2. $F \approx_{obs} F'$ for all $F,F' \in C_\lambda(T)$ (H2, follows from Property P1).

B3. No ROM oracle computes any functional of R_F, g, or H^- (Def. 6).

B4. $\theta \mapsto T_s$ is constant (Property P1, verified by Proposition 1 and Lemma 4).

B5. V certifies $C(T) \neq \emptyset$ without determining which $F \in C(T)$ is true (Cor. 1).

B6. Stable discretization (H4, verified by Theorem 11 for explicit parameters).

B7. Perfect completeness and negl(λ)-soundness simultaneously (Theorem 16).

B8. ε satisfies Lemma 4 ($\varepsilon \neq 0$; recommended $\varepsilon = 2^{-\lambda}$).

B9. $\dim S_{3/2}(\Gamma) \geq \lambda$ (Lemma 5, with $\Gamma = \Gamma_0(4M(\lambda))$).

Appendix C — Parameter Recommendations

λ = 128: m=128, N=256, r=64, u=32, t=64, p=192, M=∏_{p≤257}p, ε=2^{-128}

λ = 192: m=192, N=384, r=96, u=48, t=96, p=256, M=∏_{p≤383}p, ε=2^{-192}

λ = 256: m=256, N=512, r=128, u=64, t=128, p=384, M=∏_{p≤512}p, ε=2^{-256}

Tolerances: η_{coeff} = η_{eval} = 2^{-λ}, η_{proj} = 2^{-λ/2}, η_{lap} = 2^{-λ}. All are protocol constants, not implementation choices. Note: the M values given are sufficient but not tight; significantly smaller M values may achieve dim $S_{3/2}(\Gamma_0(4M)) \geq \lambda$ in practice, to be determined by direct computation using dimension formulas [DS05].

References

[ABC+20] D.J. Bernstein et al. "Classic McEliece." NIST PQC Round 3, 2020.

[BBBV97] C.H. Bennett, E. Bernstein, G. Brassard, U. Vazirani. "Strengths and weaknesses of quantum computing." SIAM J. Comput. 26(5):1510–1523, 1997.

[BDK+18] J. Bos, L. Ducas, E. Kiltz et al. "CRYSTALS-Kyber." EuroS&P 2018.

[BF04] J.H. Bruinier, J. Funke. "On two geometric theta lifts." Duke Math. J. 125(1):45–90, 2004.

[BGI+01] B. Barak et al. "On the (im)possibility of obfuscating programs." CRYPTO 2001.

[BGW88] M. Ben-Or, S. Goldwasser, A. Wigderson. "Completeness theorems for non-cryptographic fault-tolerant distributed computation." STOC 1988.

[BHK+19] D.J. Bernstein et al. "SPHINCS+." NIST PQC Round 3, 2019.

[BHMT02] G. Brassard, P. Høyer, M. Mosca, A. Tapp. "Quantum amplitude amplification and estimation." AMS Contemp. Math. 305:53–74, 2002.

[Bon90] A. Boneh. Notes on quantum HSP, 1990.

[BR93] M. Bellare, P. Rogaway. "Random oracles are practical." CCS 1993.

[BR94] M. Bellare, P. Rogaway. "Entity authentication and key distribution." CRYPTO 1993.

[BR96] M. Bellare, P. Rogaway. "The exact security of digital signatures." EUROCRYPT 1996.

[Can01] R. Canetti. "Universally composable security." FOCS 2001.

[CCD88] D. Chaum, C. Crépeau, I. Damgård. "Multiparty unconditionally secure protocols." STOC 1988.

[CD22] W. Castryck, T. Decru. "An efficient key recovery attack on SIDH." EUROCRYPT 2023.

[DH76] W. Diffie, M. Hellman. "New directions in cryptography." IEEE Trans. Inf. Theory 22(6):644–654, 1976.

[DKL+18] L. Ducas et al. "CRYSTALS-Dilithium." TCHES 2018.

[DLRW23] L. De Feo et al. "SQIsignHD." EUROCRYPT 2024.

[DS05] F. Diamond, J. Shurman. "A First Course in Modular Forms." Springer, 2005.

[EHKS04] M. Ettinger et al. "The quantum query complexity of the hidden subgroup problem." IPL 91(1):43–48, 2004.

[GGM86] O. Goldreich, S. Goldwasser, S. Micali. "How to construct random functions." J. ACM 33(4):792–807, 1986.

[GHP18] F. Giacon, F. Heuer, B. Poettering. "KEM Combiners." PKC 2018.

[GM84] S. Goldwasser, S. Micali. "Probabilistic encryption." JCSS 28(2):270–299, 1984.

[Göd31] K. Gödel. "Über formal unentscheidbare Sätze." Monatsh. Math. Phys. 38:173–198, 1931.

[GPV08] C. Gentry, C. Peikert, V. Vaikuntanathan. "Trapdoors for hard lattices." STOC 2008.

[Gro96] L.K. Grover. "A fast quantum mechanical algorithm for database search." STOC 1996, 212–219.

[HHK17] D. Hofheinz, K. Hövelmanns, E. Kiltz. "A modular analysis of the Fujisaki-Okamoto transformation." TCC 2017.

[HRT03] S. Hallgren, A. Russell, A. Ta-Shma. "The hidden subgroup problem and quantum computation." SIAM J. Comput. 32(4):916–934, 2003.

[Iwa97] H. Iwaniec. "Topics in Classical Automorphic Forms." AMS, 1997.

[JF11] D. Jao, L. De Feo. "Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies." PQCrypto 2011.

[Kit95] A.Yu. Kitaev. "Quantum measurements and the Abelian stabilizer problem." arXiv:quant-ph/9511026, 1995.

[KL21] J. Katz, Y. Lindell. "Introduction to Modern Cryptography." 3rd ed. CRC Press, 2021.

[Kob87] N. Koblitz. "Elliptic curve cryptosystems." Math. Comp. 48:203–209, 1987.

[LPR13] V. Lyubashevsky, C. Peikert, O. Regev. "A toolkit for ring-LWE cryptography." EUROCRYPT 2013.

[Mil86] V.S. Miller. "Use of elliptic curves in cryptography." CRYPTO 1985.

[MMPPW22] L. Maino et al. "A direct key recovery attack on SIDH." EUROCRYPT 2023.

[NIST24] NIST. "Post-Quantum Cryptography Standardization: FIPS 203, 204, 205." 2024.

[Ono09] K. Ono. "Unearthing the visions of a master: Harmonic Maass forms and number theory." Harvard-MIT Curr. Dev. Math., 2008.

[Pei09] C. Peikert. "Public-key cryptosystems from the worst-case shortest vector problem." STOC 2009.

[PFH+20] T. Prest et al. "FALCON." NIST PQC Round 3, 2020.

[Ram20] S. Ramanujan. "Collected Papers." Cambridge, 1927 (repr. 1962).

[Reg05] O. Regev. "On lattices, learning with errors, random linear codes, and cryptography." STOC 2005.

[Rob23] D. Robert. "Breaking SIDH in polynomial time." EUROCRYPT 2023.

[RSA78] R.L. Rivest, A. Shamir, L. Adleman. "A method for obtaining digital signatures." Commun. ACM 21(2):120–126, 1978.

[Scu15] D. Galindo, J. Herranz, J. Villar. "Tight security bounds for key encapsulation mechanisms." ESORICS 2010.

[Sha49] C.E. Shannon. "Communication theory of secrecy systems." Bell Syst. Tech. J. 28(4):656–715, 1949.

[Shi73] G. Shimura. "Introduction to the Arithmetic Theory of Automorphic Functions." Princeton, 1973.

[Sho94] P.W. Shor. "Algorithms for quantum computation." FOCS 1994, 124–134.

[Tur36] A.M. Turing. "On computable numbers." Proc. London Math. Soc. 42:230–265, 1936.

[Zag09] D. Zagier. "Ramanujan’s mock theta functions." Sém. Bourbaki 986, 2007–2008.

[Zwe02] S. Zwegers. "Mock Theta Functions." Ph.D. thesis, Utrecht University, 2002.