Information-Theoretic Boundaries of the Rosario Primitive

Information-Theoretic Boundaries of the Rosario Primitive: Observable-Channel Isolation, Inference Limits, and Session-Bound Residues

Author(s) *(to be completed)*
Affiliation(s) *(to be completed)*
Corresponding author: *(email to be completed)*

Received: *(date)* · **Accepted:** *(date)* · **Published:** *(date)*

Abstract

We formalize the information-theoretic boundaries of the Rosario Primitive, with emphasis on the observable channel and the limits of adversarial inference. The central claim is not merely that the system is hard to break computationally, but that, in the ideal regime, the observable channel is constructed so that it carries zero mutual information about the hidden state beyond a nonce-scoped event-validity residue. In the non-ideal regime, the architecture admits explicit leakage bounds that remain small, measurable, and non-accumulative under proper implementation constraints.

The analysis focuses on five linked ideas: (1) the distinction between hidden state and observable residue, (2) the role of microsecond-bound geometry in destroying replay value, (3) the Mechanical Displacement Axiom and its entropy consequences for interface observation, (4) the factorial uncertainty induced by a private bijective morphism, and (5) a formal real-vs-simulated transcript security game that captures distinguishability and inference advantage in both classical and quantum settings [16,17]. The result is a framework for stating and auditing the system's information-theoretic guarantees in language suitable for technical review.

The core conclusion is precise: the public channel may emit a binary event outcome, but that output is a historical residue tied to a unique hidden geometric alignment. It is not a reusable credential [6,7], not a coordinate of the secret [11], and not a predictive bit for future sessions [5].

This distinction is the key to understanding the Rosario Primitive as an information-theoretic authentication and attestation architecture rather than a conventional hardness-only cryptosystem.

Keywords: information-theoretic security; authentication; observable channel; mutual information; perfect secrecy; Rosario Primitive; session-bound residue; Mechanical Displacement Axiom; private bijective morphism; replay resistance; credential stuffing; passkey; WebAuthn.

---

1. Introduction

Many security systems are evaluated by a single question: how difficult is it for an attacker to compute the secret? That framing is incomplete. A stronger and more useful question is: what information does the attacker actually observe, and does the observable channel reduce uncertainty about the hidden state?

This paper adopts the second framing.

The Rosario Primitive is designed around an information-theoretic objective. The observable channel should not reveal the hidden state coordinates required to reconstruct a valid proof event. Even if an attacker records every visible artifact, the recorded data should not collapse the hypothesis space of possible secrets, private mappings, or future-valid witness alignments. The attacker may learn that an event succeeded at a specific moment. But that success must be non-predictive and non-transferable.

That is why the architecture is best understood through Shannon theory [1,2].

The system does not rely solely on computational hardness in the familiar sense of factoring [37], discrete logarithms [38], or algebraic inversion. Instead, it is structured so that the transcript channel is conditionally independent of the hidden state under the public challenge context. In the ideal regime, this is expressed as zero mutual information [2,3]. In real deployments, the target is a bounded, auditable deviation from that ideal, with no useful accumulation across sessions because the hidden geometry is re-randomized at microsecond granularity [12].

The paper proceeds as follows. Section 2 defines the hidden state, the observable channel, and the event-validity residue. Section 3 introduces the ideal and non-ideal inference-advantage framework. Section 4 analyzes the Mechanical Displacement Axiom and the entropy of biometric observation. Section 5 formalizes hypothesis space expansion under a private bijective morphism. Section 6 treats the event-validity residue and temporal fossilization. Section 7 gives a formal security game (real vs. simulated transcripts). Section 8 addresses quantum attacks and why the main defense is information absence, not trapdoor hardness. Section 9 summarizes implementation conditions for maintaining the information-theoretic boundary. Section 10 provides a compact summary of the three defensive layers. Section 11 concludes with the practical meaning of the residue: public auditability without secret leakage.

---

2. System Model and Information Variables

2.1 Hidden State

Let the hidden state be denoted by:

$$X = (K, \Phi, E, \tau, \Theta)$$

where:

- K is the active secret, or active selector-path witness chosen from a family of committed secrets;
- Φ is the private bijective morphism (the user's private synonym map);
- E is the hidden entropy field embedded in the verifier/prover environment;
- τ is the microsecond-bound time coefficient used in session geometry derivation;
- Θ denotes additional hidden geometric parameters (ring offsets, leaf assignments, internal derivation state, and related manifold parameters).

The adversary does not observe X directly.

2.2 Observable Channel

Let the observable channel be denoted by:

$$O = (C, W, Y)$$

where:

- C is the public session context (challenge metadata, protocol framing, public labels, timing envelope, and any publicly visible interface state);
- W is the observable witness representation, if any (for some implementations, this may be visible labels, interface actions, or synonym declarations);
- Y is the public validation residue, typically a binary success/failure output.

In some deployments, the public witness component W is reduced or hidden, in which case the channel is:

$$O = (C, Y)$$

This paper covers both forms. The stronger target is the residue-centric channel.

2.3 Session Geometry

The hidden session geometry for round r is derived as:

$$G_r = f(E, \tau, r, \Theta)$$

The exact implementation of f may vary, but the critical property is invariance of derivation between prover and verifier and fresh re-randomization of geometry with each microsecond-scoped τ [12].

This implies a core operational fact:

- A witness aligned to $G_r(\tau)$ is generally misaligned to $G_r(\tau')$ for $\tau' \neq \tau$.

This is the root of non-replayability.

2.4 Event-Validity Residue

The verifier evaluates internal membership or consistency predicates across rounds and accumulates them into one internal aggregate predicate. Let:

$$\Lambda = \bigwedge_{r=1}^{R} M_r$$

where M_r is the hidden round predicate for round r.

The public output is then:

$$Y = \quad \Lambda = 1$$

This output is deliberately minimal. The public channel does not expose the vector $(M_1, \ldots, M_R)$, the failure index, or partial correctness. The output Y is a historical statement only: the submitted witness array either validated or failed for the unique hidden state and session geometry associated with that event.

This distinction is central. The output is often called a "bit," but in this architecture it is better described as an **event-validity residue**. It is a record of a completed event, not a coordinate of the hidden state.

---

3. Quantitative Inference-Advantage Framework

3.1 Adversarial Objective

Consider a passive adversary attempting to infer a predicate $\psi(X)$ of the hidden state from the observable channel. Examples include:

- "Which secret K_j from a secret family was active?"
- "What is the private morphism Φ?"
- "Which target zone was intended for a given interface movement?"

- “Can a future-valid witness be predicted from prior observations?”

The proper metric is not only computational cost but also **inference advantage**: how much does observing O improve the adversary’s posterior over $\psi(X)$? [3,4]

3.2 Ideal Regime

In the ideal regime, the architecture targets:

$$I(X; O \mid C) = 0$$

Equivalently, for any predicate ψ over X, the posterior remains equal to the prior:

$$\Pr[\psi(X) = 1 \mid O, C] = \Pr[\psi(X) = 1 \mid C]$$

This is the strongest form of the claim: conditioned on the public challenge context, the observable channel does not reduce uncertainty about the hidden state [2,3,4].

In this regime:

- transcripts are indistinguishable from simulated transcripts generated without access to X,
- interface observations remain uniformly ambiguous across hyperplane leaves,
- synonym observations do not reveal the private morphism,
- the event-validity residue Y does not identify the active secret in a family.

3.3 Non-Ideal Regime

Real implementations are never perfectly ideal. Timing jitter, rendering asymmetries, human behavior variance, and hardware effects may introduce bounded leakage. The correct formulation is then:

$$I(X; O \mid C) \leq \varepsilon$$

for a small implementation-dependent ε.

Likewise, distinguishability and inference advantages are bounded by a small function of implementation balance parameters. Let δ denote a balance deviation term (capturing non-uniformity in zone transitions, interface timing bias, or imperfect symmetry). Then one may express:

$$\mathrm{Adv}^{\mathrm{dist}} \leq g(\delta)$$

$$\mathrm{Adv}^{\mathrm{inf}}(\psi) \leq h(\delta)$$

with $g(\delta), h(\delta) \to 0$ as $\delta \to 0$.

The architecture is therefore not “all or nothing.” It supports a measurable path from theoretical target to implementation bound.

3.4 Inference-Advantage Interpretation

The practical interpretation of the table is:

- In the ideal regime, the adversary's inference advantage is zero.
- In the non-ideal regime, any advantage is bounded and linked to observable imbalance, not to the cryptographic core.
- Because the session geometry re-randomizes with τ, leakage does not accumulate cleanly across sessions unless the implementation repeatedly violates the same balance assumptions.

This is important for peer review. It states a testable boundary. The system's security can be audited in terms of measurable deviations from symmetry and uniformity [9,20].

---

4. Mechanical Displacement Axiom and Maximum-Entropy Interface Observation

4.1 The Problem of Biometric Side Observation

In many interactive systems, an observer can learn by watching:

- gaze direction,
- dwell time,
- finger movement,
- cursor motion,
- keystroke cadence,
- body orientation.

Even if the cryptography is strong, the interface can leak the target coordinate. This is a common weakness in practical authentication systems [33,34].

The Rosario architecture addresses this with a structural rule: the **Mechanical Displacement Axiom**.

4.2 Mechanical Displacement Axiom

The axiom states that any legal interface movement causes a simultaneous, balanced transition across all hyperplane leaves:

- exactly one value enters each leaf,
- exactly one value leaves each leaf,
- all leaves transition in lockstep.

This is the "simultaneous 1-in / 1-out" property.

The consequence is immediate: observing that “something changed” in the interface does not identify which leaf is the target, because every leaf changed.

4.3 Entropy of the Visible Leaf Target

Let Z be the random variable representing the target hyperplane leaf among H leaves. Under balanced displacement and absent knowledge of the private morphism, the observer’s posterior over leaves remains uniform:

$$\Pr[Z = z] = \frac{1}{H}, \quad z \in 1, \ldots, H$$

The Shannon entropy is therefore maximal [1]:

$$H(Z) = -\sum_{z=1}^{H} \Pr[Z = z] \log_2 \Pr[Z = z] = -H \cdot \frac{1}{H} \log_2 \frac{1}{H} = \log_2 H$$

This is not merely “high entropy.” It is the maximum entropy possible for a system with H leaves.

For a six-leaf manifold:

$$H(Z) = \log_2 6 \approx 2.585 \text{ bits}$$

This means each visible movement is consistent with all six target leaves. The observer sees activity but gains no directional narrowing.

4.4 Why This Matters for Gaze and Dwell-Time Attacks

Suppose an observer tracks eye focus or cursor dwell and claims to infer a “target zone.” In a conventional interface, this can be informative because only one region changes meaningfully at a time [33,34].

Under the Mechanical Displacement Axiom, that logic fails. Every region undergoes a synchronized state transition. A dwell event may correlate with user cognition or motor action, but not with a unique target leaf. The observed behavior is equally compatible with any leaf because the visible manifold updates globally.

Thus, the information available from a biometric side channel is capped by the uniform leaf distribution [33,34]. The adversary’s posterior remains flat at the leaf level before even considering the private morphism.

4.5 Non-Ideal Balance and Entropy Deviation

If implementation imbalance causes slight bias, then the distribution becomes:

$$\Pr[Z = z] = \frac{1}{H} + \epsilon_z \quad \text{with} \quad \sum_{z=1}^{H} \epsilon_z = 0$$

The entropy becomes:

$$H(Z) = -\sum_{z=1}^{H} \left(\frac{1}{H} + \epsilon_z\right) \log_2 \left(\frac{1}{H} + \epsilon_z\right)$$

As long as $|\epsilon_z|$ is small for all z, the entropy remains close to $\log_2 H$. This gives a concrete audit target: measure transition bias and bound entropy loss.

---

5. Hypothesis Space Expansion Under the Private Bijective Morphism

5.1 Extrinsic Witness Principle

The observable witness, when present, is not a direct statement of the target leaf. It is an **extrinsic declaration** in a synonym space. The user (or proving agent) internally applies a private bijective morphism:

$$\Phi : \quad \to \Sigma$$

where:

- is the set of hyperplane leaves (public-facing labels, zones, or color leaves),
- Σ is a disjoint synonym set used for witness declaration.

Because Φ is private, the observer cannot interpret a synonym declaration as a leaf index without first recovering the permutation [16,21,22].

5.2 Factorial Ambiguity

If there are H leaves and the synonym map is bijective, then the number of possible private morphisms is:

$$|\mathrm{Bij}(\quad, \Sigma)| = H!$$

Assuming no side information about Φ, the observer's prior over mappings is uniform:

$$\Pr[\Phi = \phi] = \frac{1}{H!}$$

The entropy of the mapping space is:

$$H(\Phi) = \log_2(H!)$$

For $H = 6$:

$$H(\Phi) = \log_2(6!) = \log_2(720) \approx 9.49 \text{ bits}$$

This is larger than the leaf entropy. That is the key point: the witness declaration does not collapse uncertainty; it **expands** the hypothesis space because the observer must now infer both the target leaf and the hidden permutation [3,4].

5.3 Combined Uncertainty

At a minimum, the observer faces uncertainty over:

1. target leaf Z,
2. private morphism Φ,
3. active secret K (possibly one of many equal-length secrets),
4. session geometry G_r induced by hidden entropy and microsecond time.

A simplified lower bound on uncertainty in the ideal regime is additive when variables are independent under the conditioning assumptions:

$$H(Z, \Phi, K \mid C) \geq H(Z \mid C) + H(\Phi \mid C) + H(K \mid C)$$

The witness declaration cannot be treated as collapsing this space, because it is expressed in the extrinsic synonym space. Without Φ, the declaration does not identify Z.

5.4 Multi-Secret Families and Scenario Expansion

The architecture permits multiple committed secrets under a single private morphism. Let:

$$= K_1, \ldots, K_m$$

be a family of secrets, and let J be the secret index. If the active secret is chosen from the family and equal-length formatting is maintained, then the observable event shape can remain invariant across J.

The desired secrecy condition is:

$$I(J; O \mid C) = 0$$

In plain language, the observer cannot tell which secret in the family was used, even after seeing the transcript and residue.

This is a powerful “many-universes” effect. Each observed event is consistent with multiple hidden-state scenarios, and the system is built to preserve that ambiguity. The hypothesis space does not collapse.

6. Event-Validity Residue and Temporal Fossilization

6.1 The “Bit Leakage” Misinterpretation

A frequent objection is that the system emits a pass/fail result, so it must leak one bit of information. That statement is mathematically incomplete [21,22]. The correct question is: one bit about what variable?

The event-validity residue Y does reveal one bit about a historical event:

- the submitted witness array validated (1) or failed (0)

But the security claim does **not** require Y to be independent of the event itself. It requires that Y not reveal hidden-state coordinates useful for future inference, replay, or secret reconstruction [11,15,23].

6.2 Scenario-Scoped Information

The residue is scenario-scoped to the unique tuple:

$$(K, \Phi, G(\tau))$$

where $G(\tau)$ abbreviates the session geometry induced by hidden entropy and the microsecond time coefficient. The residue says only that the event was valid for that tuple at that instant.

Because the geometry is re-randomized at microsecond granularity, the hidden alignment is destroyed immediately after the event. The residue is therefore a **temporal fossil**: evidence of a past alignment, not a handle on a future one [13,14].

6.3 Formal Non-Predictiveness

Let Y_t be the residue at time $\tau = t$, and let $X_{t'}$ be hidden state or geometry coordinates relevant to a future session at $\tau = t'$. The target property is:

$$I(X_{t'}; Y_t \mid C_t, C_{t'}) = 0 \quad \text{for} \quad t' \neq t$$

This is the practical meaning of non-replayability at the information level. A success residue at time t does not improve prediction for the hidden geometry or valid witness coordinates at time t'.

6.4 Ledger Semantics

This property is especially important when residues are recorded to a ledger or audit log [5]. The log remains useful for governance because it records that a proof event occurred and whether it succeeded. Yet the log does not become a secret-bearing artifact.

That is a major distinction from many systems where logs, traces, or error records become attack material [10,35].

---

7. Formal Security Game: Real vs Simulated Transcripts

7.1 Purpose of the Game

To make the information-theoretic claim precise, define a distinguishing game between [16,17,19]:

- a real transcript generator using the hidden state X,
- a simulator that has no access to X but knows the public session context and output rules.

If an adversary cannot distinguish these games, the observable channel does not carry exploitable structure tied to the hidden state.

7.2 Real Game real

In the real game, the challenger samples hidden state X and for each session i:

1. derives session geometry $G^{(i)}$ from hidden entropy and τ_i,
2. obtains or generates witness behavior consistent with the hidden state,
3. runs the internal accumulator and emits an observable transcript O_i.

The adversary sees the sequence $(O_1, \ldots, O_n)$.

7.3 Simulator Game sim

In the simulator game, a simulator receives only public session contexts $(C_1, \ldots, C_n)$ and the aggregate reporting rules. It does not know X.

It generates synthetic observables $(O'_1, \ldots, O'_n)$ consistent with the public channel format and residue policy.

If the channel is truly zero-information with respect to hidden-state coordinates, then the simulator can match the real transcript distribution without access to X [17,18].

7.4 Distinguishing Advantage

Define the adversary's distinguishing advantage as:

$$ ^{}*{} = | ————————————————————————–

| $$

The ideal target is:

$$\mathrm{Adv}^{\mathrm{dist}} = 0$$

The non-ideal implementation target is:

$$\mathrm{Adv}^{\mathrm{dist}} * \quad \leq \varepsilon * \mathrm{dist}$$

for a small measured bound $\varepsilon_{\mathrm{dist}}$.

7.5 Inference Advantage from Distinguishability

Distinguishability and inference advantage are linked. If real and simulated transcripts are indistinguishable, then no adversary can extract hidden-state predicates from the observable channel with nontrivial advantage.

Thus, for any predicate $\psi(X)$:

$$\mathrm{Adv}^{\mathrm{inf}} * \quad (\psi) \leq \mathrm{Adv}^{\mathrm{dist}} * \quad + \varepsilon_{\mathrm{model}}$$

where $\varepsilon_{\mathrm{model}}$ captures approximation terms from the modeling assumptions.

This is useful for review because it provides a clean proof strategy:

- first bound distinguishability of transcript channels,
- then inherit bounds on inference of secrets, mappings, and future-valid states.

---

8. Quantum Perspective: Why Information Absence Beats Speed

8.1 The Wrong Question About Quantum Attacks

A common quantum-era question is whether Shor [30] or Grover [31] breaks the system. That is the right question for trapdoor cryptosystems and symmetric-key search spaces. It is incomplete here.

The primary defense of the Rosario Primitive is not “the hidden value is too expensive to compute.” It is “the observable channel does not contain the information needed to define the target state.”

This changes the meaning of quantum speedup.

8.2 Shor-Type Attacks

Shor’s algorithm applies to algebraic structures such as integer factorization and discrete logarithms [30]. The Rosario information boundary does not depend on such a trapdoor. There is no exposed algebraic key relation in the observable channel for Shor to exploit [37,38].

So the relevant advantage is:

$$\mathrm{Adv}^{\mathrm{Shor}} \approx 0$$

not because the system is "harder than RSA," but because the attack surface is structurally different [24,26,28,40].

8.3 Grover-Type Attacks

Grover's algorithm provides a quadratic speedup for unstructured search [31]. If an attacker attempts brute-force search over a forgery space of size N, Grover reduces the complexity from $O(N)$ to $O(\sqrt{N})$.

In the Rosario framework, this is still not enough for two reasons:

1. The search space includes factorial and combinatorial uncertainty (private morphisms, secret families, geometry alignment), so N can remain very large.
2. The target itself moves with microsecond-bound geometry. A valid alignment for time τ is generally invalid for time $\tau + \Delta$.

Thus, even if the attacker searches faster, they are chasing a time-varying target. Computation speed does not solve missing information or temporal invalidation [23,32].

8.4 Information-Theoretic Barrier to Inference

Most importantly, Grover does not help against a channel with zero mutual information [2,3]. If:

$$I(X; O \mid C) = 0$$

then no algorithm, classical or quantum, can extract hidden-state information from O because the information is not present.

This is the strongest part of the security story. It is not a statement about computational limits. It is a statement about channel content.

---

9. Implementation Conditions for Maintaining the Information Boundary

The ideal regime is a design target. To preserve it in practice, implementation discipline is required. The most important conditions are structural, not cosmetic.

9.1 Balance of Hyperplane Transitions

The Mechanical Displacement Axiom must be implemented faithfully. Any systematic imbalance in leaf transitions can create bias and reduce entropy [9].

Operational requirement:

- synchronized, uniform leaf updates for each legal movement,
- no hidden UI asymmetries that privilege one leaf.

Audit metric:

- estimate ϵ_z deviations and compute entropy loss from $\log_2 H$.

9.2 Suppression of Round-Level Diagnostics

The black-box accumulator must not emit partial correctness or failure localization.

Operational requirement:

- only full-event residue Y is public,
- no "which round failed," "how many rounds passed," or timing-derived equivalents.

Audit metric:

- inspect API responses, logs, telemetry, and UI feedback for hidden gradients.

9.3 Strict Microsecond Geometry Binding

The verifier and prover must derive geometry from hidden entropy and microsecond-bound τ consistently, and stale geometries must be rejected [5,12].

Operational requirement:

- enforce nonce/time freshness windows,
- bind residue and verification to the exact session context.

Audit metric:

- measure replay acceptance probability across shifted τ.

9.4 Private Morphism Isolation

The private bijective morphism Φ must remain isolated from the observable channel and from analytics or telemetry streams.

Operational requirement:

- no logging of mapping hints,
- no debug traces of internal leaf-to-synonym correspondence,
- no UI artifacts that encode Φ persistently.

Audit metric:

- simulator-based distinguishability testing on witness declarations.

9.5 Equal-Length Secret Family Policy (if used)

If multiple secrets are supported and secret-index privacy is desired, event shapes should remain invariant across active secrets.

Operational requirement:

- equal-length or shape-normalized secret families,
- fixed round count and uniform witness formatting across secret alternatives.

Audit metric:

- test $I(J; O \mid C)$ empirically via classifier indistinguishability.

---

10. Summary of Information-Theoretic Isolation

The architecture can be summarized as three observable categories and three defenses:

10.1 Interface/Biometric Observation

- **Raw observable:** motion, gaze, dwell, keystroke timing
- **Defense:** Mechanical Displacement Axiom
- **Effect:** target leaf uncertainty remains near-uniform, with entropy near $\log_2 H$

10.2 Witness Observation

- **Raw observable:** synonym or response symbol (if visible)
- **Defense:** private bijective morphism Φ
- **Effect:** factorial hypothesis expansion, entropy contribution $\log_2(H!)$

10.3 Validation Outcome Observation

- **Raw observable:** binary pass/fail residue Y
- **Defense:** microsecond-bound hidden geometry and black-box accumulation
- **Effect:** historical residue only; no predictive information for future sessions

These three defenses work together. No single mechanism carries the full burden. The Mechanical Displacement Axiom preserves leaf ambiguity. The private morphism severs semantic interpretation. The microsecond geometry binding destroys replay value and future predictiveness. The black-box accumulator prevents incremental leakage.

That layered design is why the information-theoretic boundary is robust.

---

11. Conclusion

The Rosario Primitive's strongest claim is not merely computational strength. It is a structured information-theoretic boundary between hidden state and observable channel.

In the ideal regime, the observable channel satisfies:

$$I(X; O \mid C) = 0$$

meaning the public artifacts of a proof event do not reduce uncertainty about the hidden state coordinates required to reconstruct secrets, private mappings, or future-valid witness alignments [1,2,3]. In the non-ideal regime, the architecture admits explicit leakage bounds tied to measurable implementation imbalance rather than to the cryptographic core.

This distinction matters. It means the system can be evaluated and improved using concrete entropy and distinguishability metrics. It also means the public validation output can remain useful for audit and governance without becoming a secret-bearing artifact.

The event-validity residue is therefore best understood as a historical fossil: a record that a specific hidden-state alignment succeeded under a specific microsecond-bound geometry. It is not a reusable credential [6,7,10], not a revealing bit of the secret [11,15], and not a predictor of future validity [5]. That is the central information-theoretic insight of the Rosario Primitive.

For authentication and attestation systems, this is a meaningful shift in design philosophy. Security is no longer defined only by how hard it is to compute the secret. It is defined by whether the observable channel ever carries the secret's coordinates in the first place [5,6,8].

Appendix A: Compact Formula Reference

Hidden State and Observable Channel

$$X = (K, \Phi, E, \tau, \Theta)$$

$$O = (C, W, Y) \quad \text{or} \quad O = (C, Y)$$

Session Geometry

$$G_r = f(E, \tau, r, \Theta)$$

Internal Accumulation and Public Residue

$$\Lambda = \prod_{r=1}^{R} M_r$$

$$Y = \mathbb{1}_{\Lambda = 1}$$

Ideal and Bounded Leakage

$$I(X; O \mid C) = 0$$

$$I(X; O \mid C) \leq \varepsilon$$

Leaf Entropy Under Balanced Displacement

$$H(Z) = \log_2 H$$

Private Morphism Entropy

$$H(\Phi) = \log_2(H!)$$

Secret Family Index Hiding

$$I(J; O \mid C) = 0$$

Distinguishing Advantage

$$ ^{}*{} = | ————————————————————–

| $$

Appendix B: Plain-English Reading of the Main Claim

The public channel can tell an observer that an event succeeded. It does not tell the observer what secret was used, how the private map was configured, which hidden geometry made it valid, or how to repeat that success later [2,3]. The system is designed so that every visible clue points to many possible hidden explanations, and the one thing the observer can learn (success/failure) expires immediately as a useful signal because the hidden geometry changes with time [6,10,35].

Acknowledgments

(To be completed. Acknowledgements of funders, reviewers, colleagues, or institutions that supported the work.)

Author Contributions

(To be completed. CRediT-style: conceptualization, formal analysis, writing—original draft, writing—review and editing, etc. Example: "A.B.: conceptualization, formal analysis, writing—original draft. C.D.: validation, writing—review and editing.")

Funding

(To be completed. "This work received no external funding" or list grants, project numbers, and funding agencies.)

Competing Interests

The author(s) declare no competing interests. *(Or specify: patent applications, advisory roles, etc.)*

Data Availability

Data sharing is not applicable to this article—no datasets were generated or analyzed. The paper presents a theoretical and formal analysis with no empirical data.

References

[1] Claude E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal* (1948). URL: https://archive.org/details/bstj27-3-379 (Foundational source for entropy, information, and channel capacity concepts used in the secrecy analysis.)

[2] Claude E. Shannon, "Communication Theory of Secrecy Systems," *Bell System Technical Journal* (1949). URL: https://archive.org/details/bstj28-4-656 (Primary source for perfect secrecy and the formal condition behind Shannon-style hiding.)

[3] Thomas M. Cover and Joy A. Thomas, *Elements of Information Theory* (2nd ed., Wiley). URL: https://onlinelibrary.wiley.com/doi/book/10.1002/047174882X (Standard graduate-level reference for entropy, mutual information, conditional entropy, and KL divergence.)

[4] David J.C. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press). URL: https://www.inference.org.uk/mackay/itila/ (Useful for intuitive and formal treatment of information-theoretic reasoning, including uncertainty and inference limits.)

[5] U.S. National Institute of Standards and Technology (NIST), *Digital Identity Guidelines: Authentication and Lifecycle Management* (SP 800-63B). URL:

https://csrc.nist.gov/pubs/sp/800/63/b/final (Authoritative guidance on authentication, memorized secrets, replay resistance considerations, and modern identity policy.)

[6] W3C, *Web Authentication: An API for Accessing Public Key Credentials Level 2 (WebAuthn Level 2)*. URL: https://www.w3.org/TR/webauthn-2/ (Primary standard for modern passkey/FIDO-style authentication flows; useful for comparison against reusable-credential models.)

[7] W3C, *Web Authentication: An API for Accessing Public Key Credentials Level 3 (WebAuthn Level 3)*. URL: https://www.w3.org/TR/webauthn-3/ (Current evolution of the WebAuthn standard; useful for current-state attestation/auth comparison.)

[8] FIDO Alliance, *UX Guidelines for Passkey Creation and Sign-ins* (PDF). URL: https://fidoalliance.org/wp-content/uploads/2023/05/FIDO-Alliance-UX-Guidelines-for-Passkey-Creation-and-Sign-ins.pdf (Helpful for contrasting ENI6MA onboarding and mnemonic design with mainstream authentication UX patterns.)

[9] OWASP Cheat Sheet Series, *Authentication Cheat Sheet*. URL: https://cheatsheetseries.owasp.org/cheatsheets/Authentication_Cheat_Sheet.html (Industry-practical guidance on brute force mitigation, auth controls, and operational protections.)

[10] OWASP Cheat Sheet Series, *Credential Stuffing Prevention Cheat Sheet*. URL: https://cheatsheetseries.owasp.org/cheatsheets/Credential_Stuffing_Prevention_Cheat_Sheet.html (Direct support for claims about credential stuffing and the weaknesses of reusable credentials.)

[11] OWASP Cheat Sheet Series, *Password Storage Cheat Sheet*. URL: https://cheatsheetseries.owasp.org/cheatsheets/Password_Storage_Cheat_Sheet.html (Useful for comparison with hash-based password storage and why ENI6MA differs from hash-comparison architectures.)

[12] IETF RFC 4086, *Randomness Requirements for Security* (Best Current Practice). URL: https://datatracker.ietf.org/doc/html/rfc4086 (Core reference for entropy sourcing, entropy pools, randomness pitfalls, and cryptographic unpredictability requirements.)

[13] IETF RFC 4226, *HOTP: An HMAC-Based One-Time Password Algorithm*. URL: https://www.ietf.org/rfc/rfc4226.txt (Useful contrast source: OTP systems are time/nonce-like at the output layer but still depend on long-lived shared secrets.)

[14] IETF RFC 6238, *TOTP: Time-Based One-Time Password Algorithm*. URL: https://www.ietf.org/rfc/rfc6238.txt (Contrast source for time-based OTP authentication and shared-seed architectures.)

[15] IETF RFC 9106, *Argon2 Memory-Hard Function for Password Hashing and Proof-of-Work Applications*. URL: https://datatracker.ietf.org/doc/rfc9106/ (Strong modern password hashing reference; useful in comparisons against ENI6MA's non-hash, witness-based model.)

[16] Oded Goldreich (ed.), *Foundations of Cryptography* (Cambridge University Press). URL: https://www.wisdom.weizmann.ac.il/~oded/foc.html (Foundational theoretical cryptography reference for commitment, proof systems, and formal security games.)

[17] Shafi Goldwasser, Silvio Micali, Charles Rackoff, "The Knowledge Complexity of Interactive Proof Systems" (STOC 1985). URL: https://dl.acm.org/doi/10.1145/22145.22178 (Canonical source for interactive proof systems and knowledge complexity; baseline reference for ZK comparison.)

[18] Silvio Micali (hosted copy), "The Knowledge Complexity of Interactive Proof Systems" (PDF reprint). URL: https://people.csail.mit.edu/silvio/Selected%20Scientific%20Papers/Proof%20Systems/The_Knowledge_Complexity_Of_Interactive_Proof_Systems.pdf (Convenient PDF access to the same core result, for reviewers who want direct text.)

[19] Boneh and Shoup, *A Graduate Course in Applied Cryptography* (free draft textbook). URL: https://crypto.stanford.edu/~dabo/cryptobook/ (Strong practical/theoretical source for commitments, security notions, and protocol design patterns.)

[20] Menezes, van Oorschot, and Vanstone, *Handbook of Applied Cryptography* (free online). URL: https://cacr.uwaterloo.ca/hac/ (Reference handbook for classical cryptographic primitives, protocol mechanics, and security terminology.)

[21] Torben P. Pedersen, "Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing," *CRYPTO '91* (Springer LNCS). URL: https://link.springer.com/chapter/10.1007/3-540-46766-1_9 (Important background reference on information-theoretic security properties and commitment-style reasoning; useful for binding/hiding discussions.)

[22] Moni Naor, "Bit Commitment Using Pseudorandomness," *Journal of Cryptology* / CRYPTO-era references. URL: https://link.springer.com/chapter/10.1007/0-387-34805-0_19 (Classic commitment reference, useful for contrast with Rosario commitment semantics.)

[23] Philippe Oechslin, "Making a Faster Cryptanalytic Time-Memory Trade-Off," *CRYPTO 2003* (Rainbow Tables). URL: https://link.springer.com/chapter/10.1007/978-3-540-45146-4_36 (Standard citation for rainbow tables and precomputation attacks on static hashed secrets.)

[24] NIST CSRC, *FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard (ML-KEM)*. URL: https://csrc.nist.gov/pubs/fips/203/final (Official NIST PQC standard page; useful for contrast against lattice-based computational assumptions.)

[25] NIST (PDF), *FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard*. URL: https://nvlpubs.nist.gov/nistpubs/fips/nist.fips.203.pdf (Direct PDF for technical reviewers.)

[26] NIST CSRC, *FIPS 204: Module-Lattice-Based Digital Signature Standard (ML-DSA)*. URL: https://csrc.nist.gov/pubs/fips/204/final (Official NIST PQC signature standard page.)

[27] NIST (PDF), *FIPS 204: Module-Lattice-Based Digital Signature Standard*. URL: https://nvlpubs.nist.gov/nistpubs/fips/nist.fips.204.pdf (Direct PDF for technical reviewers.)

[28] NIST CSRC, *FIPS 205: Stateless Hash-Based Digital Signature Standard (SLH-DSA)*. URL: https://csrc.nist.gov/pubs/fips/205/final (Official hash-based PQC standard page; useful because "stateless" semantics are often confused with stronger information-theoretic hiding claims.)

[29] NIST (PDF), *FIPS 205: Stateless Hash-Based Digital Signature Standard*. URL: https://nvlpubs.nist.gov/nistpubs/fips/nist.fips.205.pdf (Direct PDF for technical reviewers.)

[30] Peter W. Shor, "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer" (1994/1995). URL: https://arxiv.org/abs/quant-ph/9508027 (Core citation for quantum attacks on RSA/ECC-style algebraic trapdoor systems.)

[31] Lov K. Grover, "A Fast Quantum Mechanical Algorithm for Database Search" (1996). URL: https://arxiv.org/abs/quant-ph/9605043 (Core citation for quadratic speedup in unstructured search; useful for analyzing residual brute-force search spaces.)

[32] P. C. van Oorschot and M. J. Wiener, "Parallel Collision Search with Cryptanalytic Applications" (Journal of Cryptology / related literature). URL: https://link.springer.com/article/10.1007/PL00003816 (Useful for classical search economics and attack scaling comparisons.)

[33] Song, Wagner, and Tian, "Timing Analysis of Keystrokes and Timing Attacks on SSH" (UC Berkeley). URL: https://people.eecs.berkeley.edu/~daw/papers/ssh-use01.pdf (Strong support for side-channel leakage claims and the need for interaction designs that avoid direct coordinate/keystroke inference.)

[34] Eiband et al. / Systematic Review source, "Shoulder Surfing Experiments: A Systematic Literature Review," *Computers & Security* (ScienceDirect). URL: https://www.sciencedirect.com/science/article/pii/S0167404820302960 (Useful support for observational attack models and interface-side surveillance threats.)

[35] OWASP, *Credential Stuffing* (Attack overview). URL: https://owasp.org/www-community/attacks/Credential_stuffing (High-level attack taxonomy source; pairs well with the prevention cheat sheet.)

[36] RFC Editor / IETF Datatracker, *BCP 106 (RFC 4086) History and Status* (metadata page). URL: https://datatracker.ietf.org/doc/bcp (Useful metadata reference for standards status and best-current-practice classification.)

[37] Google Patents, *US4405829A — Cryptographic communications system and method (RSA)*. URL: https://patents.google.com/patent/US4405829A/en (Useful historical contrast reference for trapdoor-based public-key systems.)

[38] Google Patents, *US4200770A — Cryptographic apparatus and method (Diffie-Hellman / public key exchange era)*. URL: https://patents.google.com/patent/US4200770A/en (Useful historical contrast reference for key-exchange paradigms centered on reusable key material.)

[39] W3C / FIDO ecosystem background (Passkeys and FIDO deployment context) — FIDO Alliance resources hub. URL: https://fidoalliance.org/ (Useful for standards/ecosystem context and implementation adoption references.)

[40] NIST CSRC, *Post-Quantum Cryptography Project* (program page). URL: https://csrc.nist.gov/projects/post-quantum-cryptography (Program-level reference for PQC transition context and standards timeline.)