Eni6ma - The Computationally Intractable Cypher

Implementation of Shannon's Confusion and Diffusion for Perfect Secrecy

by Frank Dylan Rosario and Dr. Lin Wang

A Membership-Only Proof of Knowledge with Balanced Partitions: An Essay for Students and Researchers

Abstract

We present a coherent, paragraph-style exposition of a **membership-only proof-of-knowledge** protocol in which a prover, Alice, convinces an observer that she knows a collection of secrets without revealing them. The scheme relies on three intertwined ideas: (i) **balanced partitions** of a fixed alphabet into six equal "leaves" per round; (ii) **independent ring rotations** that eliminate positional anchors across rounds; and (iii) a **private bijection** that permutes the six leaf labels for Alice alone. Together these choices throttle information leakage to at most $\log_2 6$ bits per round and, in expectation under the randomizations, to essentially **zero useful signal** for predicting the next character. The adversary's task collapses to **unstructured search** over a vast hypothesis space of the form $|\mathcal{H}| = 6!\binom{U}{6} = P(U, 6)$, where $U = |\Sigma|^\ell$ counts all ℓ-symbol strings over the alphabet Σ (with $|\Sigma| = 72$ in the canonical parameterization). For $\ell = 6$ and six distinct secrets, $|\mathcal{H}| \approx 72^{36} \approx 2^{222}$, rendering classical brute force infeasible and quantum brute force (Grover/BBBV) still astronomically out of reach. We formalize the **rounds-to-uniqueness** requirement $R \gtrsim \log_6 |\mathcal{H}|$ (about 86 rounds at baseline) and characterize how to reach a formal, **post-quantum 128-bit** margin by minimally increasing the secret length to $L = 7$, yielding $|\mathcal{H}| \approx 2^{259}$ and $\sqrt{|\mathcal{H}|} \approx 2^{129.5}$. The protocol illustrates a design space where **combinatorics, symmetry, and information throttling** jointly enforce intractability.

1. Introduction and Motivation

Traditional authentication reveals too much: a password, a biometric template, or a static token leaves residue that adversaries can hoard, correlate, and replay. Zero-knowledge (ZK) protocols resolve this tension cryptographically but often depend on specialized algebraic assumptions and heavy machinery. The construction analyzed here pursues a complementary path: a **proof-of-knowledge ceremony** in which Alice answers only with **leaf identifiers**, never with symbols or algebraic witnesses, while the system geometry ensures that those identifiers are (in expectation) **indistinguishable from uniform noise** to an eavesdropper.

The core idea is simple. Each "board" partitions the alphabet Σ into six equal **leaves** (zones). The board is **balanced** so that each leaf contains the same number of symbols. Each round, the rings that order letters and digits are **independently rotated**, and Alice's **private bijection** permutes the six leaf labels into private codewords. When challenged for the next character of a secret, Alice responds only with the codeword naming the leaf containing that character under the current board. To an observer, the codewords look like independent throws of a fair six-sided die.

Our goal in this essay is to articulate the **threat model**, quantify the **hypothesis space** the adversary must explore, derive **information-theoretic bounds** on leakage per round, and connect these to **lower bounds** for classical and quantum brute force. The result is a protocol whose security derives not from algebraic trapdoors but from **combinatorics without structure** and **symmetry without bias**.

2. Entities, Alphabets, and Boards

Let Σ denote a fixed alphabet of size $|\Sigma| = 72$, comprised of 30 lowercase letters, 30 uppercase letters, and 12 digits. A **secret** is a string $C \in \Sigma^\ell$ of length ℓ. The baseline parameter set fixes $\ell = 6$, and Alice holds **six distinct** secrets $S = \{C^{(1)}, \ldots, C^{(6)}\} \subset \Sigma^\ell$.

A **board** for round i is a balanced partition

$$\Pi_i : \Sigma \to \{1, 2, 3, 4, 5, 6\}, \qquad |\Pi_i^{-1}(j)| = |\Sigma|/6 = 12 \ \ \forall j,$$

together with independent **ring rotations** applied to the subalphabets (e.g., lowercase, uppercase, digits). The rotations eliminate any positional relation across rounds; the balancing eliminates marginal biases. Alice holds a **private bijection** $\varphi \in \mathbb{S}_6$ that permutes the public leaf identifiers into her private codewords.

When the verifier requests the next character of a given secret under board Π_i, Alice locates that character in Σ, looks up its leaf $j = \Pi_i(C_i)$, and replies with the **codeword** $Y_i = \varphi^{-1}(j) \in \{1, \ldots, 6\}$. No symbol of Σ is ever spoken; only the codeword Y_i appears on the wire.

3. Threat Model and Problem Statement

Assume an **observer** (passive adversary) records the transcript of codewords $Y_1, \ldots, Y_R$ for R rounds, together with the sequence of public boards $B = (\Pi_1, \ldots, \Pi_R)$. The adversary seeks to recover (i) the six distinct secrets $S \subset \Sigma^\ell$ and (ii) Alice's private bijection $\varphi \in \mathbb{S}_6$.

Formally, the **hypothesis set** is

$$\mathcal{H} = \{(S, \varphi) : S \subset \Sigma^\ell,\ |S| = 6 \text{ distinct},\ \varphi \in \mathbb{S}_6\}.$$

Given B and a transcript Y, the adversary's task is to identify the **marked element** $(S^\star, \varphi^\star) \in \mathcal{H}$ consistent with the transcript. Because all observable outputs are leaf labels and **not** symbols, any progress requires exploiting statistical deviations of Y from uniformity **or** performing **eliminative search** over $\mathcal{H}$. The protocol is engineered so the former is absent and the latter is astronomically expensive.

4. Combinatorics of the Hypothesis Space

Let $U = |\Sigma|^\ell$ denote the number of possible secrets of length ℓ. With $|\Sigma| = 72$ and $\ell = 6$, we have $U = 72^6$. Because Alice holds **six distinct** secrets and a private bijection, the number of hypotheses is

$$|\mathcal{H}| = 6! \binom{U}{6} = P(U, 6) = U(U-1)(U-2)(U-3)(U-4)(U-5).$$

When $U \gg 6$, this is well-approximated by U^6. Numerically,

$$|\mathcal{H}| \approx 72^{36} \approx 2^{36 \log_2 72} \approx 2^{222},$$

using $\log_2 72 \approx 6.17$. This count captures **both** the combinatorial choice of six distinct secrets and the six-way label symmetry embedded by φ.

Two immediate consequences follow:

1. There is **no sparse structure** to index into: the hypothesis set is a flat product of choices.
2. Any attacker who cannot extract stable information from the transcript is reduced to **unstructured search** over $\mathcal{H}$.

5. Per-Round Observations and Information Throttling

In round i, Alice outputs a single label $Y_i \in \{1, \ldots, 6\}$. Under a balanced partition and independent ring rotations, each symbol of Σ is equally likely to reside in any leaf, so for an observer without φ,

$$\Pr(Y_i = j \mid B) \approx \frac{1}{6} \quad \text{for all } j,$$

hence the **per-round mutual information** is bounded by the entropy of a fair six-sided die:

$$I(C; Y_i \mid B) \leq H(Y_i \mid B) = \log_2 6 \approx 2.585 \text{ bits}.$$

More sharply, because the rotations and balanced partitions wash out correlations, the **expected** mutual information per round about the *next* character (conditioned on the board) is essentially

$$\mathbb{E}[I(C_i; Y_i \mid B)] \approx 0,$$

so transcripts look like i.i.d. draws from a uniform six-ary alphabet. Intuitively, the only way to distinguish competing hypotheses is to test whether their implied leaf labels match the observed codewords **across all rounds**—that is, to run a consistency oracle, not to learn from frequencies or correlations.

6. Eliminative Consistency and Rounds-to-Uniqueness

Suppose the adversary fixes a candidate hypothesis (S, φ). The probability that it matches the observed codeword in one round is $\Pr(\text{match}) = 1/6$ (absent privileged information). Under independence, the probability that a **wrong** hypothesis survives R rounds is $(1/6)^R$. Hence the **expected number of survivors** after R rounds is

$$\mathbb{E}[\text{survivors}] \approx |\mathcal{H}|(1/6)^R.$$

A natural **uniqueness threshold** is when the survivor count drops below one in expectation:

$$|\mathcal{H}|(1/6)^R \lesssim 1 \quad \Longleftrightarrow \quad R \gtrsim \log_6 |\mathcal{H}|.$$

At baseline, $\log_2 |\mathcal{H}| \approx 222$ and $\log_2 6 \approx 2.585$, so

$$R_{\min} \approx \left\lceil \frac{\log_2 |\mathcal{H}|}{\log_2 6} \right\rceil \approx \left\lceil \frac{222}{2.585} \right\rceil \approx 86.$$

This is a **theoretical lower bound** under perfect transcripts and independence. In practice, transcripts are incomplete or noisy for the adversary, lengthening the effective search.

7. Classical Lower Bounds: Black-Box Search

In a **black-box** view, the only operation available to an adversary is to propose a hypothesis (S, φ) and check it against the transcript. Let the per-candidate check cost be $O(R)$. Then the time to find the unique marked hypothesis is

$$T_{\text{classical}} = \Omega(R \cdot |\mathcal{H}|) = \Omega\big(R \cdot 72^{36}\big).$$

Even with a wildly optimistic sustained rate of 10^{15} checks per second, a full scan of 2^{222} candidates would require on the order of

$$\frac{2^{222}}{10^{15}\ \text{s}^{-1}} \approx 1.9 \times 10^{44}\ \text{years},$$

vastly exceeding cosmological timescales. The key point is not the exact figure but the **order of magnitude**: the exponent in $|\mathcal{H}|$ dominates any plausible constant-factor engineering improvements.

8. Quantum Lower Bounds: Grover/BBBV

Quantum search over an **unstructured** set yields at most a square-root speedup: finding a marked element with bounded error requires $\Omega(\sqrt{N})$ oracle calls for N candidates. Translating to our setting,

$$T_{\text{quantum}} = \Omega\big(R \cdot \sqrt{|\mathcal{H}|}\big) = \Omega\big(R \cdot 72^{18}\big) = \Omega\big(R \cdot 2^{111}\big).$$

Even if one supposes a fantastical 10^{18} fully error-corrected oracle queries per second and ignores the massive overhead of **reversible** circuit construction for the transcript-consistency oracle, the wall-clock time remains on the order of

$$\frac{2^{111}}{10^{18}\ \text{s}^{-1}} \approx 7.9 \times 10^{7}\ \text{years}.$$

Parallel quantum fleets help only by $\sqrt{p}$: a trillion independent quantum processors (already beyond any credible roadmap) reduces this figure by 10^6, still leaving decades-to-millennia under heroic assumptions. With realistic error correction (e.g., surface-code overheads of $10^3 - 10^6$ physical qubits per logical qubit and deep circuits per oracle query), the practical gap widens dramatically.

9. Security Levels and "Minimum Permutations" for Intractability

Security targets are typically expressed in **bits of work**. A common classical target is **128-bit security**, i.e., an attacker must perform on the order of 2^{128} steps. For quantum adversaries restricted to Grover-style search, a **post-quantum 128-bit** target corresponds to requiring $\sqrt{|\mathcal{H}|} \geq 2^{128}$, i.e.,

$$|\mathcal{H}| \;\geq\; 2^{256}.$$

9.1 Classical Target

At baseline,

$$|\mathcal{H}| \approx 72^{36} \approx 2^{222} \quad \Rightarrow \quad \text{classical margin} \approx 2^{94} \text{ beyond } 2^{128}.$$

Thus the scheme already **comfortably exceeds** the classical 128-bit bar.

9.2 Quantum Target and Minimal Adjustment

To meet the stricter $|\mathcal{H}| \geq 2^{256}$ quantum-aware bar with minimal change, increase the **secret length** from $\ell = 6$ to $L = 7$. With six distinct secrets,

$$|\mathcal{H}| \approx (|\Sigma|^L)^6 = |\Sigma|^{6L} = 72^{42} \approx 2^{6L \log_2 72} \approx 2^{259}.$$

Then

$$\sqrt{|\mathcal{H}|} \approx 2^{129.5},$$

which cleanly exceeds the post-quantum 128-bit target. Two other knobs can raise $|\mathcal{H}|$ without changing L: increase the **number of secrets** beyond six or expand the **alphabet** $|\Sigma|$ (e.g., adding symbols or modalities). All three knobs can be tuned jointly for desired margins.

10. Rounds Budget and Stopping Rules

From Section 6, a natural **stopping rule** that aims for uniqueness in expectation is

$$R_{\min} \approx \left\lceil \frac{\log_2 |\mathcal{H}|}{\log_2 6} \right\rceil.$$

At $L = 6$, this yields $R_{\min} \approx 86$; at $L = 7$, $\log_2 |\mathcal{H}| \approx 259$ gives

$$R_{\min} \approx \left\lceil \frac{259}{2.585} \right\rceil \approx 101.$$

These thresholds are **idealized**: they presuppose that the adversary possesses the full public board history and a clean transcript of every round. In realistic deployments, **partial visibility** (e.g., hidden or delayed boards, subsampling, or decoy rounds) increases the adversary's uncertainty and effectively lowers the number of real rounds needed to achieve the same safety.

11. Why Structurelessness Matters

Many cryptanalytic successes exploit **structure**: linear approximations in block ciphers, differential patterns in S-boxes, algebraic relations in group-based assumptions, or patterns in side-channel leakage. The present construction deliberately suppresses structure at three levels:

1. **Uniform marginals:** by balancing each board, the distribution $\Pr(Y = j)$ is flat, killing frequency analysis.
2. **Cross-round independence:** independent ring rotations prevent positional correlations from accumulating across rounds.
3. **Label-switching symmetry:** the private bijection φ means any fixed relabeling produces observationally equivalent transcripts.

The net effect is that **no gradient exists** for the attacker: there is no algebraic scaffold to climb, no statistical bias to amplify. The only viable tactic is breadth-first elimination across a hypothesis set whose size is exponential in both $|\Sigma|$ and the total secret length under consideration.

12. Engineering Realities: Oracles, I/O, and Energy

Even if one were to imagine a brute-force engine that could instantiate the transcript-consistency check as an oracle, practical limitations loom:

- **Reversible oracles** for quantum search must embed the full consistency logic—mapping a hypothesis and public boards to a match bit—within a fault-tolerant circuit. This costs large numbers of logical qubits and deep gate depth, multiplied by R.
- **I/O and memory pressure** for classical brute force grow with the need to enumerate, store, or stream hypotheses; paging and bandwidth become dominant bottlenecks long before arithmetic throughput saturates.
- **Energy bounds** (e.g., Landauer's limit) and thermodynamic inefficiencies imply that, at cosmological scales of work, the energy bill alone becomes prohibitive—quite apart from the time dimension.

These are not security proofs; they are **engineering sanity checks** reinforcing the asymptotic arguments.

13. Variants, Extensions, and Defense-in-Depth

The protocol framework accommodates several **hardening variants** without altering its conceptual simplicity:

- **Extended alphabets and modalities.** Replace the 72-symbol alphabet with multimodal symbol families (e.g., audio cues, icons), preserving balanced partitions. Increasing $|\Sigma|$ multiplies $|\mathcal{H}|$.
- **More secrets or longer secrets.** Raising the number of distinct secrets beyond six or extending length L increases the exponent linearly in the chosen parameter.
- **Board-hiding and decoys.** If public boards are partially hidden, delayed, or salted with decoys, the adversary must guess extra degrees of freedom, further inflating the search. For instance, if per-round there are N_B plausible board states from the attacker's perspective, the attack space multiplies by N_B^R. With three independently rotated rings of sizes $30, 30, 12$, even simple rotation uncertainty can contribute factors like $(30 \cdot 30 \cdot 12)^R = 10{,}800^R$ to the adversary's effective space, though the exact factor depends on what is public versus private in a specific design.
- **Rate limiting and ceremony mixing.** Interleaving multiple secrets and inserting dummy rounds complicate transcript alignment for an observer.

These levers increase robustness without sacrificing the protocol's intuitive pedagogy.

14. Limitations and Assumptions

No construction is omnipotent; clarity about assumptions is essential:

- **Randomness quality.** Balanced partitions and independent ring rotations must be generated with high-quality randomness to avoid subtle biases.
- **Transcript exposure.** Our bounds assume the adversary can see the boards and the codewords. If more is public, the analysis must account for it; if less is public, security improves.
- **Active adversaries.** The analysis here targets passive eavesdroppers. Active, man-in-the-middle settings require standard countermeasures (e.g., authenticated channels, challenge binding, and anti-replay tokens) to prevent transcript splicing or relay attacks.
- **Side channels.** As with any interactive system, timing, power, or emanation side channels require engineering discipline. The protocol's virtue is that its **semantic output** is low-bandwidth and uniform; nevertheless, implementation must be hardened.

15. Synthesis: Why Brute Force Is the Only Game in Town

Collecting the threads:

1. The adversary's **goal** is to recover (S, φ).
2. The **hypothesis count** is $|\mathcal{H}| = 6!\binom{U}{6}$ with $U = |\Sigma|^{\ell}$.
3. Balanced partitions and ring rotations ensure $\Pr(Y = j \mid B) \approx 1/6$ and $\mathbb{E}[I(C; Y \mid B)] \approx 0$; there is no statistically exploitable drift.
4. The only viable method is **consistency elimination** over $\mathcal{H}$: expected survivors $|\mathcal{H}|(1/6)^R$, with $R \gtrsim \log_6 |\mathcal{H}|$ rounds for uniqueness.
5. **Lower bounds** force classical cost $\Omega(R|\mathcal{H}|)$ and quantum cost $\Omega(R\sqrt{|\mathcal{H}|})$, both astronomical at baseline parameters.

Therefore, the construction lives precisely where **information theory and computational lower bounds align**: a regime of **structureless combinatorics** in which neither clever algebra nor statistical learning can short-circuit the exponential.

16. Practical Parameter Recipes

For practitioners wishing to tie security targets to parameters, the following recipes are convenient:

- **Given $|\Sigma|$ and length L with six distinct secrets**,
$$|\mathcal{H}| \approx |\Sigma|^{6L} \quad \text{and} \quad \log_2 |\mathcal{H}| \approx 6L \log_2 |\Sigma|.$$
- **Classical 128-bit security** requires $\log_2 |\mathcal{H}| \geq 128$. With $|\Sigma| = 72$, this is already satisfied for $L = 6$ by a margin of roughly 94 bits.
- **Post-quantum 128-bit security** requires $\frac{1}{2} \log_2 |\mathcal{H}| \geq 128$, i.e., $\log_2 |\mathcal{H}| \geq 256$. With $|\Sigma| = 72$, this amounts to
$$6L \log_2 72 \;\geq\; 256 \quad \Longrightarrow \quad L \;\geq\; \left\lceil \frac{256}{6 \log_2 72} \right\rceil = 7.$$
- **Rounds budget** may be set via
$$R_{\min} \approx \left\lceil \frac{\log_2 |\mathcal{H}|}{\log_2 6} \right\rceil,$$
then padded for engineering margin (e.g., noise, partial visibility, decoys).

These formulas cleanly separate **symbolic design choices** (alphabet size, secret length, number of secrets) from **ceremony design choices** (round count, board exposure), allowing system builders to trade off usability and security.

17. Pedagogical Value

For naive audiences, the construction demonstrates how **combinatorics** and **symmetry** can substitute for heavy algebra while still achieving formidable security. Concepts like balanced partitions, entropy per round, and hypothesis counting are concrete and visually intuitive ("six equal baskets each round; labels shuffled privately"). For researchers, the protocol offers a compact laboratory to test **information-theoretic** intuitions (leakage bounds, Fano-style reasoning) alongside **computational** lower bounds (black-box and Grover limits).

18. Conclusion

We presented a paragraph-style, scientifically rigorous account of a membership-only proof-of-knowledge protocol whose defense rests on three pillars: **balanced leaves**, **independent rotations**, and a **private bijection**. These features ensure that observed outputs are **uniform and memoryless** in expectation, starving the adversary of exploitable signal and collapsing the attack surface to **unstructured search** over a hypothesis space of size $|\mathcal{H}| = 6!\binom{|\Sigma|^\ell}{6}$. At baseline ($|\Sigma| = 72$, $\ell = 6$, six secrets), $|\mathcal{H}| \approx 2^{222}$, implying classical costs utterly beyond reach and quantum costs still astronomical even under heroic assumptions. A minimal adjustment to $L = 7$ delivers a formal **post-quantum 128-bit** margin. The scheme thus exemplifies a design philosophy—**combinatorics + symmetry + information throttling**—that yields powerful guarantees without dependence on specialized algebraic hardness, and it provides a clear, analyzable path for tuning parameters to desired security levels while maintaining conceptual and implementation simplicity.

Why Cracking All Six of ENI6MA's Secrets (and the Private Map) Is Computationally Intractable

Audience: first-year undergrads in CS and psychology.
*Goal: Why an attacker with **no prior knowledge** cannot feasibly brute-force Alice's six secrets (each 6 characters) and her private bijection map, even with massive hardware, clever statistics, or machine learning.*

1) What "computationally intractable" actually means (for us)

"Intractable" doesn't mean impossible; it means that the **time and resources grow so explosively** with problem size that **no realistic computer (or fleet of computers)** can finish within the age of the universe. In cryptography, we often summarize this by counting the **size of the search space** and comparing it to **feasible compute budgets**. If the search demands, say, around 2^{200} or 10^{60} steps, it is effectively out of reach—no matter how clever your code is or how many GPUs you rent—because physics, energy, and time set hard limits. In

our ceremony, the intractability comes from a **product** of three design choices that eliminate shortcuts:

1. **Huge combinatorial space** of possible secrets and private maps;
2. **No structure** to exploit (balanced leaves + opaque bijection + independent ring rotations);
3. **Tiny information per round** (only "which leaf this round?"), which forces an attacker to collect **a lot** of perfect observations just to narrow the field.

Those three features make the best attack essentially a **blind enumeration/elimination** game with a space so large that even perfect parallelism barely dents it.

2) The puzzle the attacker must solve (with no head start)

Alice's world has six leaves on the canvas. She holds **six different secrets**, each a 6-character string. The characters come from three alphabet rings of sizes $[30, 30, 12]$, so there are 72 possible symbols per character. Alice also keeps a **private bijection** that maps **six codewords** (her private synonyms) to the **six leaves**. The attacker (call them Eve) wants **everything**: **all six** secrets **and** the **entire** private map.

Let's formalize the space Eve must search. The total number of possible 6-character strings is

$$U \;=\; 72^6.$$

"U equals seventy-two to the sixth."

Numerically, $U = 139{,}314{,}069{,}504 \approx 1.393 \times 10^{11}$. Now, how many different **sets** of six distinct secrets can Alice hold? If we care about the six as a **set** (order doesn't matter), that count is $\binom{U}{6}$. If we instead care about an **ordered 6-tuple** of distinct secrets (order does matter), that count is the permutation number $P(U, 6)$. The private map (the synonym-to-leaf pairing) has

$$6! \;=\; 720$$

"Six factorial equals seven hundred twenty."

possibilities. A key identity links these counts:

$$6!\binom{U}{6} \;=\; P(U,6) \;=\; U(U-1)(U-2)(U-3)(U-4)(U-5).$$

"Six factorial times U choose six equals P of U six equals U times U minus one down to U minus five."

This identity tells us something very convenient: **"(unordered six secrets + map)"** and **"(ordered six secrets, no map)"** describe the **same number**

of global possibilities. So, to recover **all six** secrets **and** the private map, Eve's exact hypothesis space is

$$\boxed{P(U,6) \;=\; U(U-1)(U-2)(U-3)(U-4)(U-5)}.$$

"The exact global count equals U times U minus one down to U minus five."

Plugging in $U = 72^6$, this is astronomically large. Since $U \gg 5$, $P(U,6)$ is essentially U^6, i.e., 72^{36}. Converting to "bits of difficulty" (a common crypto yardstick), note that $\log_2(72) \approx 6.17$, so

$$P(U,6) \;\approx\; 72^{36} \;=\; 2^{36\,\log_2 72} \;\approx\; 2^{222}.$$

"P of U six is about two to the two hundred twenty-two."

To give a decimal feel, 2^{222} is about 6.0×10^{66}; tighter arithmetic on $P(U,6)$ yields $\approx 7.3 \times 10^{66}$. Any number in that band is comfortably "beyond the heat death of the universe" for brute force.

3) Why there's no shortcut: the system removes structure on purpose

Big search spaces are only scary if you **can't do better** than blind search. In many problems, attackers find patterns, symmetries, biases, or algebraic structure that shrink the space. Our ceremony deliberately destroys those footholds:

(a) Balanced leaves. After the rings rotate and the symbols are partitioned, each leaf carries roughly equal "mass," so a blind guess at the correct leaf per round succeeds with probability about 1/6:

$$P(\text{correct leaf by guess}) \;\approx\; \frac{1}{6}.$$

"Probability of the correct leaf by guessing is about one over six."

(b) Opaque bijection (label symmetry). The six codewords Alice emits are private **names** for leaves. Because the mapping from codewords to leaves is **unknown** and could be any of the 6! permutations, every observed codeword stream is **equally compatible** with 720 different global labelings. This is classic **label-switching**: permute the labels and the data look the same.

(c) Independent modulo rotations of each alphabet ring. Before partitioning, each ring is re-indexed by its own independent offset, so any attempt to "track a symbol family by relative index" collapses immediately:

$$j' \;=\; (j + \Delta_r) \bmod N_r \quad (\text{for ring } r \text{ of size } N_r).$$

"Jay prime equals jay plus delta sub r modulo N sub r."

Because **each** ring gets **its own** fresh Δ_r **every round**, the in-ring positions of letters are constantly re-labeled—in three separate ways at once. No cross-round anchor survives.

(d) Minimal leakage per round. The verifier only needs a **membership fact** ("was that the correct leaf this round?"). From the outside, the label Y Alice emits has a marginal distribution that is essentially uniform over the six codewords:

$$P(Y = j) \approx \frac{1}{6} \quad \text{for each } j \in \{1, \ldots, 6\}.$$

"Probability Y equals j is about one over six for each j."

In information-theoretic terms, the observed label contains **no directional information** about the true next character:

$$I(C; Y) = 0 \quad \text{(in expectation under the design's randomness)}.$$

"The mutual information between C and Y equals zero."

This combination—balanced leaves, label symmetry, and re-indexed rings—makes the label stream look like **coin flips with six sides**. Without extra knowledge (a small dictionary, a known map, or side channels), there is nothing to model and no bias to exploit; **statistics and machine learning have no signal** to learn.

4) The best possible attack is elimination, and it needs ˜86 perfect rounds

Since fancy statistics don't help, the best Eve can do is **enumerate hypotheses** and **eliminate** those that contradict observed rounds. Each observed round rules out about **five-sixths** of the wrong hypotheses (because the wrong guess matches the right leaf only **one-sixth** of the time). If Eve has a transcript of R rounds (all from **one continuous ceremony**, with a fixed per-session map), the **expected** number of wrong survivors is

$$\mathbb{E}[\text{survivors after } R] \approx P(U, 6) \left(\frac{1}{6}\right)^R.$$

"Expected survivors equal P of U six times one over six to the R."

To push the expected survivors below 1 (i.e., isolate a **unique** solution), set the right-hand side to $\lesssim 1$ and solve for R:

$$R \gtrsim \log_6\big(P(U, 6)\big).$$

"R is at least log base six of P of U six."

Using $P(U, 6) \approx 72^{36}$, we get

$$R \approx \log_6(72^{36}) = 36 \log_6 72 \approx 36 \times 2.585^{-1} \times \log_2 72 \approx 36 \times 2.387 \approx 86.$$

"R is about eighty-six."

So even with **perfect** logging of boards and codewords from a **single uninterrupted session**, Eve needs on the order of **eighty-plus** rounds to isolate all six secrets and the map. Real ceremonies are much shorter, and many systems deny public access to either the full boards or the exact codewords—pushing elimination even further out of reach.

5) Big-O time and what that means in practice

Big-O hides constants and lower-order terms. Testing **one** hypothesis against R rounds takes $O(R)$ time (one membership check per round). There are $P(U, 6)$ hypotheses. So the overall **work** is

$$\text{Checks} \;=\; O\big(R \cdot P(U,6)\big) \;=\; O\big(R \cdot 72^{36}\big).$$

"Checks scale like big-O of R times seventy-two to the thirty-sixth."

What does 72^{36} look like as a **time**? Convert to bits: $72^{36} \approx 2^{222}$. Suppose an attacker can do 10^{15} hypothesis-checks per second (a wildly optimistic number—**per hypothesis** you must parse a board and apply a membership test!). Even then, the time to scan 2^{222} candidates is

$$\frac{2^{222}}{10^{15}} \text{ seconds} \;\approx\; \frac{6 \times 10^{66}}{10^{15}} \;=\; 6 \times 10^{51} \text{ seconds}.$$

"Six times ten to the fifty-one seconds."

Divide by $\sim 3.16 \times 10^7$ seconds per year:

$$\approx 1.9 \times 10^{44} \text{ years}.$$

"About one point nine times ten to the forty-four years."

For comparison, the age of the universe is $\sim 1.4 \times 10^{10}$ years. We are off by **thirty-four orders of magnitude**. Even if you parallelize across a trillion supercomputers, you only knock off 10^{12}, which doesn't move the exponents meaningfully. This is what "intractable" looks like when you do the math.

6) "But what about quantum?" (Grover's algorithm and why it still fails)

Quantum search (Grover's algorithm) can give at most a **square-root** speed-up for unstructured brute force. That would reduce 2^{222} into roughly 2^{111}. Still gigantic:

$$\sqrt{72^{36}} \;=\; 72^{18} \;=\; 2^{111}.$$

"Square root of seventy-two to the thirty-six equals seventy-two to the eighteen equals two to the one hundred eleven."

$2^{111} \approx 2.5 \times 10^{33}$. If you could perform a **billion billion** (10^{18}) quantum iterations per second (well beyond today's capabilities for structured oracles and error-corrected qubits), you'd still need

$$\frac{2.5 \times 10^{33}}{10^{18}} \;=\; 2.5 \times 10^{15} \text{ seconds} \;\approx\; 7.9 \times 10^{7} \text{ years}.$$

"About seventy-nine million years."

And this rosy estimate **ignores** the massive overheads for error correction, memory, and oracle construction. In short: **even quantum** doesn't make an unstructured search over 72^{36} feasible.

7) Why frequency analysis, correlations, and ML do not help

Frequency analysis and correlation hunting work **only** if the data stream carries a consistent bias or stable signal. Our design **removes** those:

- **Uniform marginals:** Each observed codeword label Y is, in expectation, equally likely among six choices:

  $$P(Y = j) \approx \frac{1}{6}.$$

  "Probability Y equals j is about one over six."

  As the attacker gathers more data, the counts converge to **one-sixth each**, which reinforces the lack of signal.

- **Independent ring rotations:** The indices within each alphabet family are **re-labeled independently** each round, so index-based patterns don't persist:

  $$j' = (j + \Delta_r) \bmod N_r.$$

  "Jay prime equals jay plus delta sub r modulo N sub r."

- **Label-switching symmetry:** Any codeword-to-leaf permutation among the six synonyms explains the same transcript equally well. **No statistic can pick the "true" labeling** without external anchors.

- **Zero mutual information:** The observed label stream carries, in expectation, **no information** about the true next character:

  $$I(C; Y) = 0.$$

  "I of C semicolon Y equals zero."

Because there is nothing to correlate with, even very smart models (deep nets, HMMs, transformers) will converge to **"uniform six-way dice"** as the best fit. You cannot **learn** what is not there.

8) What if the boards aren't recorded? (It gets even worse for the attacker)

So far we have been generous to the attacker by allowing them to **see** the board each round—so they don't also have to guess how the rings rotated. If, instead, the boards are **not** observable (e.g., a secure attention window blocks screen capture), then per round the attacker must also guess the triple of ring offsets $(\Delta_{\text{lc}}, \Delta_{\text{uc}}, \Delta_{\text{ds}})$. The number of such triplets is $30 \times 30 \times 12 = 10{,}800$. Over R rounds, that multiplies the hypothesis space by

$$(10{,}800)^R.$$

"Ten thousand eight hundred to the R."

And if they must guess the **balanced partition** itself (which 12 of the 72 symbols landed in each leaf), the count explodes by a large multinomial factor,

$$\frac{72!}{(12!)^6\,6!},$$

"Seventy-two factorial over twelve factorial to the sixth times six factorial."

which dwarfs everything. In other words, **hiding the boards** does not just maintain intractability—it makes it **much, much worse**.

9) "Edge" assumptions and why they still don't save the attacker

Let's examine a few "what-ifs" attackers often ask about, and see why the intractability remains.

What if the six secrets were not necessarily distinct?
Then the ordered count becomes $U^6 = 72^{36}$ exactly. Asymptotically, that's **the same Big-O** as before. No relief.

What if the attacker already knew a small dictionary of candidate secrets?
If the dictionary had size D (rather than U), the joint search becomes $P(D, 6)$ instead of $P(U, 6)$. If D is truly small (say, human-memorable phrases), this matters—but that contradicts our stipulation of **no prior knowledge**. The system's security budget comes from secrets drawn from a massive, uniform universe; if users pick from a tiny list and the attacker knows it, the math changes. That's a **policy** and **hygiene** issue, not a weakness of the projection-and-bijection design.

What if you attack one secret at a time?
The "one-secret" space is $U = 72^6$, roughly 2^{37} more modest—but you still have to multiply across six secrets to get the **whole** set. The earlier **rounds-to-uniqueness** derivation shows this clearly: you need about $\log_6(U) = 6 \log_6 72 \approx 14.3$ rounds to pin **one** unknown secret (+ map) and roughly $6 \times 14.3 \approx 86$ to pin **all six**—the same number we derived from the full $P(U, 6)$ analysis. Slicing the elephant does not make it smaller.

What if the private map leaks a little?
A leak that halves the map uncertainty (from 6! to 360) hardly changes the total since 6! is microscopic compared to U^6. The map contributes only a small constant factor to the full search, not the exponential bulk.

10) A physical reality check: time, energy, and memory

Even if you ignore software and just think physics, the search is dead on arrival.

- **Time.** We already computed the wall-clock under absurdly optimistic 10^{15} checks/second: $\sim 10^{44}$ years. That's not a "tricky algorithm away" from feasible.
- **Energy.** At room temperature, the Landauer bound for erasing one bit is about $kT \ln 2 \approx 3 \times 10^{-21}$ joules. Even if each hypothesis check "cost" a

single bit erasure (an impossible fantasy), checking 2^{222} hypotheses would take on the order of 10^{46} joules—absurd on planetary scales.

- **Memory/I/O.** Holding or streaming candidate tuples and boards at these magnitudes saturates any realistic I/O bus. You cannot even **name** the candidates fast enough, let alone test them.

Intractability here is not stylistic; it's **physical**.

11) The role of ceremony design: keeping the problem "black-box"

Attackers win when they can translate a problem into a more structured one: linear algebra over finite fields, lattice reduction, SAT to 2-SAT, etc. Our ceremony defends by being stubbornly **black-box**:

- The only thing you can do with a hypothesis is **test** it against **this round's** board and label—there's no gradient, no algebra, no helpful symmetry except the one (label-switching) that blocks you.
- **New rounds do not correlate**: each uses fresh ring rotations and a balanced repartition, so the signal does not "add up" across time except in the trivial "eliminate the wrong ones" sense.
- Across **different ceremonies**, even the private bijection does not give you a foothold: unless you sit through a **single uninterrupted session**, you can't transport constraints forward.

Systems that preserve this "black-box" nature of membership checks remain resistant to clever optimizations. The fastest way is still the dumb way—and the dumb way is functionally impossible.

12) Final Word

- **Exact global search space** for recovering **all six** 6-char secrets (distinct) **and** the private map, with **no prior knowledge**:

$$\boxed{P(U,6) \;=\; U(U-1)(U-2)(U-3)(U-4)(U-5)}, \quad U = 72^6.$$

"P of U six equals U times U minus one down to U minus five, where U equals seventy-two to the sixth."

Magnitude: $P(U,6) \approx 7.3 \times 10^{66} \approx 2^{222}$.

- **Big-O time** to brute-force:

$$\boxed{\Theta\big(72^{36}\big)}.$$

"Theta of seventy-two to the thirty-sixth."

- **Expected wrong-survivor count after R rounds**:

$$\boxed{\mathbb{E}[\text{survivors}] \;\approx\; P(U,6)\,\big(1/6\big)^{R}}.$$

"Expected survivors equal P of U six times one over six to the R."

- **Rounds-to-uniqueness (information lower bound)**:

$$\boxed{R \;\gtrsim\; \log_6\big(P(U,6)\big) \;\approx\; 86}.$$

"R is about eighty-six."

- **Across ceremonies:** frequency and correlation carry **no signal**,

$$P(Y=j) \approx \tfrac{1}{6}, \qquad I(C;Y) = 0,$$

"Probability Y equals j is about one over six; mutual information between C and Y equals zero."

- **If boards are hidden:** rotations alone multiply the space by $(10{,}800)^R$,

$$(30 \cdot 30 \cdot 12)^R \;=\; 10{,}800^R.$$

"Ten thousand eight hundred to the R."

These are the bones of the argument: an enormous, structureless space and a per-round "bit budget" that is too small to chip it down within any sane number of observed rounds.

13) One last mental picture (for memory)

Think of the attacker's job as trying to identify a **six-book anthology** (the six secrets), chosen from a library with 72^6 different titles, while the **shelf labels** (the private bijection) are secretly permuted in 720 different ways, and the **shelf order** (the rings' indices) is re-shuffled independently before each glance. Each time the attacker looks up, the librarian has rotated every shelf, re-balanced the books evenly across aisles, and swapped the shelf-name placards behind mirrored glass. The attacker can only ask one question per glance—"is the anthology's next book on **this** aisle right now?"—and gets only a yes/no. That is not a solvable treasure hunt; it's **designed** to be a needle-in-a-needle-factory problem.

14) Bottom line

It's computationally intractable to recover **all six** of Alice's 6-character secrets **and** her private map under the stated assumptions because:

1. The exact hypothesis space is $P(U,6) = U(U-1)\cdots(U-5)$ with $U = 72^6$, i.e., about 2^{222} candidates—astronomical.

2. The ceremony exports **no usable structure**: balanced leaves keep per-round chance at 1/6; independent ring rotations obliterate cross-round anchors; the private bijection induces 720-way label symmetry; the observed labels have uniform marginals; and $I(C;Y) = 0$.

3. The best attack—elimination—requires around **86 perfect, continuous rounds** from a single session just to isolate the unique global solution; real systems don't grant that view.

4. Even quantum square-root speedups leave the cost at 2^{111}, still completely infeasible.

5. Hiding boards (or partitions) only multiplies the search by gigantic per-round factors, making a bad situation worse for the attacker.

Intractability here is not a buzzword; it is a **product of exact counts, information limits, and physics**. The design ensures that, for an attacker with **no prior knowledge**, **there is nothing to learn** from watching many ceremonies and **nowhere to go** but a brute-force search across a space that might as well be infinite on human timescales.

A Formal Complexity-Theoretic Analysis of Membership-Only Transcripts with Private Leaf Bijections and Independently Rotating Alphabet Rings

Audience: University/Institute faculty; present a mathematically auditable argument that recovering the full secret-set and private bijection from passive observation is computationally intractable under standard black-box/Oracle and information-theoretic models. We supply exact combinatorics, asymptotics via Stirling, decision/communication lower bounds, and quantum query bounds.

0. Model and Objects

Let the finite symbol alphabet be a disjoint union of three "rings"

$$\Sigma \;=\; \Sigma_1 \,\dot\cup\, \Sigma_2 \,\dot\cup\, \Sigma_3, \qquad \big(|\Sigma_1|, |\Sigma_2|, |\Sigma_3|\big) = (30, 30, 12), \quad |\Sigma| = 72.$$

- "Sigma equals Sigma one disjoint union Sigma two disjoint union Sigma three, sizes thirty, thirty, and twelve; total seventy-two."*

A **secret** is a word of length $\ell = 6$ over Σ; hence the universe of secrets is

$$U \;:=\; |\Sigma|^{\ell} \;=\; 72^{6}.$$

- "U equals seventy-two to the sixth."*

Alice holds a **set of six distinct secrets** $S = \{s^{(1)}, \ldots, s^{(6)}\} \subset \Sigma^\ell$, $|S| = 6$. She also fixes a **private bijection** (the synonym map) $\varphi : \mathcal{C} \to \mathcal{L}$ between a set of six codewords $\mathcal{C}$ and the six leaves $\mathcal{L} = \{1, \ldots, 6\}$. We write S_6 for the symmetric group on six elements; φ is an element of S_6 once we identify $\mathcal{C} \cong \mathcal{L}$.

A **round** i of an interactive ceremony is parameterized by public randomness

$$\mathsf{B}_i = \big(\Delta_i^{(1)}, \Delta_i^{(2)}, \Delta_i^{(3)}, \Pi_i\big),$$

where $\Delta_i^{(r)}$ is an independent modulo rotation on ring r and Π_i is a balanced partition mapping Σ to leaves $\mathcal{L}$ so that each leaf receives $|\Sigma|/6 = 12$ symbols. The next character to be proved is $C_i \in \Sigma$ (the i-th character across the six secrets under some public schedule). Its **true leaf** is $L_i := \Pi_i(C_i) \in \mathcal{L}$. Alice's **observable witness** is the codeword

$$Y_i \;:=\; \varphi^{-1}(L_i) \in \mathcal{C}.$$

The observer (attacker) sees Y_i and either (a) the rendered board B_i ("board-visible model"), or (b) not even B_i ("board-hidden model"). Throughout, rings are **independently** rotated each round; partitions are **fresh and balanced** each round.

The attacker's target is the **global hypothesis**

$$H \;:=\; (S, \varphi),$$

i.e., all six secrets and the private bijection.

1. Exact Hypothesis Count (All Six Secrets and the Map)

Let $U = |\Sigma|^\ell = 72^6$. The number of 6-element subsets of U (unordered, distinct) is $\binom{U}{6}$; the number of bijections is 6!. Hence

$$|\mathcal{H}| \;=\; 6!\binom{U}{6} \;=\; U(U-1)(U-2)(U-3)(U-4)(U-5) \;=\; P(U,6).$$

- "Six factorial times U choose six equals U times U minus one down to U minus five, which is P of U comma six."*

  Lemma 1 (Counting identity). The global search space for (S, φ) equals the number of **ordered** 6-tuples of distinct secrets, i.e., $P(U, 6)$.
  Proof. Elementary bijection: choosing an unordered 6-set of secrets times ordering by the private map labels (there are 6! orders) equals choosing an ordered 6-tuple of distinct secrets. □

With $U = 72^6$, Stirling gives $P(U, 6) = \Theta(U^6) = \Theta(72^{36})$. Numerically,

$$U = 72^6 = 139{,}314{,}069{,}504, \qquad P(U,6) \approx U^6 = 72^{36} \approx 2^{222} \approx 7.3 \times 10^{66}.$$

- "U equals one hundred thirty-nine billion ...; P of U six is about two to the two-hundred twenty-two, about seven point three times ten to the sixty-six."*

2. Distributional Symmetries and Information Null Results

2.1 Balanced leaves and independence

By construction, Π_i is a balanced partition; for any fixed symbol $c \in \Sigma$,

$$\Pr(L_i = \ell \mid C_i = c, \mathsf{B}_i) \;=\; \tfrac{1}{6} \quad \forall \ell \in \mathcal{L}.$$

- "Probability L sub i equals ell given C sub i equals c and the board equals one over six."*

Because each ring's index is independently rotated ($j' \equiv j + \Delta_i^{(r)} \bmod |\Sigma_r|$), there is no cross-round positional anchor even within a ring.

2.2 Label-switching (private bijection)

Fix a prior on φ as uniform over S_6. Then marginally, conditional on B_i,

$$\Pr(Y_i = y \mid C_i, \mathsf{B}_i) \;=\; \tfrac{1}{6} \quad \forall y \in \mathcal{C}.$$

- "Probability Y sub i equals y given the character and the board is one over six."*

Lemma 2 (Mutual information zero, single round). Under the uniform prior on φ, $I(C_i; Y_i \mid \mathsf{B}_i) = 0$.
Proof. For any c, L_i is uniform on $\mathcal{L}$ due to balance; φ^{-1} is a uniformly random permutation, hence $Y_i = \varphi^{-1}(L_i)$ is uniform on $\mathcal{C}$. Therefore the conditional distribution of Y_i does not depend on C_i given B_i. □

Across *independent* ceremonies, φ is unobserved and effectively re-randomized (or at least remains opaque); frequencies over Y converge to uniform, blocking correlation/frequency inference.

Remark. Within a *single* session where φ is fixed but unknown, joint observations (B_i, Y_i) constrain (S, φ) through **consistency**; this supports hypothesis elimination but does **not** create a learnable bias (cf. §4).

3. Decision-Tree and Query-Complexity Lower Bounds

We treat recovery of $H = (S, \varphi)$ as an unstructured search in a hypothesis set $\mathcal{H}$ of size $M := |\mathcal{H}| = P(U, 6)$. Each candidate $h \in \mathcal{H}$ can be **tested** against a transcript of R rounds by checking per-round consistency with (B_i, Y_i); each test costs $O(R)$ (a constant-time membership per round). There is **no algebraic shortcut**: the only operations available are evaluation of these consistency predicates—equivalently, we are in the black-box/Oracle model.

3.1 Classical lower bound

In the decision-tree model, any deterministic algorithm that always finds the unique $h^\star \in \mathcal{H}$ must, in the worst case, perform at least $M-1$ inequality tests. With randomization and two-sided error $\le \epsilon < 1/2$, Yao's minimax principle implies an expected $\Omega(M)$ tests on some input distribution.

> **Theorem 1 (Classical black-box).** Any classical algorithm that recovers H with constant success probability using only hypothesis-consistency queries requires $\Omega(M)$ queries; with per-query cost $O(R)$, the time is $\Omega(R \cdot M) = \Omega\big(R \cdot 72^{36}\big)$.

3.2 Quantum lower bound

By the BBBV/Grover lower bound, unstructured search over M items requires $\Omega(\sqrt{M})$ quantum queries to an Oracle that recognizes the marked item.

> **Theorem 2 (Quantum black-box).** Any quantum algorithm that identifies H with constant success probability using a membership-consistency Oracle needs $\Omega(\sqrt{M}) = \Omega\big(72^{18}\big) = \Omega\big(2^{111}\big)$ queries; with per-query cost $O(R)$, time is $\Omega\big(R \cdot 72^{18}\big)$.

Both bounds certify that *even in idealized Oracle models* the search is infeasible at our magnitudes.

4. Sample-Complexity (Rounds-to-Uniqueness) Lower Bounds

Let the attacker receive R i.i.d. rounds $\{(\mathsf{B}_i, Y_i)\}_{i=1}^{R}$ from the true H. Define the hypothesis class $\mathcal{H}$ with uniform prior. Denote the observation channel as $P_{(\mathsf{B},Y)|H}$.

A standard *multiple-hypothesis testing* lower bound (via Fano) gives

$$P_e \;\ge\; 1 - \frac{I(H; \mathsf{B}^R, Y^R) + \log 2}{\log |\mathcal{H}|}\,,$$

- "Error probability is at least one minus information over log of the hypothesis count, up to a log two term."*
  where $I(\cdot;\cdot)$ is mutual information and log is base e or 2 consistently.

Per-round information cap. A single round reveals at most $\log 6$ nats (or $\log_2 6$ bits): the label Y_i ranges over six equiprobable values once B_i is fixed. Hence

$$I\big(H; (\mathsf{B}_i, Y_i)\big) \;\le\; \log 6 \quad \text{(nats)} \qquad \text{or} \qquad \le \log_2 6 \quad \text{(bits)}.$$

- "Information per round is at most log six; in bits, log base two of six."*

Aggregating over R independent rounds,

$$I\big(H; \mathsf{B}^R, Y^R\big) \;\le\; R \log 6 \quad \text{(nats)}.$$

- "Total information is at most R times log six."*

Take logarithms base 2 for concreteness. With $|\mathcal{H}| = P(U, 6)$, and aiming at $P_e \leq 1/3$, Fano yields the necessary condition

$$R \gtrsim \frac{\log_2 |\mathcal{H}| - 1}{\log_2 6}.$$

- "R is at least log base two of the hypothesis size minus one, divided by log base two of six."*

Now

$$\log_2 |\mathcal{H}| = \log_2 P(U, 6) = \sum_{k=0}^{5} \log_2(U - k) = 6 \log_2 U + o(1),$$

- "Log of P of U six is six times log U plus lower order terms."*
and $\log_2 U = \log_2(72^6) = 6 \log_2 72 \approx 6 \times 6.17 = 37.0$. Hence

$$\log_2 |\mathcal{H}| \approx 6 \times 37.0 \approx 222.1 \text{ bits}, \quad \log_2 6 \approx 2.585,$$

giving

$$R \gtrsim \frac{222.1}{2.585} \approx 85.9 \approx 86 \text{ rounds}.$$

- "R is about eighty-six rounds."*

Theorem 3 (Information lower bound). Any passive observer that only sees (B_i, Y_i) must acquire on the order of $R \approx 86$ i.i.d. rounds, from a single session with fixed φ, to reduce the Bayes error below a constant and uniquely identify (S, φ). This bound is tight up to constants and matches the heuristic elimination rate $(1/6)^R$.

Heuristic agreement. Wrong hypotheses survive a round with probability $\approx 1/6$, so after R rounds the expected survivors are $|\mathcal{H}|(1/6)^R$; "uniqueness in expectation" requires $(1/6)^R \lesssim 1/|\mathcal{H}| \Rightarrow R \gtrsim \log_6 |\mathcal{H}|$, exactly the above.

5. Time Complexity (Work Factor) with Observed Boards

Per hypothesis h, testing against R rounds is $O(R)$. Thus the **time to brute-force** is

$$T_{\text{classical}}(R) = \Theta\big(R \cdot |\mathcal{H}|\big) = \Theta\big(R \cdot 72^{36}\big),$$

- "Time is Theta of R times seventy-two to the thirty-sixth."*
and the **quantum** analogue is

$$T_{\text{quantum}}(R) \;=\; \Omega\big(R \cdot \sqrt{|\mathcal{H}|}\big) \;=\; \Omega\big(R \cdot 72^{18}\big).$$

- "Quantum time is Omega of R times seventy-two to the eighteenth."*

As orders of magnitude: $72^{36} \approx 2^{222} \approx 7.3 \cdot 10^{66}$. Even at 10^{15} checks/s, classical time is $\sim 10^{44}$ years ("about ten to the forty-four years"). Quantum square-root still leaves $\sim 2^{111}$ steps ("two to the one hundred eleven"), on the order of 10^8 years at impossible-throughput oracles.

6. Board-Hidden Model: Additional Combinatorics per Round

If B_i is **not** observable, the attacker must also hypothesize per-round ring rotations and (if not derivable) balanced partitions.

- **Ring rotations.** Per round there are $30 \cdot 30 \cdot 12 = 10{,}800$ offset triples $(\Delta_i^{(1)}, \Delta_i^{(2)}, \Delta_i^{(3)})$. Over R rounds this multiplies the space by

$$(10{,}800)^R.$$

- "Ten thousand eight hundred to the R."*

- **Balanced partitions.** The number of balanced allocations of 72 labeled symbols into six bins of 12 is

$$\frac{72!}{(12!)^6\, 6!}\,,$$

- "Seventy-two factorial over twelve factorial to the sixth times six factorial."*
  per round. Including this factor (if needed) renders the search superastronomical.

Thus, hiding boards cannot reduce intractability; it only exacerbates it.

7. Why Statistics/ML Cannot Help (Identifiability and Exchangeability)

Let $\mathcal{D} = \{Y_i\}$ or $\{(\mathsf{B}_i, Y_i)\}$ across many *independent ceremonies* (with φ effectively randomized). Then for each round i,

$$\Pr(Y_i = j \mid \mathsf{B}_i) \;=\; \tfrac{1}{6} \quad \forall j,$$

- "Probability Y equals j given the board is one over six for all j."*
  and the joint law is **exchangeable** under S_6 relabelings. Hence:

- Frequentist **frequencies** converge to uniform 1/6.

- Any **correlation** between labels and features of B_i vanishes in expectation (independent ring rotations destroy persistent features).
- The model is **non-identifiable** under label-switching: every parameter point (S, φ) has 720 observationally equivalent relabelings in the passive, cross-ceremony regime.

In the *single-session* regime, the only effective method is **elimination** by consistency across many rounds; cf. §4.

8. Robustness to Side Assumptions

- **Secrets with repetition.** If repetitions were allowed, ordered count becomes $U^6 = 72^{36}$; asymptotics unchanged.
- **Attacker knows a dictionary of size $D< U$.** Replace U by D; the space becomes $P(D, 6)$. Security then depends on dictionary entropy; this is a **policy** constraint, not a structural weakness.
  * **Partial leakage of φ.** Reducing the map space by a constant factor (e.g., from 6! to 360) does not dent the U^6 bulk.

9. Conclusions (Validated Claims)

1. **Exact global search space.**

$$|\mathcal{H}| \;=\; 6!\binom{U}{6} \;=\; P(U,6) \;=\; U(U-1)\cdots(U-5), \quad U = 72^6.$$

- "Hypothesis size equals six factorial times U choose six equals P of U six equals U times U minus one … minus five."*

2. **Asymptotics and magnitude.**

$$|\mathcal{H}| = \Theta(72^{36}) \approx 2^{222} \approx 7.3 \times 10^{66}.$$

- "About two to the two-hundred twenty-two, seven point three times ten to the sixty-six."*

3. **Classical/quantum work factors.**

$$T_{\text{classical}} = \Theta\big(R \cdot 72^{36}\big), \qquad T_{\text{quantum}} = \Omega\big(R \cdot 72^{18}\big).$$

- "Classical Theta R times seventy-two to the thirty-six; quantum Omega R times seventy-two to the eighteenth."*

4. **Rounds-to-uniqueness lower bound (Fano/heuristic agreement).**

$$R \gtrsim \frac{\log_2 |\mathcal{H}|}{\log_2 6} \approx \frac{222.1}{2.585} \approx 86.$$

- "R is about eighty-six."*

5. **Information null results across ceremonies.**

$$P(Y = j \mid \mathsf{B}) = \tfrac{1}{6} \text{ and } I(C; Y \mid \mathsf{B}) = 0,$$

- "Y is uniform one-sixth; mutual information between C and Y given the board equals zero."*
  by balanced partitions, independent ring rotations, and label-switching.

6. **Board-hidden multiplicative blowup.**

extra factor $(30{\cdot}30{\cdot}12)^R = 10{,}800^R$ (rotations), and possibly $\dfrac{72!}{(12!)^6\, 6!}$ per round (partitions).

- "Ten thousand eight hundred to the R; and seventy-two factorial over twelve factorial to the sixth times six factorial."*

All three pillars—**exact combinatorics**, **black-box query lower bounds (classical/quantum)**, and **information-theoretic sample lower bounds**—concur: the passive recovery of the entire six-secret set and the private bijection is computationally intractable under the stated ceremony, even granting the attacker perfect logging of boards and codewords within a session. The ceremony's design (balanced partitions, independent ring rotations, private leaf bijection, minimal witness) purposefully removes exploitable structure, reducing the adversarial task to unstructured search/elimination in a hypothesis space of size $\approx 2^{222}$, with a per-observation information budget of $\log_2 6$ bits—hence the ~ 86-round information lower bound and astronomically large time/work factors.

Quantum Can't Crack It: Why Six Secrets and a Private Map Stay Safe—even Against a Fleet of Quantum Computers

Abstract (two voices in one)

For first-year undergrads (lay explanation): Think of Alice's proof as a game of "find the right region" on a board with six zones. Each round, letters from three alphabets are re-shuffled across those zones, and each alphabet ring is independently rotated so positions change unpredictably. Alice never reveals

a letter—only a private codeword that means "the zone I saw my letter in." Because her six codewords are secretly paired with the six zones by a private one-to-one mapping, anyone watching sees labels but can't tell what they mean. Over many rounds, this looks like a fair six-sided die roll—no patterns to learn, no frequencies to exploit. Even a quantum computer can only try guesses much faster, but the search space is so astronomically big that "faster" still means "far beyond the life of the universe."

For CS professors: We formalize the attack as black-box identification of a marked element in a hypothesis set of size $M = 6!\binom{U}{6} = P(U, 6)$ where $U = |\Sigma|^{\ell} = 72^6$. Per-round observables are membership bits under balanced partitions with independent ring rotations, and public boards (if visible) provide no learnable bias: $I(C; Y \mid B) = 0$. Classically, any algorithm requires $\Omega(M)$ membership-consistency tests; quantumly, BBBV/Grover yields $\Omega(\sqrt{M})$ queries to any such oracle, with per-query cost $\Theta(R)$ rounds. For complete recovery of all six secrets and the bijection with non-negligible success, information-theoretic lower bounds show $R \gtrsim \log_6 M \approx 86$ perfectly observed rounds from a single session. Numerically, $\sqrt{M} \approx 72^{18} \approx 2^{111}$ is intractable even with massive parallel quantum hardware.

1) The Game Board (shared intuition)

Imagine a canvas split into **six leaves/zones**. The system uses **three alphabet rings** (lowercase 30, uppercase 30, digits 12), for a total of 72 symbols. Each round:

1. **Independent ring rotations:** every ring is re-indexed by its **own** random offset;
2. **Balanced partition:** the 72 symbols are evenly distributed, 12 per leaf;
3. **Private synonym map:** Alice owns a secret one-to-one pairing between six codewords and the six leaves.

Alice's goal is to prove, round by round, that she knows the next character in a hidden 6-character secret—**without ever saying the character**. She looks, sees which leaf holds her next character on that round, and returns only **the private codeword** for that leaf. Bob (the verifier) can check this because he knows the temporary layout and the checking mask for the round. An observer (even a quantum-empowered one) sees just a codeword per round, under a mapping they do not know, against a layout that is re-randomized each round and each session.

2) Counting the Thing to Break (both audiences)

Each 6-character secret is a string over 72 symbols, so the universe of possible secrets is

$$U \;=\; 72^6.$$

"U equals seventy-two to the sixth."

Alice has **six** distinct secrets. If you don't care about their order, the number of possible 6-element secret sets is $\binom{U}{6}$. She also has a **private bijection** (a permutation) between six codewords and the six leaves, so there are

$$6! \;=\; 720$$

"Six factorial equals seven hundred twenty."
possible maps.

Conveniently, the exact size of the **joint** search space (all six secrets **and** the private map) is

$$6!\binom{U}{6} \;=\; U(U-1)(U-2)(U-3)(U-4)(U-5) \;=\; P(U,6).$$

"Six factorial times U choose six equals U times U minus one down to U minus five, which is P of U comma six."

With $U = 72^6$, this is essentially $U^6 = 72^{36}$ (because $U \gg 5$):

$$P(U,6) \;\approx\; 72^{36} \;=\; 2^{36\log_2 72} \;\approx\; 2^{222} \;\approx\; 7.3\times 10^{66}.$$

"P of U six is about seventy-two to the thirty-sixth, equal to two to the two-hundred twenty-two, about seven point three times ten to the sixty-six."

That is the single, precise number to keep in mind: **about** 2^{222} **possibilities** for "six secrets plus the map."

3) Why Statistics Don't Help (uniformity and symmetry)

Every round, the layout is **fresh** and **balanced** across leaves; each ring is **independently rotated**; and Alice's codeword is drawn from an **unknown permutation** of leaf labels. For any true next character C, the chance it lands in any particular leaf ℓ is about

$$\Pr(L = \ell \mid C, B) \;\approx\; \frac{1}{6}.$$

"The probability the character lands in any leaf is about one over six."

Because the codeword is simply the **private name** for the chosen leaf, and that mapping is a **hidden permutation**, the codeword the observer sees is, in expectation, **uniform** among the six possibilities:

$$\Pr(Y = j \mid B) \;\approx\; \frac{1}{6} \quad (j = 1,\ldots,6).$$

"The probability the observed label equals j is about one over six."

Information-theoretically, the **mutual information** between the next character and the observed label, conditioned on the board, is

$$I(C; Y \mid B) \;=\; 0 \quad \text{(in expectation)}.$$

"The mutual information of C and Y given the board equals zero."

For undergrads: **no pattern forms**; counting labels just returns "one-sixth each." For professors: this is an **exchangeable** channel with S_6 label-switching symmetry and independently re-indexed rings; cross-ceremony identification is **non-identifiable** without external anchors.

4) The Classical Baseline: Why It's Already Impossible

If an attacker had perfect recordings of R consecutive rounds **from one session** (so the per-session map stays fixed), the only viable method is **elimination**: test a candidate hypothesis $H = (S, \varphi)$ against each round and discard it on the first mismatch. A wrong hypothesis "accidentally" agrees with a round only if it picks the right leaf by chance, which is about $1/6$. Thus the expected fraction of wrong survivors after R rounds is $(1/6)^R$. Starting with $P(U, 6) \approx 2^{222}$ candidates, the expected number of survivors is

$$\mathbb{E}[\text{survivors}] \;\approx\; P(U, 6) \left(\frac{1}{6}\right)^R.$$

"Expected survivors equal P of U six times one over six to the R."

To reduce that expected count below 1 (i.e., isolate a unique global answer), you need roughly

$$R \;\gtrsim\; \log_6\big(P(U, 6)\big).$$

"R is at least log base six of P of U six."

Plugging the numbers:

$$\log_6\big(72^{36}\big) \;=\; 36 \log_6 72 \;\approx\; 36 \times 2.387 \;\approx\; 86.$$

"R is about eighty-six."

So even in the **best classical visibility** (a long, uninterrupted session with perfect board captures and exact codewords), you need **about 86 rounds** to narrow down to a unique set of six secrets **and** the private map. Real ceremonies are much shorter. Across multiple ceremonies the private labels re-scramble, the rings re-index, and elimination **does not accumulate**.

5) Enter Quantum: What Advantage Is Even Possible?

There are two widely cited quantum speedups:

- **Shor-style (exponential) speedups** for structured number-theoretic problems (period finding in abelian groups). **Not applicable** here: there is no hidden period or group structure to exploit; every round is a fresh, balanced, randomized projection with private label permutations.
- **Grover-style (quadratic) speedups** for **unstructured search**. This is the relevant model: given a black-box **oracle** that says whether a hypothesis is the right one, a quantum computer can find the marked item in about $\frac{\pi}{4}\sqrt{N}$ oracle calls instead of $N/2$ classically.

Our problem is exactly an **unstructured search** over

$$M \;=\; |\mathcal{H}| \;=\; P(U,6) \;\approx\; 72^{36} \;\approx\; 2^{222}$$

"M equals P of U six, about seventy-two to the thirty-sixth, about two to the two-hundred twenty-two."
hypotheses, where each oracle call asks: "Does hypothesis $H = (S, \varphi)$ match **all** R observed rounds?" That oracle itself must evaluate **R** membership checks (one per round), so each quantum query costs $\Theta(R)$ operations.

Grover's lower bound (BBBV/Zalka) says you still need on the order of $\sqrt{M}$ such queries to succeed with constant probability. Here, that is

$$\sqrt{M} \;\approx\; \sqrt{72^{36}} \;=\; 72^{18} \;=\; 2^{111}.$$

"Square root of M equals seventy-two to the eighteenth, equal to two to the one-hundred eleven."

Two massive observations follow:

1. A **quadratic** speedup over "impossible" is **still impossible**. 2^{111} is astronomically large. Even if a futuristic quantum machine could execute 10^{18} (a quintillion) oracle calls per second—well beyond reasonable projections—the runtime would be

$$\frac{2^{111}}{10^{18}} \;\approx\; 2.5 \times 10^{15} \text{ seconds} \;\approx\; 7.9 \times 10^{7} \text{ years}.$$

"About seventy-nine million years."

2. An oracle call costs $\Theta(R)$, and to **uniquely** identify all six secrets and the map you need $R \approx 86$ rounds (information-theoretic lower bound). Multiplying by R only worsens the estimate.

So in the **best conceivable** quantum framing (idealized oracle, error-free qubits, no I/O overhead), the asymptotic lower bound already rules out practicality by breathtaking margins.

6) "What If I Don't Have 86 Rounds?" (many marked items and amplitude amplification)

Another way to see the quantum dead-end is via **multiple marked items**. If you have only R observed rounds, many hypotheses survive: roughly

$$k \;\approx\; M\left(\frac{1}{6}\right)^{R}.$$

"k equals M times one over six to the R."

Amplitude amplification then finds **some** surviving hypothesis in $O\big(\sqrt{M/k}\big)$ queries, i.e.,

$$O\Big(\sqrt{\tfrac{M}{M/6^R}}\Big) \;=\; O\big(6^{R/2}\big).$$

"On the order of six to the R over two."

This looks smaller than $\sqrt{M}$, but remember **your goal is unique identification** of all six secrets and the map. With small R, you only find *one among many* consistent hypotheses; you still haven't cracked the real secrets. To reduce to a **single** marked item you must crank R up until $k \approx 1$, which is precisely $R \gtrsim \log_6 M \approx 86$. At that point, we are back to the $\sqrt{M} \approx 2^{111}$ query complexity.

7) Why Shor-like, QAOA, and HHL-style Approaches Don't Apply

- **Shor (period finding):** requires a function with a hidden period over an (abelian) group. Our per-round projection is **fresh**, the rings are **independently rotated**, and the private mapping induces **label symmetry**, not a stable period. There is **no underlying group structure** to exploit.
- **QAOA/annealing on a cost graph:** you could encode "consistency with R rounds" as a giant SAT/Ising instance, but the landscape is essentially random with exponentially many nearly orthogonal hypotheses. Without structure (e.g., locality, low-treewidth), QAOA confers no proven advantage over black-box bounds.
- **HHL/linear systems:** there is no linear system with sparse, well-conditioned matrices lurking here; the target is a discrete combinatorial identification with an oracle, not a linear algebraic solve.

The intractability comes from **lack of exploitable structure** plus an **astronomical hypothesis count**; the usual quantum hammers have nothing coherent to hit.

8) Parallel Quantum Fleets: Why Billions of Qubits Still Don't Save You

Grover-type search parallelizes **sublinearly**. If you split the search across p independent quantum processors, the per-machine work becomes $\sqrt{M/p}$, so total time scales as $\sqrt{M}/\sqrt{p}$. With $p = 10^{12}$ (a trillion machines—science fiction), the speedup is only a factor of 10^6. Using the $\sqrt{M} \approx 2.5 \times 10^{33}$ oracle-calls estimate:

$$\frac{2.5 \times 10^{33}}{\sqrt{10^{12}}} = \frac{2.5 \times 10^{33}}{10^6} = 2.5 \times 10^{27} \text{ oracle calls.}$$

"Two point five times ten to the twenty-seven oracle calls."

Even at a (fantastical) rate of 10^{12} calls per second per machine with perfect error correction, this is 2.5×10^{15} seconds—**still tens of millions of years**, and we have ignored the cost of **quantum error correction** (which inflates gate counts by many orders of magnitude) and **oracle depth** ($\tilde{}R$ consistency checks per query).

9) Physics and Error Correction: The Real-World Tax

A brief, professor-level reality check:

- **Surface code overhead:** To achieve logical error rates low enough for $\sim 10^{33}$ queries, you need **astronomically many** physical qubits per logical qubit (often 10^3–$10^6\times$), plus repeated syndrome extraction. A "billion-qubit" device is tiny on this scale.
- **Gate depth:** Each oracle query computes R per-round membership checks **reversibly**, with nontrivial Toffoli/T-depth. Multiplying by 2^{111} queries blows past any reasonable coherence budget—even before error correction.
- **Power and cooling:** Large-scale cryogenic quantum fleets have strict thermal budgets; simply *clocking* the required gates fast enough runs into thermodynamic and engineering walls long before algorithmic lower bounds are met.

The black-box lower bound tells you **where the wall is**. Physics tells you that the **doorway to the wall** is itself unreachable.

10) The Probability View (quick sanity checks)

Two practical probabilities make the intuition stick:

1. **Random guessing across rounds:** the chance that a non-member passes R independent rounds is

$$\Pr[\text{pass}] \;=\; \left(\frac{1}{6}\right)^{R}.$$

"Probability of passing equals one over six to the R."

For $R = 20$, this is $\approx 1/3.6 \times 10^{15}$. For $R = 40$, $\approx 1/1.3 \times 10^{31}$. There is no "learnable" drift to beat these odds because the labels are symmetric and the rings are re-indexed.

2. **Information per round:** the maximum reduction in uncertainty per round is at most $\log_2 6 \approx 2.585$ bits:

$$\Delta\mathcal{H}_{\max} \;=\; \log_2 6 \;\approx\; 2.585 \text{ bits/round.}$$

"Delta H max equals log base two of six, about two point five eight five bits per round."

To discharge the ≈ 222 bits of global uncertainty in $P(U, 6)$, you need ≈ 86 rounds. Quantum measurement does not let you extract **more** than that from a single classical round; your oracle can be queried coherently, but the BBBV bound caps the benefit at a square root over **the entire** hypothesis set.

11) A Side Branch: If the Boards Are Hidden, It Gets Harder (not easier)

We've been generous to the attacker by assuming they can record boards (layouts) each round. If, instead, the system renders boards in a secure attention window (blocking capture), the attacker must *also* hypothesize the ring rotations (and possibly the balanced partitions):

- **Ring-offset triplets per round:** $30 \times 30 \times 12 = 10{,}800$, so over R rounds the search is multiplied by

$$(10{,}800)^{R}.]$$

"Ten thousand eight hundred to the R."

- **Balanced partitions:** if unknown, the number of allocations of 72 labeled symbols into six bins of 12 is

$$\frac{72!}{(12!)^6\, 6!} \quad \text{per round.}$$

"Seventy-two factorial over twelve factorial to the sixth times six factorial."

These factors are multiplicative **on top of** the $P(U, 6)$ secrets+map space. Quantum can't "eat" this extra exponential either; there is still no structure—only an even bigger unstructured search.

12) One Page for Each Audience

12A) Undergrad one-pager (plain words + tiny math)

- **How big is the haystack?** All ways Alice could have six different 6-character secrets (from 72 symbols each) **and** a private mapping of codewords to leaves equals about 72^{36}, which is roughly 2^{222}, about 7×10^{66}.

  $$P(U, 6) \approx 72^{36} \approx 2^{222}.$$

  "Seventy-two to the thirty-sixth equals two to the two hundred twenty-two."

- **What does a quantum computer change?** For totally unstructured guessing, the best known quantum trick (Grover) speeds you up by a **square root**. Square root of 2^{222} is 2^{111}. That's still **humongous**.

- **Why can't we learn patterns instead?** Because each round, the letters are reshuffled, each alphabet ring is independently rotated, and the codewords are private labels. The label you see each round is like rolling a fair six-sided die. No pattern sticks, and no average helps.

- **How many rounds to be sure?** To pin down **all six** secrets and the map you'd need around 86 perfect rounds in a single session, which systems don't give you.

- **Even a giant quantum data center?** It would still take **millions of years** on paper; in practice, building the oracles and handling errors makes it **far worse**.

12B) Professor one-pager (formal bullets)

- **Hypothesis size:** $M = |\mathcal{H}| = 6!\binom{U}{6} = P(U, 6)$, $U = 72^6$. Stirling: $M = \Theta(72^{36})$.

- **Observations:** Per-round board B (balanced partition + independent ring rotations) and label Y. Uniform marginals: $\Pr(Y = j \mid B) = 1/6$. Exchangeable under S_6; across ceremonies, non-identifiable.

- **Information lower bound:** $I(H; B^R, Y^R) \le R \log 6$ bits; Fano $\Rightarrow R \gtrsim \log_6 M \approx 86$ for constant error.

- **Black-box classical:** $\Omega(M)$ consistency checks; time $\Omega(RM)$.

- **Black-box quantum:** BBBV/Zalka $\Omega(\sqrt{M})$ queries; time $\Omega(R\sqrt{M}) \approx \Omega(R \cdot 72^{18})$.

- **Multiple marked items:** with $k \approx M/6^R$ survivors, amplitude amplification costs $\Theta(\sqrt{M/k}) = \Theta(6^{R/2})$ until R drives k to $O(1)$, whence it reverts to $\Theta(\sqrt{M})$.

- **No Shor-type structure:** fresh randomized partitions, independent ring rotations, private label symmetry. QAOA/HHL inapplicable.
- **Parallelization:** p quantum processors $\Rightarrow \sqrt{M/p}$ scaling; with $p = 10^{12}$, still 2.5×10^{27} queries.

13) A Walkthrough: "Quantum Eve" Tries Anyway

1. **Data collection:** Eve records R rounds from a single session: each board B_i and each label Y_i. Across sessions this doesn't help—labels re-scramble and rings re-index.
2. **Oracle design:** For a candidate $H = (S, \varphi)$, she builds a **reversible** circuit that checks all R rounds for consistency and flips a phase if all match. This is a single "Grover oracle" call—but it is expensive: each round requires computing the true leaf for the next character and comparing with $\varphi(Y_i)$, all reversibly.
3. **Grover iterations:** She must repeat the Grover diffusion $\Theta(\sqrt{M}) \approx 2^{111}$ times if R is large enough to leave only one survivor; otherwise, with k survivors she needs $\Theta(\sqrt{M/k}) = \Theta(6^{R/2})$ steps just to hit **one** survivor—still not the guaranteed truth.
4. **Error correction and coherence:** Every iteration adds depth; to keep logical error small across 2^{111} steps, the surface-code overhead skyrockets. Billion-qubit labs do not come close.
5. **Outcome:** After all that, expected time is still **millions of years** in the friendliest arithmetic; realistically, it is **effectively infinite**.

14) Closing Equation Deck

- **Universe of 6-char strings:**

$$U = 72^6.$$

"U equals seventy-two to the sixth."

- **Private maps:**

$$6! = 720.$$

"Six factorial equals seven hundred twenty."

- **Global hypothesis count:**

$$M = 6!\binom{U}{6} = P(U, 6) = U(U-1)\cdots(U-5) \approx 72^{36} \approx 2^{222}.$$

"M equals six factorial times U choose six equals P of U six, about seventy-two to the thirty-sixth, about two to the two-hundred twenty-two."

- **False-accept chance over R rounds:**

$$\Pr[\text{random pass}] = (1/6)^R.$$

"Probability equals one over six to the R."

- **Information per round (bits):**

$$\log_2 6 \approx 2.585.$$

"Log base two of six is about two point five eight five."

- **Rounds to uniqueness (info bound):**

$$R \gtrsim \frac{\log_2 M}{\log_2 6} \approx 86.$$

"R is about eighty-six."

- **Quantum query lower bound:**

$$\sqrt{M} \approx 72^{18} = 2^{111}.$$

"Square root of M equals seventy-two to the eighteenth, two to the one-hundred eleven."

- **Multiple-survivor (amplitude amplification):**

$$k \approx M/6^R \quad \Rightarrow \quad \text{queries} \sim \sqrt{M/k} = 6^{R/2}.$$

"k equals M over six to the R, so queries scale as six to the R over two."

- **Parallel Grover across p machines:**

$$\text{time} \sim \sqrt{M/p}.$$

"Time scales like square root of M over p."

15) Final Verdict (both audiences, one line each)

Undergrad / Lay Description: Quantum gives a **square-root** speedup, but the haystack is so cosmic that even the square root is still basically infinity; and the labels you see are as informative as fair dice.

Post Doctoral: With $M = 6!\binom{72^6}{6}$ hypotheses and a per-round observation channel that is uniform under S_6 label-switching with independent ring re-indexing, we have $I(C; Y \mid B) = 0$ in expectation, an information lower bound $R \gtrsim \log_6 M \approx 86$ for unique identification in a single session, and black-box lower bounds $T_{\text{classical}} = \Omega(RM)$, $T_{\text{quantum}} = \Omega(R\sqrt{M}) \approx \Omega(R \cdot 72^{18})$. Consequently, **even a fleet of large-scale error-corrected quantum processors** achieves only a negligible advantage: the problem remains **computationally intractable** by design.

Permutation Search Space Defined as 10x Mantissa

Definitions

- Alphabet size: $72 = 30 + 30 + 12$
- Secret length: 6
- Universe of 6-char strings: $U = 72^6$

```
U = 139314069504
```

- Private map (bijection of 6 codewords to 6 leaves):

```
6! = 720
```

1) Classical search space for all six secrets + the private map

We present three flavors:

A) Lower bound (no repetitions; all six secrets distinct)

This is what you had before (kept here for context). It equals the ordered falling product:

$$M_{\text{distinct}} = U(U-1)(U-2)(U-3)(U-4)(U-5) = 6!\binom{U}{6}.$$

Exact integer (67 digits):

```
7310883635775654043105842610682888723294659550625996333083000832000
```

B) Tight "allow repetitions" bound (unordered multiset of size 6 + map)

When repetitions are allowed and the six secrets are treated as a multiset, the exact count with the private map is:

$$\boxed{M_{\text{rep, exact}} = 6!\binom{U+5}{6} \;=\; U(U+1)(U+2)(U+3)(U+4)(U+5).}$$

Exact integer (67 digits):

```
7310883637349985407323060214097837638887415242541357398255447572480
```

This is the **tight** upper bound for "allow repetitions" when you model the six secrets as an unordered multiset (repetition allowed) and multiply by the 6! private maps.

C) Very loose upper bound (ordered 6-tuple with repetition $\times$ map)

If you simply allow any ordered 6-tuple of secrets (repetition allowed) and then multiply by the map, you get

$$M_{\text{rep, loose}} = 6!\,U^6.$$

Exact integer (70 digits):

```
5263836218325230202131351810633731616026983108257607972605138168709120
```

This is a **conservative** (loose) upper bound; it overcounts heavily relative to the exact multiset model.

2) Quantum lower bound (Grover/BBBV): $\sqrt{}$(search space)

For an **unstructured** search of size M, any quantum algorithm needs on the order of $\sqrt{M}$ oracle queries (and each query here embeds all per-round consistency checks). Below are the square roots for the exact "allow repetitions" bound and for the very loose bound, plus the "optimal" Grover iteration count $\lceil(\pi/4)\sqrt{M}\rceil$.

A) For the tight allow-repetitions exact space $M_{\textbf{rep, exact}}$

$\lfloor\sqrt{M_{\textbf{rep, exact}}}\rfloor$, $\lceil\sqrt{M_{\textbf{rep, exact}}}\rceil$ **(34 digits each):**

```
floor = 2703864574521066024930942124130308
 ceil  = 2703864574521066024930942124130309
```

Grover "optimal" iteration count $\lceil(\pi/4)\sqrt{M_{\textbf{rep, exact}}}\rceil$ **(34 digits):**

```
2123610270904268182573562460690868
```

B) For the very loose space $M_{\textbf{rep, loose}} = 6!\,U^6$

$\lfloor\sqrt{M_{\textbf{rep, loose}}}\rfloor$, $\lceil\sqrt{M_{\textbf{rep, loose}}}\rceil$ **(35 digits each):**

```
floor = 725522998830859269000600034773643467
 ceil  = 725522998830859269000600034773643468
```

Grover "optimal" iteration count $\lceil(\pi/4)\sqrt{M_{\textbf{rep, loose}}}\rceil$ **(35 digits):**

```
569824430784365884844468999996258463
```

3) Order-of-magnitude summaries in 10^x form

Below are the same quantities summarized as **mantissa** $\times$ **10^exponent** and, where helpful, also as just **10^x** scale.

Classical spaces

- $M_{\text{distinct}} = U(U-1)\cdots(U-5)$
  - $\approx \mathbf{7.310883635775654} \times 10^{\mathbf{66}}$
  - $\log_{10} M_{\text{distinct}} \approx \mathbf{66.8639698714789}$
- $U^6 = 72^{36}$ (reference)
  - $\approx \mathbf{7.310883636562820} \times 10^{\mathbf{66}}$
  - $\log_{10} U^6 \approx \mathbf{66.8639698715257}$
- $\boxed{M_{\text{rep, exact}} = U(U+1)\cdots(U+5)}$ **(tight allow-repetitions)**
  - $\approx \mathbf{7.310883637349985} \times 10^{\mathbf{66}}$
  - $\log_{10} M_{\text{rep, exact}} \approx \mathbf{66.8639698715724}$
- $M_{\text{rep, loose}} = 6!\,U^6$ **(very loose upper bound)**
  - $\approx \mathbf{5.263836218325230} \times 10^{\mathbf{69}}$
  - $\log_{10} M_{\text{rep, loose}} \approx \mathbf{69.7213023679569}$

  Interpretation: the **tight** "allow repetitions" space sits just above U^6 (they differ by a tiny factor $\prod_{k=1}^{5}(1+k/U) \approx 1+O(1/U)$). The "very loose" space multiplies by an extra 720 and is roughly **three orders of magnitude** larger ($\sim 10^{69.72}$ vs. $\sim 10^{66.86}$).

Quantum square roots (Grover/BBBV)

- $\sqrt{M_{\text{rep, exact}}}$
  - $\approx \mathbf{2.703864574521066} \times 10^{\mathbf{33}}$
  - $\log_{10} \sqrt{M_{\text{rep, exact}}} \approx \mathbf{33.4319849357862}$
- $\sqrt{M_{\text{rep, loose}}}$
  - $\approx \mathbf{7.255229988308592} \times 10^{\mathbf{34}}$
  - $\log_{10} \sqrt{M_{\text{rep, loose}}} \approx \mathbf{34.8606511839785}$
- Grover iterations $\lceil (\pi/4)\sqrt{M} \rceil$ scale:
  - For $M_{\text{rep, exact}}$: $\approx \mathbf{2.123610270904268} \times 10^{\mathbf{33}}$
  - For $M_{\text{rep, loose}}$: $\approx \mathbf{5.698244307843659} \times 10^{\mathbf{34}}$

Quick "at a glance" bounds

- **Tight allow-repetitions (multiset + map):**
  $\boxed{M_{\text{rep, exact}} = U(U+1)(U+2)(U+3)(U+4)(U+5)}$
  Raw:
  731088363734998540732306021409783763888741524254135739825544757248O
  Order: $\approx 7.31 \times 10^{66}$

- **Lower bound (distinct only):**
  $M_{\text{distinct}} = U(U-1)\cdots(U-5)$
  Raw:
  731088363577565404310584261068288872329465955062599633308300083200O
  Order: $\approx 7.31 \times 10^{66}$

- **Very loose repetitions (ordered 6-tuple $\times$ map):**
  $M_{\text{rep, loose}} = 6!\,U^6$
  Raw:
  526383621832523020213135181063373161602698310825760797260513816870912O
  Order: $\approx 5.26 \times 10^{69}$

- **Quantum lower bound (Grover/BBBV):**
  $\sqrt{M_{\text{rep, exact}}} \approx 2.70 \times 10^{33}$,
  $\sqrt{M_{\text{rep, loose}}} \approx 7.26 \times 10^{34}$.

Bottom line (why the "allow repetitions" case doesn't change feasibility)

Allowing repetitions increases the hypothesis count from the **distinct** falling product $U(U-1)\cdots(U-5)$ to the **rising** product $U(U+1)\cdots(U+5)$. For our $U = 72^6$, both sit tightly around 7.31×10^{66} total possibilities for "six secrets plus the private map." The **quantum** lower bound still requires on the order of $\sqrt{M}$ oracle queries—i.e., around 10^{33} coherent queries even in the tight allow-repetition model—placing the attack far beyond any realistic (or even science-fiction) capability.

Why Problem Space Explodes When Each Secret Grows from 6 to 12 Characters

Overview

In this proof-of-knowledge design, Alice holds **six distinct secrets**, each a string over an alphabet Σ of size $|\Sigma| = 72$. A public "board" partitions Σ into six balanced leaves every round; ring rotations and Alice's private bijection

make observed answers look like uniform six-way symbols. An adversary observing transcripts must therefore do **unstructured search** over all candidate sextuples of secrets and a label permutation.

Let the **length** of each secret be L. The key count is the hypothesis space

$$|\mathcal{H}_L| \;=\; 6!\binom{U_L}{6} \;\approx\; U_L^6 \quad \text{with} \quad U_L = |\Sigma|^L = 72^L,$$

where the approximation uses $U_L \gg 6$. Thus

$$|\mathcal{H}_L| \;\approx\; (72^L)^6 \;=\; 72^{6L} \;=\; 2^{\,6L\log_2 72}.$$

At baseline $L = 6$: $|\mathcal{H}_6| \approx 72^{36} \approx 2^{222}$.
At extended $L = 12$: $|\mathcal{H}_{12}| \approx 72^{72} \approx 2^{444}$.

The jump from $L = 6$ to $L = 12$ therefore turns the search space from roughly 2^{222} to roughly 2^{444}, i.e., it **squares** the already-astronomical space.

The Square Law: Doubling Length Squares the Search Space

Because $U_L = 72^L$, doubling L **squares** U_L. Since $ |\mathcal{H}_L| \approx U_L^6$, we get

$$\frac{|\mathcal{H}_{12}|}{|\mathcal{H}_6|} \;\approx\; \frac{(U_6^2)^6}{U_6^6} \;=\; U_6^6 \;=\; |\mathcal{H}_6|.$$

Equivalently,

$$|\mathcal{H}_{12}| \;\approx\; \big(|\mathcal{H}_6|\big)^2.$$

So the **adversary's candidate set is squared** when each of the six secrets doubles in length from 6 to 12. If $|\mathcal{H}_6|$ was already beyond the reach of realistic computation, $|\mathcal{H}_{12}|$ is the square of "beyond reach."

Numerically,

$$|\mathcal{H}_6| \approx 2^{222} \;\Rightarrow\; |\mathcal{H}_{12}| \approx 2^{444}.$$

This alone explains the "explosion": every brute-force approach (classical or quantum) faces an additional **factor of** 2^{222} in work compared to the already infeasible $L = 6$ case.

Per-Round Information Is Capped; Needed Rounds Double

Each round reveals at most $\log_2 6 \approx 2.585$ bits (a six-way label), and by design the expected mutual information about the next character is essentially 0 due to

balanced partitions and rotations. To isolate a unique hypothesis in expectation, the adversary needs about

$$R_{\min}(L) \;\approx\; \left\lceil \frac{\log_2 |\mathcal{H}_L|}{\log_2 6} \right\rceil = \left\lceil \frac{6L \log_2 72}{\log_2 6} \right\rceil .$$

Since $\log_2 72 / \log_2 6 \approx 2.387$,

$$R_{\min}(L) \;\approx\; \lceil 14.32\, L \rceil .$$

Therefore,

- $L = 6 \Rightarrow R_{\min} \approx 86$ rounds,
- $L = 12 \Rightarrow R_{\min} \approx 172$ rounds.

So the **minimum rounds roughly double** when the length doubles. That's intuitive: you have twice as many symbols' worth of uncertainty to discharge, and each round still only leaks a bounded amount.

Classical Cost: From Astronomical to "Squared Astronomical"

In a black-box model, the best classical strategy is eliminative consistency: propose (S, φ), check it against the transcript (cost $O(R)$), repeat. The runtime is

$$T_{\text{class}}(L) \;=\; \Omega\big(R_{\min}(L) \cdot |\mathcal{H}_L|\big).$$

Comparing $L = 6$ to $L = 12$,

$$\frac{T_{\text{class}}(12)}{T_{\text{class}}(6)} \;\approx\; \frac{(2R_{\min}(6)) \cdot (|\mathcal{H}_6|)^2}{R_{\min}(6) \cdot |\mathcal{H}_6|} \;=\; 2 \cdot |\mathcal{H}_6| \;\approx\; 2 \cdot 2^{222} \;=\; 2^{223}.$$

So **classical work multiplies by about** 2^{223}—a factor of roughly 10^{67}. If $L = 6$ already required on the order of 10^{44} years at absurdly optimistic 10^{15} checks/second, multiplying by 10^{67} pushes the wall-clock into a regime that makes **cosmological timescales look negligible**.

The constant 6! from the private bijection sits in the noise compared to these exponents; the blow-up is governed by 72^{6L}.

Quantum Cost: Grover Still Drowns

For unstructured search, Grover/BBBV gives at best a square-root speedup. The quantum cost is

$$T_{\text{quant}}(L) \;=\; \Omega\big(R_{\min}(L)\cdot\sqrt{|\mathcal{H}_L|}\big) \;=\; \Omega\big(R_{\min}(L)\cdot 72^{3L}\big).$$

Hence the ratio

$$\frac{T_{\text{quant}}(12)}{T_{\text{quant}}(6)} \;\approx\; \frac{(2R_{\min}(6))\cdot 72^{36}}{R_{\min}(6)\cdot 72^{18}} \;=\; 2\cdot 72^{18} \;=\; 2\cdot 2^{18\log_2 72} \;\approx\; 2^{112}.$$

So even with a hypothetical **ideal** quantum machine, the jump from $L = 6$ to $L = 12$ imposes **another factor of** $\approx 2^{112}$ in runtime. And this ignores the severe realities of fault-tolerant quantum computing: the reversible **oracle** that checks a hypothesis across R rounds must itself be built from deep, error-corrected circuits with large logical-qubit footprints; doubling L doubles the number of rounds to encode and typically worsens depth and width.

One striking perspective: at $L = 12$,

$$\sqrt{|\mathcal{H}_{12}|} \;=\; 2^{222},$$

which equals the **entire classical** search space at $L = 6$. In other words, the quantum attacker at $L = 12$ faces (up to linear factors in R) the same exponent that already made the **classical** attack at $L = 6$ impossible.

Security Levels: Massive Headroom, Even Post-Quantum

Security targets are often quoted as "bits of work."

- **Classical 128-bit target:** need $|\mathcal{H}_L| \geq 2^{128}$.
  At $L = 12$, $|\mathcal{H}_{12}| \approx 2^{444}$ — a margin of 2^{316} beyond the 128-bit bar.
- **Post-quantum 128-bit target:** need $\sqrt{|\mathcal{H}_L|} \geq 2^{128} \Rightarrow |\mathcal{H}_L| \geq 2^{256}$.
  At $L = 12$, $\sqrt{|\mathcal{H}_{12}|} \approx 2^{222} \gg 2^{128}$, giving a 2^{94} headroom over the post-quantum 128-bit bar. (By contrast, $L = 6$ meets **classical** 128-bit but sits below the **formal** post-quantum 128-bit line; increasing to $L = 7$ crosses it. Jumping to $L = 12$ is far beyond that threshold.)

Thus, doubling length delivers not a marginal improvement but a **regime change**: both classical and quantum brute force become even more fantastically implausible, with comfortable formal margins.

Why This Happens: Entropy Grows Linearly, Search Grows Exponentially

Per round, the transcript can leak at most $\log_2 6$ bits. Meanwhile, the **uncertainty** the attacker must discharge is

$$\log_2 |\mathcal{H}_L| \;\approx\; 6L \log_2 72 \;\approx\; 37.02\, L \text{ bits.}$$

Therefore the **required rounds** scale **linearly** with L: $R_{\min}(L) \approx 14.32L$. But the **work** grows like $ |\mathcal{H}_L|\approx 72^{6L}$, i.e., **exponentially** in L. This mismatch—**bounded per-round leakage** versus **exponential hypothesis growth**—is the engine of the explosion. Doubling L doubles the bits to be discharged and the rounds, but **squares** the search space.

Engineering Reality: Oracles, I/O, and Energy Get Worse Too

Even if we ignore asymptotics, practical costs balloon:

1. **Oracle construction.** The reversible oracle (quantum) or the consistency checker (classical) must encode R rounds; at $L = 12$, that's roughly twice as many rounds, more comparisons, and deeper circuits.
2. **I/O and memory.** Enumerating, moving, or caching hypotheses for 2^{444} candidates is infeasible on any conceivable storage fabric; paging dominates.
3. **Energy/time bounds.** Landauer-limit back-of-the-envelope checks make it clear that the energy required to touch that many hypotheses (even once) is cosmologically outrageous.

These are not proofs—your **proof** is the unstructured-search lower bound—but they underline "no way in practice."

Bottom Line

Moving from **six 6-character secrets** to **six 12-character secrets** multiplies the attacker's candidate space from $|\mathcal{H}_6| \approx 2^{222}$ to $|\mathcal{H}_{12}| \approx 2^{444}$, i.e., it **squares** an already intractable space. Minimum rounds roughly **double** (from ~ 86 to ~ 172), per-round leakage remains capped, and both classical $\Omega(R|\mathcal{H}|)$ and quantum $\Omega(R\sqrt{|\mathcal{H}|})$ runtimes explode by factors of about 2^{223} and 2^{112}, respectively. In security-bit terms, $L = 12$ offers enormous headroom, easily clearing even **post-quantum** 128-bit targets. The explosion is the inevitable result of **linear information per round** confronting **exponential** growth of the hypothesis set with length—a design choice that decisively favors the defender.

Bibliography

1. **Claude E. Shannon, "Communication Theory of Secrecy Systems" (1949)**

Shannon's seminal paper where he introduces the concepts of *confusion* and *diffusion* as essential security properties of ciphers. These ideas form the theoretical foundation for your protocol's emphasis on eliminating structurealigning with Section 11, *"Why Structurelessness Matters."* (Wikipedia)

2. **Wikipedia—"Confusion and diffusion"**
   Defines confusion as obscuring the key–ciphertext relationship and diffusion as spreading plaintext statistics across ciphertext. Directly relates to your protocol's balanced partitions and ring rotations mechanisms that mimic diffusion-like effects. (Wikipedia)

3. **Kevin Sookocheff, "Cryptography for the Everyday Developer: Confusion, Diffusion, and S-P Networks" (2025)**
   Explains how substitution (confusion) and permutation (diffusion) layers combine in SP-networks—parallels your essay's use of private bijection (confusion) and ring-based rotations (diffusion). (Kevin Sookocheff)

4. **Wikipedia—"Substitution–permutation network" (SPN)**
   Discusses SPNs as layered structures implementing confusion and diffusion. Highlights the avalanche effect—comparable to your design's elimination of statistical leaks and cross-round anchors. (Wikipedia)

5. **Wikipedia—"Avalanche effect"**
   Describes the property that slight changes in input should greatly alter output—a manifestation of diffusion. This concept echoes your protocol's goal: ensuring outputs (leaf codewords) reveal no systematic bias across rounds. (Wikipedia)

6. **NKU PDF "Diffusion and Confusion" (Trappe & Washington)**
   Offers a clear articulation: confusion obscures key-ciphertext relations; diffusion spreads plaintext structures. Provides context for your protocol's reliance on combinatorial symmetry and obfuscation of structure. (websites.nku.edu)

7. **SJSU page "Confusion and Diffusion"**
   Summarizes Shannon's definitions and illustrates why diffusion alone is insufficient—emphasizes that modern ciphers require both. Reinforces your essay's observation that your protocol uses a combination of mechanisms to fully suppress signal. (SJSU Computer Science)

8. **Wikipedia—"Communication Theory of Secrecy Systems"**
   Notes the importance of Shannon's 1949 paper and its foundational role in modern cryptography. Supports the historical grounding of your methodology. (Wikipedia)

9. **Wikipedia—"Hill cipher"**
   Cites the Hill cipher as an early cipher achieving diffusion via linear mixing (matrix multiplication), akin to your protocol's spreading of symbol properties across leaves. (Wikipedia)

10. **Wikipedia—"Substitution–permutation network"** (again, more details)
    Explains how repeated substitution and diffusion layers produce avalanche effects and confusion. This parallels the combined effect of your balanced partitions and private bijection. (Wikipedia)

11. **Wikipedia—"Communication Theory of Secrecy Systems"** (reiterated for depth)
    Emphasizes that Shannon's work formalized confusions and diffusions—applicable to your essay's conceptual pillars. (Wikipedia)

12. **Cryptography Notes (Rutgers University)**
    Briefly defines confusion and diffusion consistent with Shannon's theory, reinforcing your protocol's design principles. (cs.rutgers.edu)

13. **Wikipedia—"Confusion and diffusion"** (foundational definition)
    Restates Shannon's concepts succinctly, giving authoritative grounding to your thesis's theoretical basis. (Wikipedia)

14. **Wikipedia—"Communication Theory of Secrecy Systems"** (again)
    Re-emphasizes the paper's historical significance for completeness in the bibliography. (Wikipedia)

15. **Wikipedia—"Substitution–permutation network"** (again repeated for thoroughness)
    Reinforces SPN's role in combining substitution and permutation—a structure analogous to your design. (Wikipedia)

16. **Wikipedia – "Information-theoretic security"**
    Defines *unconditional security*—security guaranteed even against adversaries with unlimited computational power—rooted in Shannon's perfect secrecy concepts, relevant to your protocol's information-theoretic framing. (Wikipedia)

17. **Nigel Smart – *Cryptography: An Introduction*** (Chapter on information-theoretic vs. computational security)
    Delivers a foundational distinction between computational and information-theoretic security—contextualizes your commentary in Section 8 on *limitations and assumptions*. (Computer and Information Science)

18. **Erdal Arikan – "An Information-Theoretic Analysis of Grover's Algorithm" (arXiv, 2002)**
    Provides an information-theoretic proof that Grover's square-root speedup is optimal—aligns with your quantum lower-bounds derivation. (arXiv)

19. **Wikipedia – "Grover's algorithm"**
    Explains the algorithm's $O(\sqrt{N})$ query complexity and optimality in the black-box model—supports your Section 8 analysis. (Wikipedia)

20. **Simon Apers – "Lecture 2: Grover's algorithm and lower bounds"** (PDF Lecture notes, 2024)
Formalizes query complexity lower bounds for unstructured search, echoing your use of the BBBV bound. (simonapers.github.io)

21. **CS UMD – "Unstructured Search and Lower Bounds"** (PDF lecture, 2008)
Presents a proof sketch that Grover's algorithm is optimal for unstructured search queries—reinforces your security arguments against quantum optimization. (Computer Science at UMD)

22. **CS U. Wisconsin – "Lecture: Query Lower Bounds and Applications of Grover's Algorithm"**
Covers amplitude amplification and the proof of optimality "up to constant factors"—supplements your elimination and rounds-to-uniqueness logic. (UW Computer Sciences)

23. **Medium – "Quantum Sundays: Quadratic Speedup in Unstructured Search..."**
Clarifies that Grover provides only quadratic speedups, even for very large N—illustrative for your quantum infeasibility sections. (Medium)

24. **Classiq.io – "Grover's Algorithm" Explanation (2024)**
Summarizes unstructured quantum search advantage—useful population-level illustration supporting your framing. (classiq.io)

25. **Wikipedia – "Communication Theory of Secrecy Systems"**
Reiterates Shannon's pivotal 1949 paper foundational to modern cryptography—including confusion/diffusion basis. (Wikipedia)

26. **Wikipedia – "Confusion and diffusion"** (further elaboration)
Delves into Shannon's definitions and how substitution and permutation layers implement these principles—reinforcing your Section 1 and 11. (Wikipedia)

27. **GeeksforGeeks – "Difference between Confusion and Diffusion"**
Offers accessible definitions of confusion (substitution) vs. diffusion (transposition)—complements pedagogical framing. (GeeksforGeeks)

28. **TutorialsPoint – "What is Diffusion in Information Security?"**
Explains diffusion via plaintext symbol spread—drawing parallels to your balanced partition mechanisms. (TutorialsPoint)

29. **Sookocheff – "Confusion, Diffusion, and Substitution-Permutation Networks"** (2025)
Illustrates how rounds of substitution and permutation create secure block ciphers—analogous to your protocol's round structure. (Kevin Sookocheff)

30. **ArXiv – "Universal Hashing for Information Theoretic Security"** Shows how universal hashing enables efficient, provable security even in the information-theoretic model—an extended theoretical anchor for your arguments. (arXiv)