The Observer's Bit Budget in a Full Rosario Cipher Interaction

A Quantitative Information-Theoretic Model of Shoulder-Surfing Across Rings, Hyperplanes, the Witness Synonym, and the Pass/Fail Bit

2026

F.D. ROSARIO
DR. LIN WANG
DR. DAVID SERENA
DR. FRANCISCO PEREZ

WWW.ENI6MA.COM

contact@eni6ma.co

Abstract

We give a quantitative, first-principles information-theoretic analysis of how much an adversary can learn by observing a complete authentication ceremony in the implemented ENI6MA / Rosario–Wang construction. The construction comprises six colored zones, six bearings, three disjoint rotating alphabet rings of sizes 26/26/10, a foliation into six leaves, a private $6! = 720$ color–bearing bijection, and a single accumulator bit. We separate two quantities that the word "bits" routinely conflates: *surface bits* (the gross Shannon entropy of the symbols an observer can physically record) and *effective bits* (the secret-relevant mutual information those recordings actually pin down). We prove the central pair of results

$$H_{\text{surface}}(\text{one character}) = \log_2 6 \approx 2.585 \text{ bits}, \qquad I_{\text{effective}}(\text{secret}\,;\,\text{observation}) = 0 \text{ bits},$$

and we show the second identity is preserved across an L-round ceremony and across any number n of observed ceremonies. The zero is *structural and unconditional* in the Shannon sense — the information is absent from the observation rather than merely expensive to extract — so it survives unlimited offline compute and quantum cryptanalysis, leaving only rate-limited brute force against the verifier. Three independent mechanisms enforce the result: a lossy hyperplane collapse that destroys ring/position/alphabet structure, a private witness synonym (one of 720 unknown bijections) that severs the bearing transcript from the colored challenge, and a per-round one-time XOR pad that uniformizes the wire payload by Vernam's theorem. We present a self-contained six-lemma proof that $I(C;O) = 0$, relate it explicitly to Shannon's 1949 perfect-secrecy theorem, give a holographic (bulk/boundary/dictionary) reading of the dimensional collapse, and situate the construction against modern computational and NIST post-quantum cryptography (ML-KEM, ML-DSA, SLH-DSA).

Keywords / Index Terms

Information-theoretic security; perfect secrecy; Shannon entropy; mutual information; one-time pad (Vernam cipher); zero-knowledge / proof of knowledge; shoulder-surfing resistance; authentication; post-quantum cryptography; holographic principle; data-processing inequality.

A threat-model analysis for security engineers and architects. All claims are grounded in the

implemented ENI6MA / Rosario–Wang construction: six colored zones, six bearings, three rotating alphabet rings of sizes 26/26/10, foliation into six leaves, the $6! = 720$ *color–bearing bijection, and the single accumulator bit.*

Editorial Note — Figure Program and Illustration Style

To make the argument legible at a glance, the manuscript carries a coordinated figure program: **one principal illustration per section (Figures 1–13, where Figure** N **accompanies Section** N**)**, with additional sub-figures permitted where a single image cannot carry the load. All figures use **monochrome wireframe line art** suitable for black-and-white print. Because the medium is black-and-white, the six *colors* of the challenge are rendered throughout as **labeled tiles with distinct line-hatch patterns** (e.g., diagonal, cross-hatch, dotted, stippled, vertical, blank) rather than as filled colors, and the six *bearings* as **labeled directional arrows**.

---

0. The Challenge Question: How Many Bits Can the Observer Actually Use?

This document answers one operational question that any security team must ask before trusting an authentication scheme against shoulder-surfing, screen capture, malware overlays, keyloggers, and on-wire interception:

> ***Challenge question.*** An adversary observes a full authentication ceremony with the best instrumentation available, a camera over the user's shoulder, a screen recorder, a network tap on the wire, and unlimited storage and compute to analyze it all afterward. How many bits of the user's secret can that adversary actually **use**?

The answer, stated up front so there is no suspense, is a single number:

> ***Bottom line.*** *The observer can* ***record*** *about* 2.585 *bits per character entered (one of six bearings), accumulating to roughly* $2.585 \times L$ *bits over an* L*-character ceremony plus one final accept/reject bit, but the number of those bits that are* ***usable*** *to recover, narrow, or even slightly bias a guess at the secret is* ***exactly zero****. Recording is not the same as learning. Capture is not the same as compromise.*

This is the distinction that the proof in this document provides. In operational terms it means a captured session, by camera, by screen recorder, by network sniffer, by malware, is a **forensic dead end**: it can prove that *an* authentication happened, but it cannot be replayed, cannot be correlated across sessions, and cannot be mined by any present or future analysis to yield the credential. The apparent paradox ("you saw everything, yet you got nothing usable") dissolves once we separate **two completely different quantities** that the word "bits" can denote, and a precise threat model must keep them distinctly apart:

1. **Surface bits** (also: *observable bits*, *gross transcript entropy*): the raw Shannon entropy of the symbols an observer can physically record, which bearing was pressed, what the screen displayed. This number is strictly positive and grows linearly with the length of the witness ceremony.
2. **Effective bits** (also: *leaked bits*, *secret-relevant mutual information*): the number of bits of *the user's secret* that those recordings actually pin down. This number is the mutual information between the secret and the observation, and the design goal is that it equals exactly **zero**.

The central result of this analysis is the following pair of equations, which we derive from first principles:

$$H_{\text{surface}}(\text{one character}) = \log_2 |\mathcal{B}| = \log_2 6 \approx 2.585 \text{ bits},$$

$$I_{\text{effective}}(\text{secret} \,;\, \text{observation}) = 0 \text{ bits}.$$

Watching a single character being entered yields about 2.585 bits of *something to write down* (one of six bearings), but those 2.585 bits contain zero bits of information about the secret character the user was actually proving. The observer's notebook fills up; the observer's knowledge of the secret does not.

Parameter breakdown. - $H_{\text{surface}}(\cdot)$, the Shannon entropy of the observable symbol for one character entry, measured in bits. - $\mathcal{B}$, the set of bearings (the response alphabet); $|\mathcal{B}| = 6$. - $\log_2 6$, the base-2 logarithm of six; the number of bits needed to name one of six equally likely outcomes. - $I_{\text{effective}}(\cdot\,;\,\cdot)$, the mutual information (in bits) between the user's secret and the observer's full record.

The remainder of the document walks through the machinery that forces the second equation to hold even while the first remains positive: the disjoint rings, the foliation collapse to colored hyperplane zones, the private witness synonym (the color–bearing map), and the one-time XOR mask applied by the witnessing agent. Each section names the real-world attack it neutralizes, so a security reviewer can map the construction directly onto a threat model.

Where this sits in the cryptographic landscape. It is worth stating up front, for a reader who lives among RSA, elliptic curves, and the new post-quantum standards, what kind of guarantee this is. Almost every deployed cipher and signature today, and almost every scheme in the post-quantum migration, rests on **computational hardness**: security holds only as long as some problem (factoring, discrete logarithms, lattice problems) stays infeasible for the adversary's hardware. That is *conditional* security, and it has an expiry risk, the arrival of a better algorithm or a large enough quantum computer. The guarantee analyzed in this document is of a fundamentally different kind: it is **information-theoretic** (also called *unconditional* or *Shannon*) security, which does not assume the attacker is computationally limited at all. When we write $I_{\text{effective}} = 0$, we mean the information is *absent* from the observation, not merely *expensive to extract*, so no future algorithm and no quantum computer can recover what was never present. Section 11 makes the link to Shannon's 1949 perfect-secrecy theorem explicit, and Section 12 places the construction in historical context and contrasts it with modern computational cryptography and the NIST post-quantum standards (ML-KEM / Kyber, ML-DSA / Dilithium, and the hash-based signatures). A reader who wants the geometric intuition first should also see the new Section 3.1, which reads the dimensional collapse through the lens of the holographic principle.

Why this distinction is the whole battlefield. Almost every classical attack on an authentication or encryption scheme is, at heart, an attempt to convert surface bits into effective bits, to take the symbols one can *see* and squeeze from them some *knowledge* of the secret. Two such conversions dominate the history of cryptanalysis. **Frequency analysis** assumes that the distribution of observed symbols is non-uniform and correlated with the underlying plaintext: in a substitution cipher the most common ciphertext letter is probably "e," and that single statistical lean is enough to unravel everything. **Correlation analysis** is the multi-variable generalization: it looks for any stable statistical dependency, between one observed symbol and another, between

an observation and a timing signal, between repeated sessions, that lets the attacker triangulate the hidden variable. Both attacks succeed precisely when surface bits and effective bits are the *same* bits viewed two ways. The entire design philosophy of the Rosario interaction is to drive a wedge between those two quantities, so that the observable distribution is perfectly flat (defeating frequency analysis), perfectly independent across rounds and sessions (defeating correlation), and the only residual strategy is undirected **brute force** against a space so large that rate limiting makes it hopeless. Each section that follows takes one structural element of the cipher and shows which of these three attack families it neutralizes and by what mechanism. The reader should keep a mental scoreboard: by the end, frequency analysis, correlation, and observation-based inference will all be reduced to the single blind guess, and brute force, throttled by the ledger, will be the only move left on the board.

---

1. Objects and Notation

We fix a vocabulary once and use it for the whole document. These objects are exactly those instantiated in the production circuit binary.

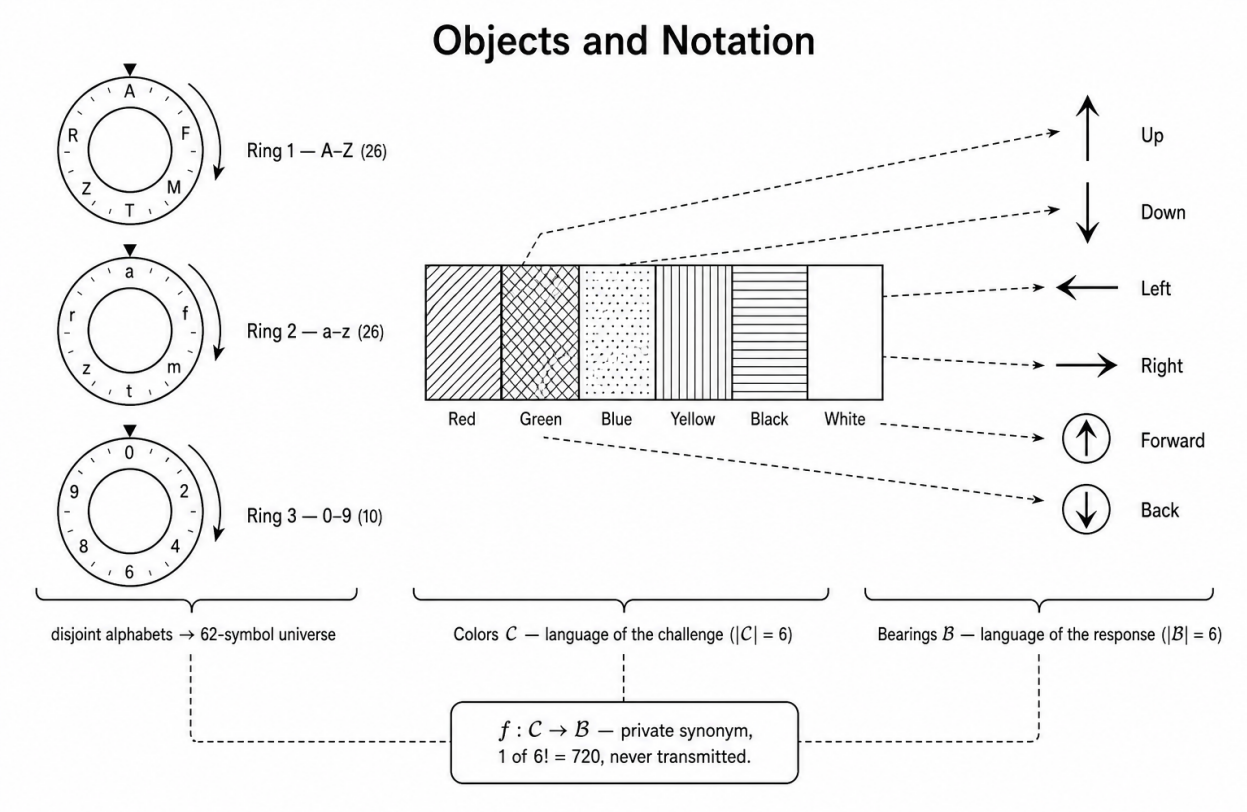

Figure 1: *Figure 1 — The cast of objects: disjoint rings, colored zones, bearings, and the private synonym*

Figure 1. The cast of objects in one frame. Three disjoint alphabet rings (uppercase 26, lowercase 26, digits 10) partition a 62-symbol universe; the screen presents six colored witness zones (the *challenge* language); the user responds with one of six bearings (the *response* language); and the private bijection $f : \mathcal{C} \to \mathcal{B}$ — one of $6! = 720$ secret synonyms — is the only bridge between the

two languages.

Why the bookkeeping matters before the attack. It is tempting to skip notation and rush to the security theorems, but in information theory the definitions *are* the security argument; a sloppy object is an exploitable object. The single most important structural choice introduced here, **disjoint alphabet rings**, is also the first and most under-appreciated line of defense against frequency analysis. Classical frequency analysis feeds on two facts: that symbols recur, and that their recurrence rates differ. ENI6MA's partition of the universe into non-overlapping rings, combined with the later foliation, guarantees that the *thing the user actually selects each round* is not a letter at all but a **zone**, and that every secret character, common or rare, "e" or "z" or "7", is funneled into exactly one of six equiprobable zones. The lopsided letter-frequency table that powered a century of codebreaking is therefore amputated at the root: the secret's identity never reaches the observable channel as a letter, only as a zone choice that is subsequently disguised. Equally, defining the secret as a *commitment that is proven and never transmitted* changes the attacker's job category entirely. There is no ciphertext-of-the-secret to analyze, because the secret is never enciphered and sent; it is only ever used as a private selector. This is the difference between attacking a message and attacking a proof of knowledge, and it is why the rest of the document can speak about *mutual information with the secret* rather than *decryption of the secret*. Get the objects right and the impossibility results fall out almost mechanically.

Disjoint alphabet rings. Let there be M rings, each a distinct alphabet:

$$A = \{A_1, A_2, \dots, A_M\}, \qquad A_i \cap A_j = \emptyset \ \text{ for } i \neq j.$$

The symbol universe is partitioned into M non-overlapping alphabets, one per ring. In the reference implementation $M = 3$: ring 1 is uppercase A–Z, ring 2 is lowercase a–z, ring 3 is digits 0–9. Disjointness guarantees that any single character belongs to exactly one ring, which removes all ambiguity when we later ask "which zone holds your character?"

Parameter breakdown. - A, the collection of all rings. - A_a, the a-th ring (an alphabet); we write $n_a = |A_a|$ for its size. - In the reference design $n_1 = 26$, $n_2 = 26$, $n_3 = 10$, so the universe size is $|S| = \sum_a n_a = 62$.

The secret commitment. The user's secret is an ordered tuple

$$C = (s_1, s_2, \dots, s_L), \qquad s_r \in A_{j(r)},$$

where $j(r)$ is the ring from which the r-th secret character is drawn.

The secret is a word of length L; the r-th character s_r lives in some ring. The user proves knowledge of one character per round, L rounds in total. The secret is **never transmitted**, only proven.

Parameter breakdown. - C, the secret commitment (the word/phrase being proven). - s_r, the r-th secret character, the *target* for round r. - L, the length of the secret, equal to the number of interactive rounds. - $j(r)$, the ring index that contains s_r.

Colors and bearings. Two six-element label sets, taken verbatim from the implementation:

$$\mathcal{C} = \{\text{Red}, \text{Green}, \text{Blue}, \text{Yellow}, \text{Black}, \text{White}\}, \qquad |\mathcal{C}| = 6,$$

$$\mathcal{B} = \{\text{Up}, \text{Down}, \text{Left}, \text{Right}, \text{Forward}, \text{Back}\}, \qquad |\mathcal{B}| = 6.$$

The screen shows six *colored* zones. The user answers with one of six *bearings* (a direction). Colors are what the challenge speaks; bearings are what the response speaks. They are deliberately different languages.

Parameter breakdown. - $\mathcal{C}$, the color set (the language of the *challenge*). - $\mathcal{B}$, the bearing set (the language of the *response*/transcript). - $|\mathcal{C}| = |\mathcal{B}| = 6$, both have six elements, which is what makes a bijection between them possible.

The witness synonym (private color–bearing map). The user's cognitive key is a bijection

$$f : \mathcal{C} \to \mathcal{B}.$$

The user privately renames each color as a bearing, "the red zone is what I call *Left*," and so on, with no two colors sharing a bearing. This is the "private synonym" the user assigns to each witness zone. It exists only in the user's memory and in the verifier's binary; it is never sent on the wire.

Parameter breakdown. - f, the private bijection (the witness synonym / cognitive key). - The number of such bijections is $6! = 720$ (Section 5), so f is one secret choice out of 720.

We will also need the session randomness:

- E, the session entropy pool (typically 256–512 bits of fresh cryptographic randomness).
- τ, the microsecond-resolution timestamp, mixed with E to make each session geometrically unique.
- $r \in \{1, \ldots, L\}$, the round index.
- $H(\cdot)$, a cryptographic hash used as the deterministic engine that turns (E, τ, r) into rotations and masks.

---

2. The Anatomy of a Single Character Entry

We now trace exactly what happens, and exactly what an observer can see, when the user enters **one** character of the witness. Five steps occur per round.

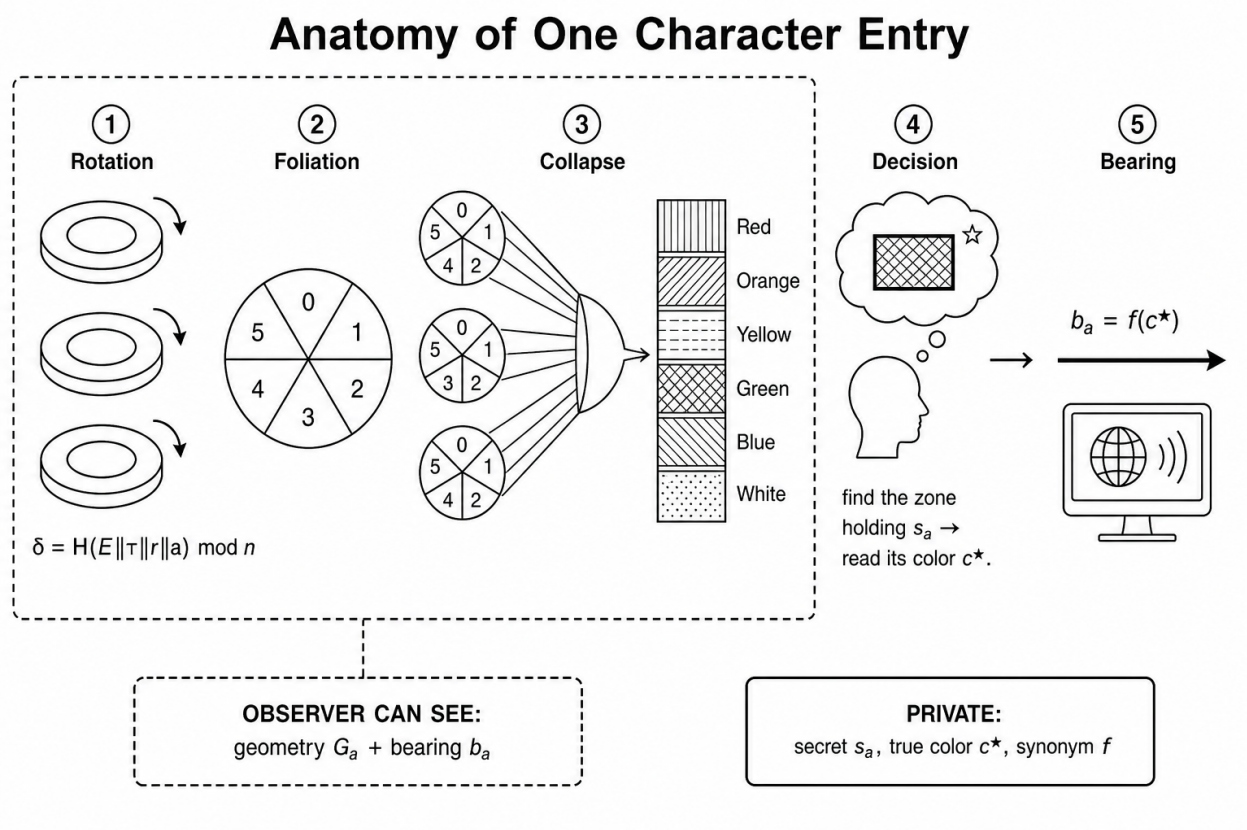

Figure 2: *Figure 2 — The five-step pipeline of a single character entry*

Figure 2. The five-step pipeline of one character entry: (1) entropy-and-time **rotation** of each ring, (2) **foliation** into six arcs, (3) **collapse** of same-index arcs into six colored zones, (4) the user's silent **decision** (the true color of the secret's zone), and (5) the **bearing response** through the private synonym. A dashed boundary marks exactly what the observer can record (geometry G_r and bearing b_r) versus what stays private (secret s_r, true color $c_r^\star$, map f).

Why a step-by-step anatomy defeats correlation. Correlation attacks live in the gaps *between* events: between this round and the next, between the position of a character now and its position a moment ago, between a session today and a session yesterday. An attacker who can establish even one stable across-round relationship, "whenever the user goes *Left* twice in a row, the second secret character tends to be a vowel", has a foothold from which the whole secret can eventually be levered open. The anatomy that follows is engineered specifically to deny every such foothold by making each round a *freshly randomized, independent experiment.* The rotation in Step 1 is reseeded every round from session entropy and a microsecond timestamp, so the board the attacker sees in round r has no learnable relationship to the board in round $r-1$; there is no carry-over geometry to correlate against. Because a single deterministic rotation moves **all** characters at once, there is also no per-character jitter that could be averaged away over many observations, a favorite technique of differential and correlation power analysis. The crucial consequence is that the (geometry, response) pairs the attacker collects form an *independent and identically distributed* sequence of non-informative samples. Correlation analysis is a search for dependency structure; this construction provides none. Each of the five steps below should be read as removing one more potential correlation channel, positional, temporal, cross-round, and cross-session, until the

per-round observation is statistically sealed off from everything before and after it.

Step 1, Rotation. Each ring is rotated by a deterministic, session-specific offset:

$$\delta_a^r = H(E \,\|\, \tau \,\|\, r \,\|\, a) \bmod n_a,$$

and a character at ring-position p moves to position $(p + \delta_a^r) \bmod n_a$.

Before the round, each combination-lock cylinder spins by an amount no observer can predict, because the amount is the output of a hash of secret session entropy and a microsecond timestamp. A single deterministic rotation moves *all* characters at once; there is no per-character randomness to correlate.

Parameter breakdown. - δ_a^r, the rotation offset (in positions) applied to ring a in round r. - $E \,\|\, \tau \,\|\, r \,\|\, a$, the concatenation of entropy, timestamp, round index, and ring index fed to the hash. - $\bmod n_a$, wraps the offset into the valid position range $[0, n_a)$ of ring a. - p, a character's index on its ring before rotation.

Step 2, Foliation. Each rotated ring is sliced into $n = 6$ contiguous arcs (the "leaves"):

$$\Pi_6(A_a, r) = \{\, \alpha_{a,0}^r, \alpha_{a,1}^r, \ldots, \alpha_{a,5}^r \,\}, \qquad \alpha_{a,z}^r = \Big\{ \text{characters of } A_a \text{ in arc } z \text{ after rotation } \delta_a^r \Big\}.$$

Cut each rotated ring like a pie into six wedges. Wedge z of ring a is the set of characters that happen to fall in that angular slice this round. Because the ring was rotated first, *which* characters land in wedge z changes every round.

Parameter breakdown. - Π_6, the foliation operator that produces six slices ("leaves") per ring. - $\alpha_{a,z}^r$, the slice of ring a, zone z, in round r (a set of characters). - $z \in \{0, 1, 2, 3, 4, 5\}$, the zone (slice) index.

Step 3, Collapse to colored hyperplane zones. The slices of equal index across all rings are unioned and assigned a color. This is the central "collapse" the question asks about, and it gets its own section (Section 3). For now:

$$W_z^r = \bigcup_{a=1}^{M} \alpha_{a,z}^r, \qquad \text{color of zone } z = \chi(z),$$

where $\chi : \{0, \ldots, 5\} \to \mathcal{C}$ assigns the six colors to the six zones.

Stack wedge z from ring 1, wedge z from ring 2, and wedge z from ring 3 on top of one another into a single colored zone. The result is six colored zones, each containing a mixture of uppercase, lowercase, and digit characters. This many-rings-into-one-zone stacking is the *hyperplane collapse*: a higher-dimensional structure (three rings) is projected down onto one colored leaf.

Parameter breakdown. - W_z^r, the *witness set* of zone z in round r (the characters shown inside the colored zone). - $\bigcup_{a=1}^{M}$, union across all M rings of the same-indexed slice. - χ, the coloring function assigning a color to each zone index.

Step 4, The user's decision. The user locates the unique zone containing the secret character and reads off its color:

$$c_r^\star = \text{the unique } c \in \mathcal{C} \text{ such that } s_r \in W^r_{\chi^{-1}(c)}.$$

The user scans the six colored zones, finds the one zone that contains this round's secret character, and notes that zone's color. Because the alphabets are disjoint and partition the universe, the secret lands in exactly one zone, so this color is unambiguous.

Parameter breakdown. - $c_r^\star$, the *true color* of the secret's zone in round r (the honest answer, in the challenge's color-language). - s_r, the secret character for this round. - $W^r_{\chi^{-1}(c)}$, the witness set of the zone whose color is c.

Step 5, The bearing response (applying the synonym). The user does *not* announce the color. Instead they apply the private synonym and emit a bearing:

$$b_r = f(c_r^\star) \in \mathcal{B}.$$

The user translates the secret's color through their private map and presses the corresponding direction. The observer sees the bearing b_r and the screen geometry G_r, but never the color $c_r^\star$ and never the map f.

Parameter breakdown. - b_r, the bearing the user emits in round r; this is the only "response" datum. - f, the private color→bearing synonym. - $c_r^\star$, the true color of the secret's zone (from Step 4).

The complete pipeline for one character is therefore:

$$(E, \tau, r) \xrightarrow{\rho} \text{rotated rings} \xrightarrow{\Pi_6} \text{slices} \xrightarrow{\text{collapse}} \{(c, W_c^r)\}_{c\in\mathcal{C}} = G_r \xrightarrow[+\ f]{+\ s_r} b_r.$$

Entropy and time generate the rotated rings; foliation slices them; the collapse builds the six colored witness zones (this collection is the challenge geometry G_r); the user combines the geometry with their secret and their private map to produce one bearing.

---

3. The Hyperplane Collapse: From Rings and Subsets Down to a Color Zone

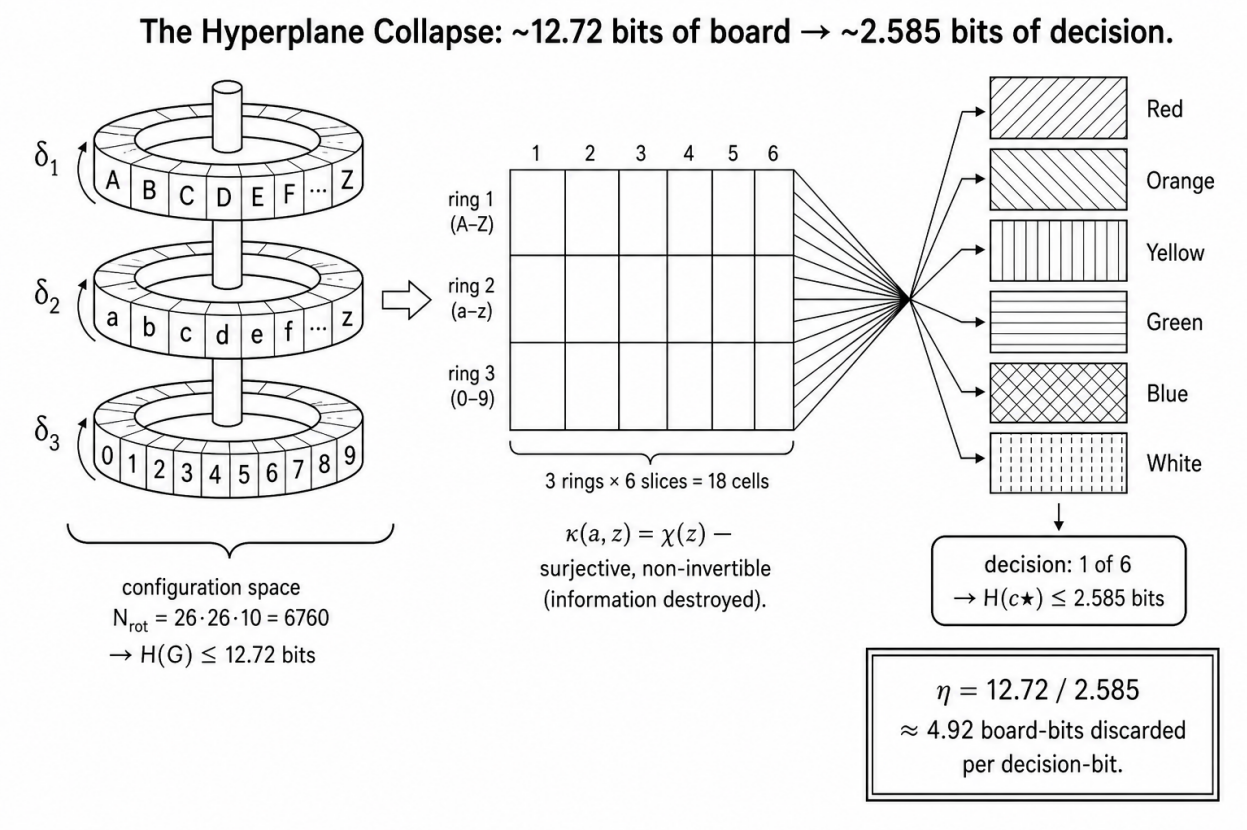

Figure 3: *Figure 3 — Rosario Cipher: Ephemeral Geometry and Hyperplane Collapse*

Figure 3. Rosario Cipher: ephemeral geometry and hyperplane collapse. Three rotated rings (uppercase, lowercase, digit alphabets) are foliated into six leaves and unioned, slice-by-slice, into six colored witness zones — a surjective, non-invertible projection that discards roughly five "bits of board" for every one "bit of decision."

The question specifically asks *how we actually collapse multiple rings and subsets of those rings, alphabet arrays, down to the hyperplane color zone.* Here is the precise dimensional accounting, because it determines how much *structure* the screen could in principle present versus how little *decision* it ultimately demands.

Why a lossy collapse is a frequency-analysis weapon. A cryptanalyst's dream is a one-to-one, structure-preserving relationship between what they observe and what they want to know; their nightmare is a *many-to-one* map that throws information away irreversibly. The hyperplane collapse described in this section is deliberately the nightmare. By unioning the same-indexed slice of three different rings into one colored zone, the construction performs a surjective, non-invertible projection: a single zone now contains a heterogeneous mixture of uppercase letters, lowercase letters, and digits, and the act of "being in the red zone" deletes all record of which ring the character came from and where on that ring it sat. This is the structural antidote to frequency analysis. Frequency analysis needs the observable categories to *track* the underlying symbol categories, it needs "e" to leave a different footprint than "q." But after the collapse, the rarest digit and the commonest letter are indistinguishable members of whatever zone the rotation

happened to drop them in this round, and the zone populations are balanced so that each of the six zones is selected with probability close to 1/6 regardless of which character is the secret. The frequency table the attacker would build over thousands of observations therefore converges to the uniform distribution: six zones, each at one-sixth, forever. There is no peak to exploit. Just as importantly, the collapse happens *before* the user makes any choice, which means the information it destroys can never be reconstructed downstream, the data-processing inequality (formalized in Section 5) guarantees that no later stage can resurrect the discarded bits. The dimensional accounting below quantifies exactly how much structure is annihilated: roughly five bits of board arrangement vaporized for every one bit of decision retained.

Configuration space before collapse. The rotated layout is determined by the joint rotation vector $(\delta_1^r, \delta_2^r, \ldots, \delta_M^r)$. The number of distinct rotational configurations is

$$N_{\text{rot}} = \prod_{a=1}^{M} n_a.$$

For the reference rings, $N_{\text{rot}} = 26 \times 26 \times 10 = 6{,}760$, so the *geometry* alone carries

$$H(G_r) \le \log_2 N_{\text{rot}} = \log_2 6760 \approx 12.72 \text{ bits}$$

of arrangement entropy per round.

Spinning three independent rings can produce 6,760 distinct board layouts; naming one of them takes about 12.7 bits. That is the maximum "richness" of the picture an observer sees. Crucially, this richness is a function of session entropy and time, not of the secret, so all 12.7 of those bits are public noise, not secret signal.

Parameter breakdown. - N_{rot}, number of distinct joint rotation states of all rings. - n_a, size of ring a; the product runs over all M rings. - $H(G_r)$, entropy of the visible geometry in round r (upper bounded by $\log_2 N_{\text{rot}}$).

The collapse map. Define the surjection that takes a (ring, slice) pair to a colored zone:

$$\kappa : \{1, \ldots, M\} \times \{0, \ldots, 5\} \longrightarrow \mathcal{C}, \qquad \kappa(a, z) = \chi(z).$$

The witness set of color c is the union of all ring-slices that κ sends to c:

$$W_c^r = \bigcup_{a\,:\,\kappa(a,z)=c} \alpha_{a,z}^r = \bigcup_{a=1}^{M} \alpha_{a,\chi^{-1}(c)}^r.$$

Every wedge from every ring is funneled into one of six colored buckets according to its slice index. The color zone forgets which ring a character came from and forgets the character's exact angular position, it remembers only "you are in the red bucket." This is the projection that collapses the three-ring, six-slice grid ($3 \times 6 = 18$ cells) down to six colored leaves.

Parameter breakdown. - κ, the collapse map from (ring, slice) cells to colors; it is *many-to-one* (surjective, not injective), which is precisely why information is destroyed. - $\chi^{-1}(c)$, the slice index assigned color c. - W_c^r, the resulting per-color witness set.

The decision-entropy after collapse. Whatever the richness of G_r, the user's actual *choice* is one of six zones. The honest answer therefore carries at most

$$H(c_r^\star) \le \log_2 |\mathcal{C}| = \log_2 6 \approx 2.585 \text{ bits.}$$

The collapse is a funnel: a 12.7-bit picture demands only a 2.585-bit decision. The screen looks complicated; the answer is one of six. This compression from ~ 12.7 bits of arrangement to ~ 2.585 bits of decision is the geometric heart of the scheme, and we are about to send even those 2.585 bits through a one-time, map-hidden channel that delivers *zero* of them to the observer.

The collapse ratio. A useful single number is the ratio of pre-collapse arrangement entropy to post-collapse decision entropy:

$$\eta = \frac{\log_2 N_{\text{rot}}}{\log_2 |\mathcal{C}|} = \frac{\log_2 6760}{\log_2 6} \approx \frac{12.72}{2.585} \approx 4.92.$$

Roughly five "bits of board" are discarded for every one "bit of choice" retained. The discarded bits are exactly the bits an observer might have hoped to exploit; the collapse throws them away before they can ever correlate with the secret.

Parameter breakdown. - η, the collapse (compression) ratio, dimensionless. - N_{rot}, $|\mathcal{C}|$, as above.

3.1 The Holographic Reading: Bulk, Boundary, and the Private Dictionary

The collapse just quantified has an illuminating interpretation that practitioners familiar with modern physics will recognize, and that gives the rest of the document a single organizing image. In physics the **holographic principle** states that all the information contained in a higher-dimensional region (the *bulk*) can be encoded on its lower-dimensional *boundary*, and that reconstructing the bulk from the boundary requires a *dictionary* relating the two descriptions. The Rosario round is holographic in exactly this sense, and naming it so makes the security argument intuitive without changing any of its content.

The correspondence is precise:

Holographic principle (physics)	Rosario interaction equivalent	Where in this document
Bulk: the high-dimensional volume of degrees of freedom	The full challenge geometry: three rings × rotation offsets × six slices × character positions × alphabet class, about 12.72 bits of arrangement ($N_{\text{rot}} = 6760$)	§3 (configuration space)
Boundary: the lower-dimensional surface that encodes the bulk	The single emitted bearing b_r (one of six), and ultimately the one accumulator bit Λ	§2, §8
Dictionary: the map relating bulk and boundary	The private synonym $f : \mathcal{C} \to \mathcal{B}$ together with the shared reconstruction of the geometry from (E, τ)	§4
Boundary is complete *with* the dictionary	The matched-twin verifier reconstructs the membership fact $s_r \in W_{c^\star}^r$ from the bearing alone	§8
Bulk is irrecoverable *without* the dictionary	An observer lacking f cannot invert the boundary token back to a zone or character, $I(C; O) = 0$	§4, §10

We can write the per-round reduction as a single boundary-encoding map:

$$\mathcal{H} : \underbrace{(s_r,\, G_r)}_{\text{bulk fact}} \;\longmapsto\; b_r \;=\; (f \circ \text{zone})(s_r, G_r),$$

where $\text{zone}(s_r, G_r) = c_r^\star$ returns the color of the zone holding the secret, and f then maps that color to a bearing.

Plain English. Everything the user actually navigated, the spinning rings, the mixed alphabets, the colored zones, is compressed into one boundary token, the bearing. That token is *sufficient* for a party who holds the dictionary (the verifier, who knows f and can regenerate the geometry) to recover the one fact that matters, namely that the secret was in the indicated zone, and *insufficient* for anyone else to recover anything at all. The interior of the challenge is sealed behind the boundary label.

Parameter / operator breakdown. - $\mathcal{H}$, the holographic boundary-encoding map for one round; takes the bulk fact to a boundary token. - (s_r, G_r), the bulk fact: the secret character together with the full round geometry it lives in. - $\text{zone}(\cdot,\cdot)$, the function that returns the color $c_r^\star$ of the zone containing s_r (the collapse of §3 applied to the secret). - f, the private synonym (the dictionary) mapping color to bearing. - $\circ$, function composition: first find the zone's color, then translate it to a bearing. - b_r, the boundary token actually emitted.

There are in fact **two nested holographic reductions**, and it is worth distinguishing them to avoid the imprecise shorthand "a single bit per character":

1. **Per round**, the bulk (≈ 12.72 bits of arrangement) is projected to a bearing carrying $\log_2 6 \approx 2.585$ *recordable* bits and, after the synonym and the mask, 0 *usable* bits.
2. **Per ceremony**, the entire boundary sequence $(b_1, \ldots, b_L)$ is further reduced by the verifier to the single membership bit Λ (§8), the accept/reject verdict. This is the literal "single bit," and it is a verdict about *work performed*, not a disclosure of the secret.

A necessary caveat for the formal reader. The holographic language here is an *explanatory analogy*, not a physics claim. The rigorous security content is unchanged: it is the surjection κ (which destroys structure), the bijection f (which is one of 720 unknown dictionaries), the compression ratio η, and the mutual-information collapse $I(C; O) = 0$ proved in Section 10. The holographic reading should be used to *picture* why a high-dimensional interior can be represented by a low-dimensional boundary that leaks nothing, and then discharged in favor of the equations whenever precision is required.

---

4. The Witness Synonym: Why the Single-Round Channel Has Zero Capacity

We have reduced one character entry to: the user knows a true color $c_r^\star$, and emits a bearing $b_r = f(c_r^\star)$. The observer sees b_r (and the geometry G_r) but not f. We now prove the central claim: **a single character entry leaks zero bits about the secret.**

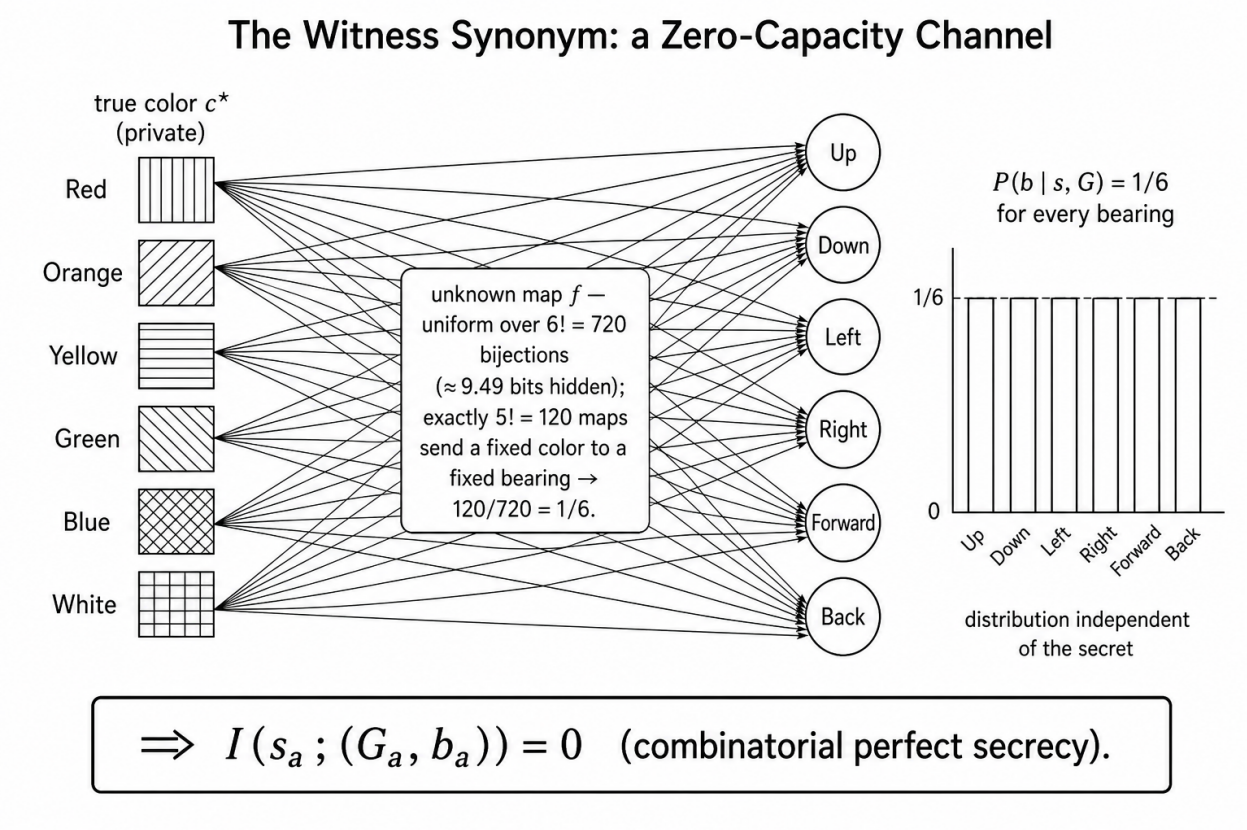

Figure 4: *Figure 4 — The synonym as a zero-capacity channel: every color smears uniformly across all six bearings*

Figure 4. Why the single-round channel has zero capacity. Because the observer is missing the synonym f (one of 720, worth $\log_2 720 \approx 9.49$ bits of "bridge"), every possible true color is spread uniformly across all six bearings: $\Pr[b_r = b \mid s_r, G_r] = \frac{1}{6}$ independent of the secret. The observed bearing histogram is therefore flat, and $I(s_r; (G_r, b_r)) = 0$.

Why the synonym is the keystone against correlation. If the collapse of Section 3 is the defense against frequency analysis, the witness synonym is the defense against *correlation between the transcript and the challenge*, the most dangerous correlation of all, because it is the one that would let an attacker align "what was on screen" with "what the user pressed" and read off the user's logic. Here is the subtle point that makes this section the heart of the analysis: even after the collapse, the user's *honest* answer (the true color $c_r^\star$) is in principle correlated with the geometry, because the secret really does live in one specific zone. If the user announced the color, an attacker who could ever learn one (geometry, secret) pair would begin to pin things down. The synonym severs that thread by translating the answer out of the challenge's language (colors) into a disjoint response language (bearings) via a private bijection the observer does not possess. The challenge speaks in colors; the transcript speaks in bearings; the only dictionary between them is the user's secret map f, and it is one of 720 equally plausible dictionaries. Because the map is unknown and uniform, every color is smeared uniformly across all six bearings, so the bearing the attacker records is statistically independent of the color the user meant, and therefore independent of the secret. This is why correlation analysis fails categorically rather than merely "by a wide margin": there is no residual statistical dependency between the transcript symbols and the challenge geometry

for the attacker to amplify across observations. The slogan "the transcript is disjoint from the challenge" is exactly the statement that this correlation has been driven to zero. We now make it a theorem.

The map is one of 720. The number of bijections between two six-element sets is

$$|\{f : \mathcal{C} \to \mathcal{B} \text{ bijective}\}| = 6! = 720.$$

There are 720 distinct, consistent ways to rename six colors as six bearings with no repeats. The user committed to one of them privately. An observer who does not know which one is missing $\log_2 720 \approx 9.49$ bits of *bridge* between the color-language of the challenge and the bearing-language of the transcript.

Parameter breakdown. - 6!, six factorial = 720, the count of bijections (permutations of six labels). - $\log_2 720 \approx 9.49$, the bits of uncertainty an observer has about the synonym.

Modeling the observer's ignorance. From the observer's standpoint, f is a random variable uniform over the 720 bijections, independent of the secret:

$$\Pr[f = \varphi] = \frac{1}{720} \quad \text{for each bijection } \varphi, \qquad f \perp (s_r, G_r).$$

Because the observer has never been told the map and the map is fixed independently of the geometry and the secret, the rational observer must treat all 720 maps as equally likely. We model this with a uniform prior on f.

The zero-information theorem (single round). Fix the geometry G_r (so the secret determines its true color $c_r^\star$). Then for every bearing $b \in \mathcal{B}$ and every secret value,

$$\Pr[\, b_r = b \mid s_r, G_r \,] \;=\; \Pr_f[\, f(c_r^\star) = b \,] \;=\; \frac{1}{6},$$

because among the 720 bijections, exactly $5! = 120$ send the fixed color $c_r^\star$ to the fixed bearing b, and $120/720 = 1/6$. Since this probability does **not depend on** s_r, the bearing is conditionally independent of the secret:

$$\Pr[\, b_r = b \mid s_r, G_r \,] = \Pr[\, b_r = b \mid G_r \,] \implies I(s_r\,;\, b_r \mid G_r) = 0.$$

And since the geometry itself is generated from session entropy and time, independently of the secret, $I(s_r; G_r) = 0$. Combining via the chain rule:

$$\boxed{I(s_r\,;\,(G_r, b_r)) = 0\,.}$$

No matter which secret character the user holds, the bearing the observer records is uniformly one of six, the map smears every possible true color uniformly across all six bearings. Knowing the bearing tells you nothing about the color, hence nothing about the secret. The geometry, being a function of fresh entropy and time, is likewise uncorrelated with the secret. So the observer's *entire* per-round record carries zero bits about the secret.

Parameter breakdown. - $\Pr[b_r = b \mid s_r, G_r]$, the chance the observer sees bearing b, conditioned on the (unknown to them) secret and the (visible) geometry. - $5! = 120$, the number of bijections fixing one specific color→bearing assignment; $120/720 = 1/6$. - $I(s_r; b_r \mid G_r)$, conditional mutual information between secret and bearing given geometry; zero means independence. - The boxed identity is the single-round instance of Shannon's perfect-secrecy condition.

This is the formal content of the slogan "the transcript is disjoint from the challenge": the challenge lives in colors, the transcript lives in bearings, and the only bridge, the synonym f, is absent from everything the observer can see. In the vocabulary of Section 11, this single-round result is **combinatorial perfect secrecy**: the response distribution is identical no matter what the secret is, which is exactly Shannon's condition $\Pr[b \mid s] = \Pr[b]$ achieved at the cognitive layer, before any encryption primitive is applied.

---

5. The XOR the Witnessing Agent Performs: A One-Time Pad on the Wire

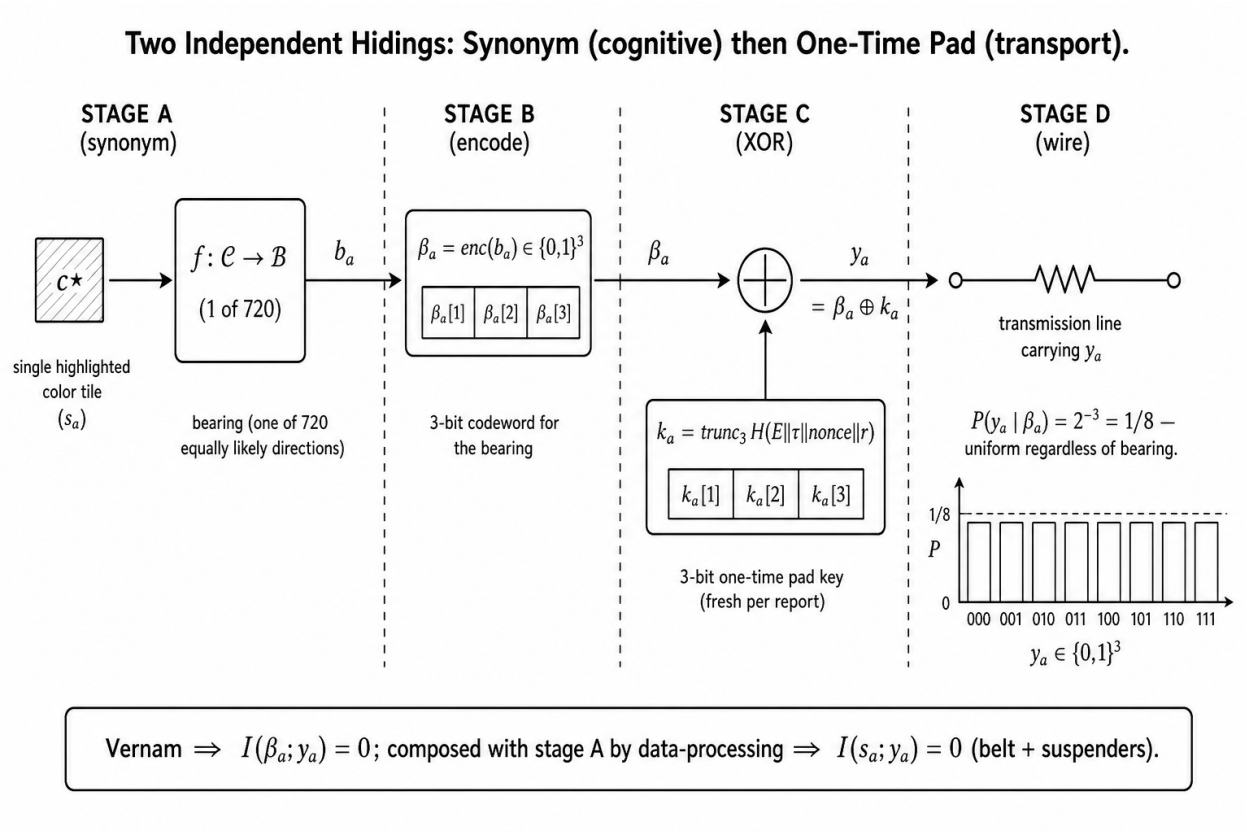

Figure 5: *Figure 5 — Rosario Cipher: Witness Synonym and Wire Masking*

Figure 5. Rosario Cipher: the witness synonym and wire masking. The cognitive map f (one of 720 bijections) translates the secret's color into a bearing, and the witnessing agent then XORs the 3-bit bearing codeword with a fresh single-use pad k_r, uniformizing the wire payload y_r by Vernam's theorem.

The question asks about *the XOR function that is performed by the agent who is providing the witnesses.* The bearing decision is a cognitive act; but what actually travels, what an on-wire

observer or interceptor records, is a **masked** payload. The witnessing agent (the prover's circuit, cooperating with the session geometry) applies a one-time XOR mask so that even the bearing index is uniformized before transmission. This is the mechanism that upgrades the combinatorial argument of Section 4 into Shannon perfect secrecy at the bit level.

The XOR mask is not an incidental engineering choice but the single most historically validated primitive in all of cryptography. It is the **one-time pad**, introduced by Gilbert Vernam in 1917 (and anticipated by Frank Miller in 1882), and it is the *only* cipher Claude Shannon proved, in 1949, to be unconditionally secure. Borrowing it here means the wire payload inherits that theorem verbatim: see Section 11 for the formal correspondence and Section 12 for why an unconditional primitive matters in a world migrating to merely *conjecturally* secure post-quantum schemes.

Why a second layer, and why XOR specifically. Sections 3 and 4 already defeat the observer who watches the *screen and the user's hand*. But a different adversary watches the *wire*, the bytes the device transmits to the verifier, and this adversary deserves an equally absolute answer. The concern is that even a "bearing" symbol, transmitted naively, has a distribution that an attacker could in principle survey for frequency structure: if some users favor *Up*, or if a particular session's encoding leaked an ordering, a patient eavesdropper might attempt a frequency or correlation attack on the raw bearing codes rather than on the secret directly. The one-time XOR pad closes this door with the strongest tool information theory offers: the Vernam cipher, the only encryption scheme ever proven to deliver perfect secrecy. By XORing each three-bit bearing code with a fresh, uniformly random, single-use pad derived from the session entropy, the witnessing agent makes the transmitted symbol a uniform draw over its range, *identically distributed no matter what the bearing was*. The frequency table of the wire payload is therefore flat by construction, every codeword appears with probability 2^{-3}, and there is nothing for frequency analysis to bite on. Correlation across rounds is likewise dead, because each round draws an independent pad, so the transmitted symbols are mutually independent even when the underlying bearings are not. The reason XOR is the right primitive (rather than, say, addition or substitution) is its group structure: XOR with a uniform key is a measure-preserving bijection on bit-strings, which is precisely the algebraic property that yields $I(\text{plaintext}; \text{ciphertext}) = 0$. This is the "belt" that complements the synonym's "suspenders": the secret is hidden once at the cognitive layer and again, independently, at the transport layer, and the two hidings *compose* by the data-processing inequality.

Bearing encoding. Index the six bearings as 3-bit words (six values fit in three bits):

$$\text{enc} : \mathcal{B} \to \{0,1\}^3, \qquad b_r \mapsto \beta_r = \text{enc}(b_r).$$

Each direction is written as a 3-bit number (000 through 101 for the six used codepoints). This is just a labeling so we can do bitwise math.

Parameter breakdown. - enc, an injective encoding of bearings into 3-bit strings. - β_r, the 3-bit codeword for the round-r bearing.

Keystream derivation. A per-round one-time mask is derived from the session secrets:

$$k_r = \text{trunc}_3\big(H(E \,\|\, \tau \,\|\, \text{nonce} \,\|\, r)\big) \in \{0,1\}^3.$$

For each round we hash the entropy pool, the timestamp, the session nonce, and the round number, then keep three bits. Because E is high-entropy and τ/nonce make every session unique, k_r behaves as a fresh, uniformly random 3-bit pad that is used exactly once.

Parameter breakdown. - k_r, the one-time 3-bit keystream chunk for round r. - trunc_3, keep three bits of the hash output. - nonce, the session's unique identifier, preventing reuse of k_r across sessions.

The masking XOR. The transmitted payload for the round is

$$y_r = \beta_r \oplus k_r,$$

where $\oplus$ is bitwise exclusive-or.

The agent XORs the bearing codeword with the one-time pad. XOR with a uniform one-time key produces a uniform output, so the bits that leave the device look like a fair coin flip per bit, regardless of which bearing was chosen.

Parameter breakdown. - y_r, the masked payload actually transmitted/recordable in round r. - β_r, the bearing codeword (the plaintext of this mini-cipher). - k_r, the one-time pad (the key). - $\oplus$, bitwise XOR.

Perfect secrecy of the wire payload. Because k_r is uniform on $\{0,1\}^3$ and independent of β_r, the output is uniform and independent of the input:

$$\Pr[\, y_r = y \mid \beta_r \,] = \Pr[\, k_r = y \oplus \beta_r \,] = 2^{-3} = \tfrac{1}{8} \quad \forall y, \qquad \Longrightarrow \qquad I(\beta_r \,;\, y_r) = 0.$$

This is Vernam's theorem: a one-time pad makes the ciphertext statistically independent of the plaintext. So the masked payload y_r reveals nothing about the bearing β_r, which (by Section 4) already revealed nothing about the secret. The two independence results compose.

Parameter breakdown. - $\Pr[y_r = y \mid \beta_r]$, distribution of the transmitted symbol given the bearing; flat at 1/8. - 2^{-3}, uniform probability over three pad bits. - $I(\beta_r; y_r) = 0$, the masked payload is independent of the bearing.

Composition (the full chain). Stacking Sections 4 and 5 by the data-processing inequality:

$$s_r \;\longrightarrow\; c_r^\star \xrightarrow{\;f\;} \beta_r \xrightarrow{\;\oplus k_r\;} y_r, \qquad I(s_r \,;\, y_r) \le \min\big(I(s_r;\beta_r),\, I(\beta_r;y_r)\big) = 0.$$

Each arrow is a stage of the pipeline. The data-processing inequality says you cannot create information about the secret by further processing; since the synonym stage already gives zero, and the XOR stage independently gives zero, the end-to-end leakage to a wire observer is zero. The XOR is the *belt* to the synonym's *suspenders*: the secret is protected at the cognitive layer (map) and again at the transport layer (pad).

Parameter breakdown. - The chain $s_r \to c_r^\star \to \beta_r \to y_r$, secret to true color to bearing codeword to masked payload. - Data-processing inequality, post-processing cannot increase mutual information.

---

6. The Observer's Per-Character Bit Budget

We can now answer the headline question for **one** character entry. The observer's complete record for round r is the triple

$$O_r = (G_r,\ b_r,\ y_r),$$

the visible geometry, the visible bearing, and the recordable masked payload. We tabulate two columns: how many bits the observer *records* (surface), and how many bits about the secret they *gain* (effective).

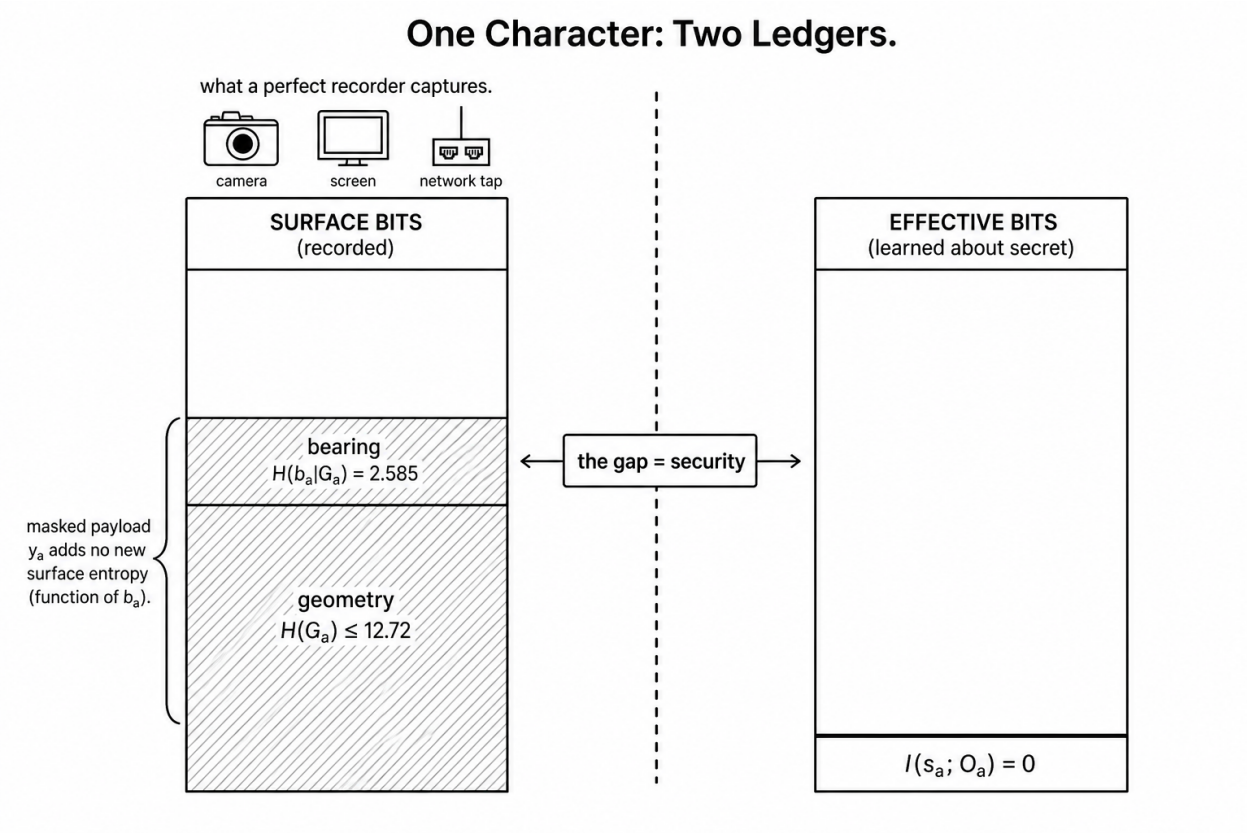

Figure 6: *Figure 6 — The per-character two-column ledger: surface bits recorded vs. effective bits learned*

Figure 6. The honest per-character accounting as two side-by-side columns. The **surface** column fills up — up to ~ 12.72 bits of board plus $\log_2 6 \approx 2.585$ bits of bearing — while the **effective** column (mutual information with the secret) stays pinned at exactly 0. Recording is not learning.

Why the two-column ledger is the honest accounting. This section is where the information-theoretic backbone becomes a single auditable number, and it is worth pausing on why that number is so often miscommunicated. Vendors who say "we leak zero bits" invite justified skepticism, because the user obviously *did* something observable, a finger moved, a symbol crossed the wire, and information theory does not let you transmit a symbol while emitting literally nothing. The resolution, made precise here, is that the positive quantity (surface entropy) and the security-relevant quantity (mutual information with the secret) are *different measures of different things*, and conflating them is the root of both the overclaim and the skepticism. Surface entropy answers "how much could a perfect recorder write down?"; mutual information answers "how much does that recording shrink the attacker's uncertainty about the secret?" The strength of an information-theoretic guarantee, as opposed to a computational one, is that the second number is *zero exactly*, not "zero unless someone finds a clever algorithm." Computational security says the secret is hidden behind a problem we believe is hard to solve; information-theoretic security says there is no problem to solve because the information is not present in the observation at all. No advance in hardware, no quantum speedup, no future cryptanalytic insight can extract a quantity that does not exist.

That is the claim this per-character ledger certifies, and it is why both frequency analysis and correlation, which are merely two algorithms for converting observations into knowledge, are not "made harder" but rendered *void*: they operate on a channel whose capacity with respect to the secret is identically zero.

Surface bits recorded per character.

$$H_{\text{surface}}(O_r) = \underbrace{H(G_r)}_{\leq 12.72} + \underbrace{H(b_r \mid G_r)}_{= \log_2 6 \approx 2.585} \quad \text{(the } y_r \text{ bits are a function of } b_r, \text{ not independent).}$$

Physically, the observer can write down the board (up to ~12.7 bits of arrangement) and the chosen bearing (~2.585 bits). If we count only the response channel, the thing a minimal shoulder-surfer actually cares about, it is $\log_2 6 \approx 2.585$ bits per character. The masked payload y_r adds no *new* surface entropy because it is determined by the bearing plus a pad the observer cannot separate.

Parameter breakdown. - O_r, the observer's full per-round record. - $H(G_r)$, geometry entropy, $\leq \log_2 6760 \approx 12.72$ bits. - $H(b_r \mid G_r) = \log_2 6 \approx 2.585$, the bearing's entropy; the operative "bits per character" of the response.

Effective bits gained per character.

$$I_{\text{effective}}(s_r\,;\, O_r) = 0 \text{ bits.}$$

By the composition of Section 4 (map hides the color) and Section 5 (pad hides the bearing), the observer's record contains zero bits about the secret character. The observer's notebook gained 2.585 bits of writing; the observer's knowledge of s_r gained nothing.

The single-character verdict.

Quantity	Symbol	Value (bits)	Meaning
Geometry arrangement	$H(G_r)$	≤ 12.72	Richness of the board (public noise)
Bearing (response channel)	$H(b_r \mid G_r)$	2.585	What a shoulder-surfer writes down
Effective leakage	$I(s_r; O_r)$	**0**	Bits of the secret obtained

Per character, the observer collects about 2.585 surface bits (one of six bearings), or up to ~15.3 bits if they also photograph the whole board, but obtains exactly 0 bits of the secret. This is the precise sense in which both "a few bits" and "zero bits" are correct answers to the challenge question: a few bits of *recording*, zero bits of *usable intelligence*.

---

7. Across the Whole Witness Ceremony

A witness ceremony is L rounds, the user enters the entire secret of length L, one character per round. We aggregate.

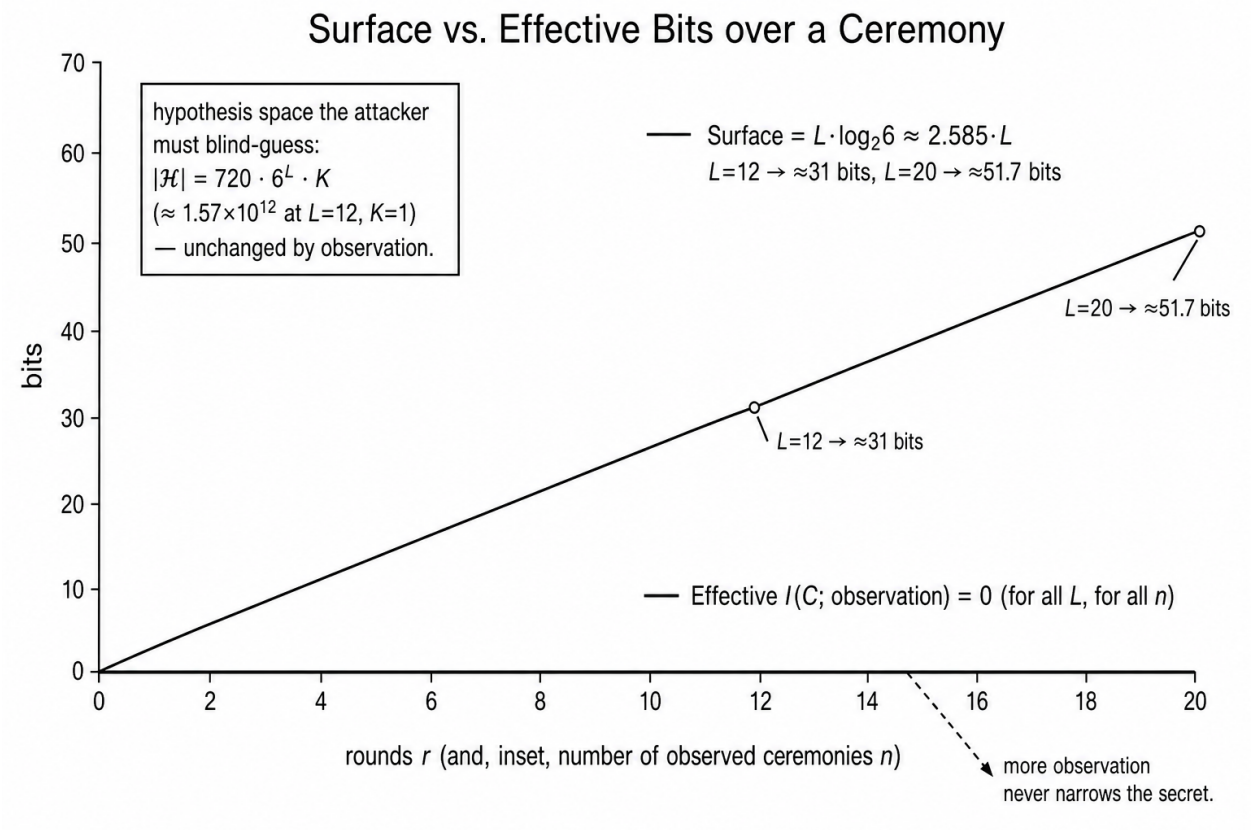

Figure 7: *Figure 7 — Across the ceremony: surface entropy grows linearly while effective leakage stays flat at zero*

Figure 7. The two ledgers over the course of a ceremony and across many sessions. The **surface** transcript grows linearly at ≈ 2.585 bits per round (the notebook gets longer); the **effective** leakage is the flat line $y = 0$ for any number of rounds L and any number of observed ceremonies n. The attacker who has watched a million sessions is in the same epistemic position as the one who has watched none.

Why aggregation is where naive schemes die, and why this one does not. The single round is rarely where real systems fail; they fail in *aggregate*, when an attacker pools many observations and lets the law of large numbers do the work that a single sample could not. This is the natural habitat of both frequency analysis (which needs a large sample to estimate a distribution) and correlation analysis (which needs many paired observations to detect a dependency). A scheme can look perfect for one round and still hemorrhage the secret over ten thousand sessions if any faint, stable signal survives averaging. The decisive question for this section is therefore not "is one round safe?", Section 4 settled that, but "does watching the *whole ceremony*, repeated across many ceremonies, ever accumulate into knowledge?" The answer is no, and the reason is structural rather than statistical. First, the per-round independence established in Section 2 means the surface transcript is a sequence of i.i.d. non-informative symbols; summing L of them grows the surface ledger linearly (the attacker's notebook gets longer) while the effective ledger stays nailed to zero (the notebook never gets more *meaningful*). Second, ENI6MA's **multiple-secret design** turns the attacker's pooled dataset into a sample from an *unidentifiable mixture*: a user authenticates with any one of several configured secrets, so even a hypothetical residual signal for a single secret

is averaged across unknown mixture weights and washed out, and different secrets can alias to identical responses. This is the specific feature that closes the theoretical gap correlation attacks rely on, there is no single stationary distribution to converge to, hence nothing for frequency or correlation estimators to lock onto. The headline consequence, derived below, is that the mutual information between the secret and the entire observed history is zero *for any number of observed ceremonies* n, so collecting more data confers no advantage whatsoever. The attacker who has watched a million sessions is in exactly the epistemic position of the attacker who has watched none.

Surface transcript entropy (linear growth). The bearings are independent fresh symbols, so

$$H_{\text{surface}}(T) = \sum_{r=1}^{L} H(b_r \mid G_r) = L \cdot \log_2 6 \approx 2.585\, L \text{ bits}, \qquad T = (b_1, \ldots, b_L).$$

The observer's transcript grows by about 2.585 bits with each additional character. After a 12-round ceremony, the observer has written down about 31 bits of bearing data; after 20 rounds, about 51.7 bits. The notebook fills linearly.

Parameter breakdown. - T, the full transcript (the bearing sequence). - L, number of rounds = secret length. - $L \cdot \log_2 6$, total surface entropy of the response channel.

Effective leakage (stays at zero). Under the design, fresh per-round geometry, the hidden synonym, the one-time XOR, and (critically) the multiple-secret mixture that ENI6MA configures so that observed sessions are samples from an unidentifiable mixture of secrets, the aggregate mutual information remains

$$I_{\text{effective}}(C\,;\, T, G_{1:L}, y_{1:L}) = 0 \text{ bits, for any number of observed ceremonies } n \geq 0.$$

No matter how many characters or how many full sessions the observer records, the bits-about-the-secret stay pinned at zero. More observation fills the notebook but never narrows the secret. Correlation analysis, frequency analysis, and machine learning all have nothing to consume, because the signal they would consume was never emitted.

Parameter breakdown. - C, the full secret commitment. - $T, G_{1:L}, y_{1:L}$, the transcript, all geometries, all masked payloads across the ceremony. - n, number of distinct ceremonies the observer has watched; the leakage is independent of n.

The attacker's hypothesis space. What the observer is *forced* into is blind enumeration over

$$|\mathcal{H}| = \underbrace{720}_{\text{maps } f} \times \underbrace{6^L}_{\text{bearing sequences}} \times \underbrace{K}_{\text{secret candidates}},$$

with single-guess success probability $1/|\mathcal{H}|$, unchanged by observation.

Because observation yields zero, the only remaining attack is to guess. The guesser must simultaneously guess one of 720 maps, one of 6^L bearing sequences, and one of K secrets. For a 12-round ceremony with $K = 1$, that is about 1.57×10^{12} equally likely possibilities, and watching the user does not eliminate even one of them.

Parameter breakdown. - $|\mathcal{H}|$, size of the (secret, map, response) hypothesis space. - $720 = 6!$, number of color–bearing maps. - 6^L, number of length-L bearing sequences. - K, number of

plausible secrets (the user may hold many; $K \geq 1$).

---

8. The Single Pass/Fail Bit

At the end of the ceremony the verifier emits exactly one bit: accept or reject. The question rightly singles this out as "the single bit that is released when pass or fail."

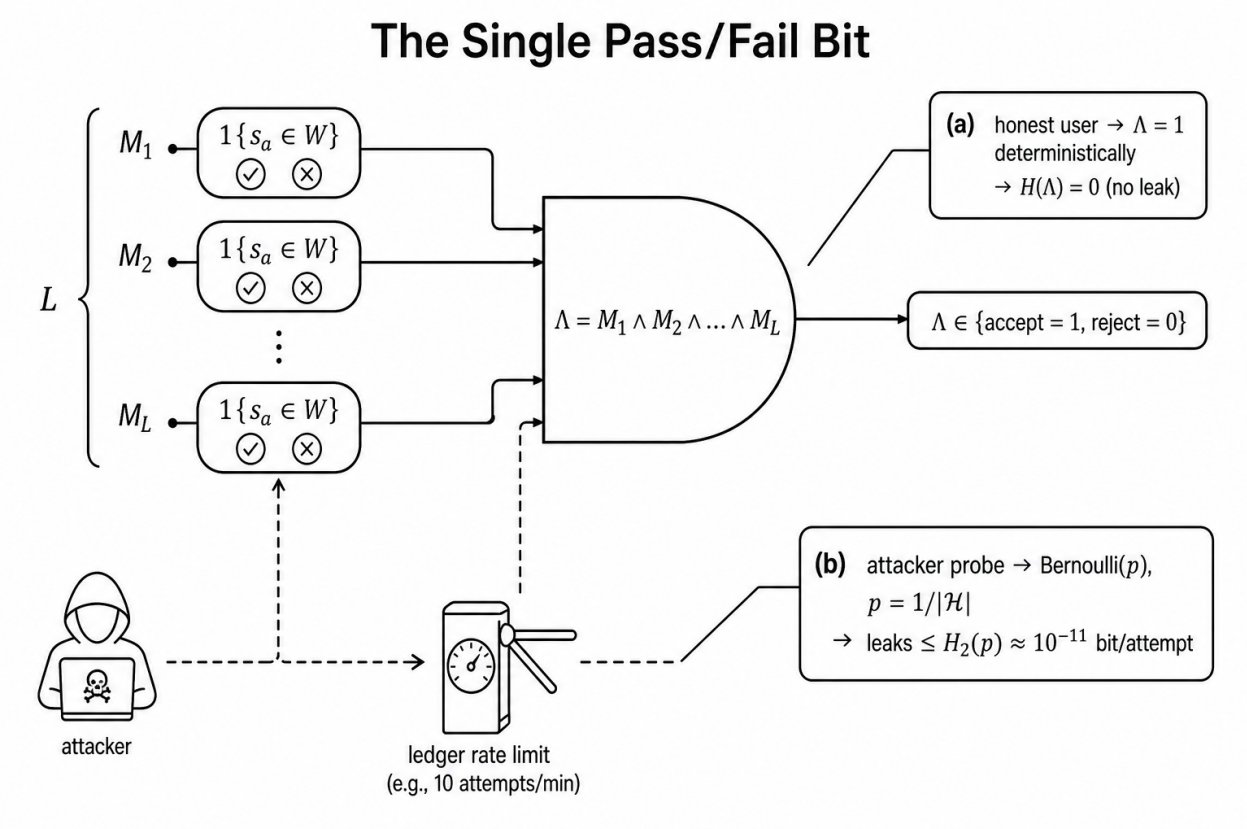

Figure 8: *Figure 8 — The accumulator: L per-round membership checks AND-ed into a single pass/fail bit*

Figure 8. The whole ceremony collapses to one bit. Each round contributes a membership indicator $M_r = \mathbf{1}\{s_r \in W^r_{z_r}\}$; the verifier ANDs them into the verdict $\Lambda = \bigwedge_r M_r$. For an honest user Λ is a foregone "accept" (entropy 0); for an attacker each probe is a Bernoulli trial leaking at most $H_2(1/|\mathcal{H}|) \approx 10^{-11}$ bits — the only surviving channel, and one the ledger rate-limits.

Why the only surviving attack is brute force, and why that is a controlled quantity. Having driven observation-based attacks to zero, we must be honest about what *does* leak, however slightly, so that the security story is complete rather than triumphal. The verifier cannot be a black hole; an authentication protocol exists precisely to produce a verdict, and a verdict is a bit. This section accounts for that bit and shows that it is the *only* channel through which an adversary can make any progress at all, and that the progress is brute force, the weakest and most easily throttled attack in the catalogue. The distinction from frequency and correlation is fundamental. Those attacks are *passive and observational*: they convert watching into knowing, and we have shown that conversion rate is zero. The pass/fail bit, by contrast, is only informative to an *active* attacker who submits a guess and learns whether it was right. For the legitimate user the verdict

is a deterministic "yes" and so carries no secret information to an eavesdropper. For an attacker, each probe is a Bernoulli trial that, on the overwhelmingly likely reject, eliminates exactly one hypothesis out of a space of size $720 \cdot 6^L \cdot K$, a vanishing fraction. This is the textbook definition of brute force: no shortcut, no statistical leverage, just sequential elimination of candidates one at a time. And brute force against an *online* verifier has a natural and brutal control: the verifier need only count attempts. The ledger's rate limit (for example, ten attempts per minute) converts the astronomically large hypothesis space into an astronomically long wall-clock time, so that the microscopic per-attempt leakage quantified below, on the order of 10^{-11} bits, can never be accumulated fast enough to matter within the lifetime of the secret, the user, or the universe. The verdict bit is thus the system's single, deliberate, and tightly metered concession to the outside world.

The membership test. Each round contributes an indicator:

$$M_r = \mathbf{1}\{\, s_r \in W^r_{z_r} \,\},$$

where z_r is the zone the verifier reconstructs from the user's (unmasked, internally) bearing via the shared map.

For each round, the verifier, who *does* hold the same map and the same geometry, translates the response back to a zone and checks whether the secret character is actually in that zone. The check returns 1 (correct) or 0 (wrong).

Parameter breakdown. - M_r, the round-r membership indicator (1 = pass). - $\mathbf{1}\{\cdot\}$, the indicator function (1 if the statement is true, else 0). - z_r, the zone the verifier attributes to the user's response in round r.

The accumulator (the one released bit). The whole session collapses to a single conjunction:

$$\Lambda = \bigwedge_{r=1}^{L} M_r \in \{0, 1\}.$$

The verifier ANDs together all L round-checks. The session passes ($\Lambda = 1$) only if *every* round was correct; one wrong round fails the whole ceremony. This single bit Λ is the only verdict released to the world.

Parameter breakdown. - Λ, the accumulator / final verdict bit. - $\bigwedge$, logical AND over all rounds. - M_r, the per-round indicators.

How much does the verdict bit leak? For an *honest* user the bit is deterministically 1, so it carries no secret information, its entropy is zero given correct play:

$$H(\Lambda \mid \text{honest user}) = 0.$$

For an *attacker's guessing attempt* the verdict is a Bernoulli trial with success probability $p = 1/|\mathcal{H}|$, and a single accept/reject leaks at most

$$I(C\,;\,\Lambda) \le H_2(p) = -p \log_2 p - (1-p)\log_2(1-p) \approx p \log_2 \tfrac{1}{p} \text{ bits}, \quad p = \tfrac{1}{|\mathcal{H}|}.$$

When the legitimate user authenticates, the "pass" bit is a foregone conclusion and therefore tells an eavesdropper nothing new. When an attacker probes by guessing, a reject (the overwhelmingly likely outcome) merely eliminates the one hypothesis they tried, a vanishing fraction $1/|\mathcal{H}|$ of the space. With $|\mathcal{H}|$ in the trillions, the verdict bit leaks on the order of 10^{-11} bits per attempt. And the ledger's rate limit caps the number of attempts, so even this microscopic dribble is throttled.

Parameter breakdown. - $H(\Lambda \mid \text{honest user})$, entropy of the verdict for a correct user (zero; deterministic pass). - $H_2(p)$, the binary entropy function evaluated at p. - $p = 1/|\mathcal{H}|$, per-attempt success probability for a blind guesser. - $I(C;\Lambda)$, mutual information between secret and the released verdict bit.

The net release. Over a full honest ceremony the *total* secret-relevant information released to any external observer is

$$\underbrace{I(C; T, G_{1:L}, y_{1:L})}_{=0 \text{ (transcript)}} + \underbrace{I(C; \Lambda \mid \text{honest})}_{=0 \text{ (verdict)}} = 0 \text{ bits.}$$

Add the transcript leakage (zero) to the verdict leakage for an honest session (zero) and the grand total of secret bits released is zero, even though the observer's notebook now holds $\sim 2.585\,L$ surface bits of bearings plus one publicly visible accept bit.

---

9. Synthesis: The Two Ledgers

The entire analysis can be summarized as two parallel ledgers that never reconcile to the same number, and that is the point. The gap between them is not a margin that a faster computer or a quantum machine could close, as it would for the computational-hardness guarantees of RSA, elliptic curves, or even the post-quantum lattice standards (Section 12); it is an *unconditional* gap, because the missing information is absent rather than merely encrypted.

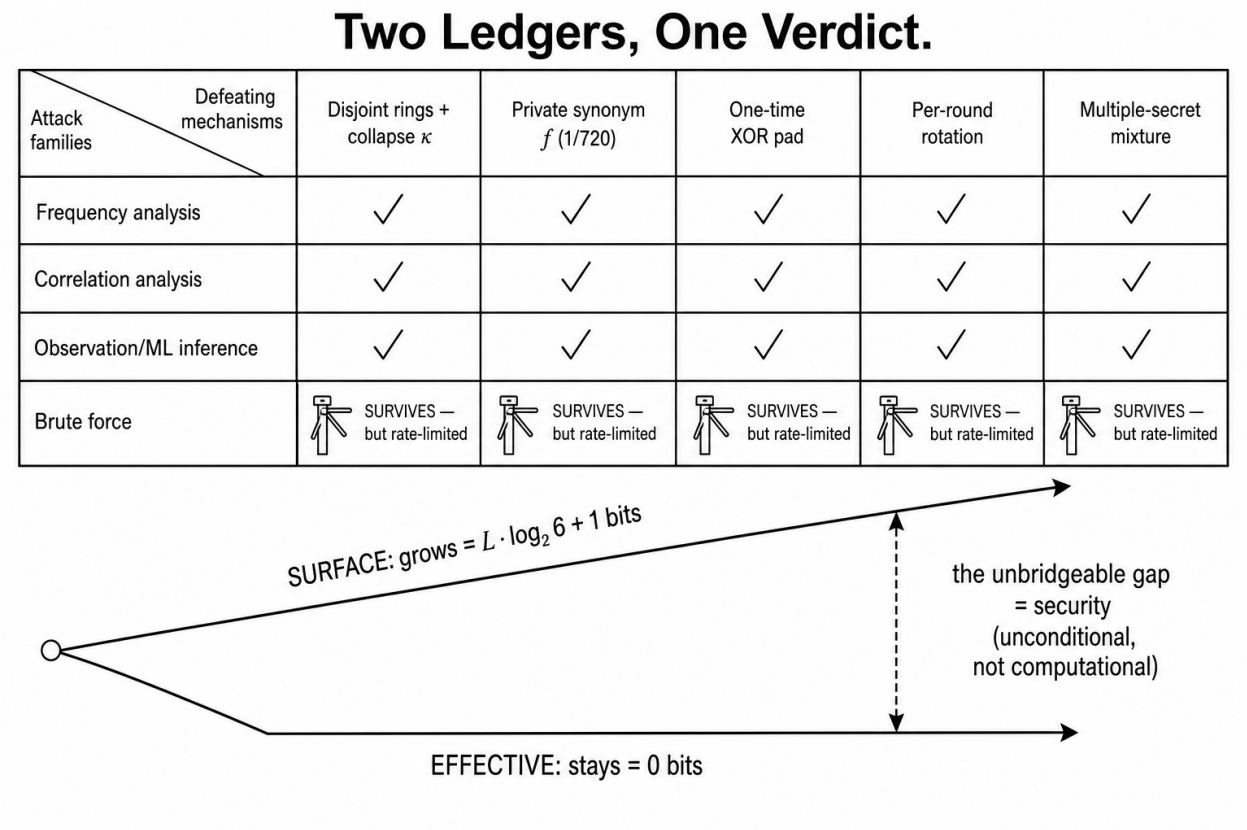

Figure 9: *Figure 9 — The attack scoreboard: three attack families reduced to throttled brute force by three mechanisms*

Figure 9. The cumulative scoreboard. Three mechanisms — the lossy **collapse**, the private **synonym**, and the one-time **XOR pad** (reinforced by the multiple-secret mixture and per-round rotation) — jointly retire frequency analysis and correlation analysis and drive observation-based inference to a structural zero, leaving only rate-limited brute force over $720 \cdot 6^L \cdot K$ candidates.

Why the two ledgers, read together, retire all three attack families. It is worth assembling the scoreboard promised in Section 0, because the power of the construction is cumulative: no single mechanism does all the work, but together they leave the adversary with no observational move and only a throttled brute-force one. Frequency analysis is defeated twice over, first by the disjoint rings and the lossy hyperplane collapse, which balance every secret into one of six equiprobable zones so that the observable distribution is flat, and again by the one-time XOR pad, which flattens the wire payload to a uniform 2^{-3} regardless of the bearing. Correlation analysis is defeated three times over, by the per-round entropy-and-time rotation, which makes successive boards independent; by the private synonym, which severs any statistical link between the bearing transcript and the colored challenge; and by the multiple-secret mixture, which dissolves any cross-session signal into an unidentifiable average. Observation-based inference of every kind collapses to the same impossibility because the mutual information between the secret and the complete observed history is identically zero, a *structural* zero that no amount of computation, classical or quantum, can mine, since the information is simply not present to be mined. What remains is brute force, and brute force is not a cryptanalytic insight but a counting problem: the attacker must traverse a hypothesis space of size $720 \cdot 6^L \cdot K$ one guess at a time, with the verdict bit as

their only feedback, against a verifier that meters their attempts. The two ledgers express this final state with a single contrast, the surface ledger grows without bound as the adversary watches, and the effective ledger never leaves zero, and that permanent, unbridgeable gap between *what can be seen* and *what can be known* is the entire security guarantee in two lines.

The surface ledger (what the observer records), per character and per ceremony:

$$\text{per char} \;=\; \log_2 6 \approx 2.585 \text{ bits (bearing)}, \qquad \text{per ceremony} \;=\; L \log_2 6{+}1 \text{ bits (transcript + verdict)}.$$

The effective ledger (what the observer learns about the secret):

$$\text{per char} = 0, \qquad \text{per ceremony} = 0, \qquad \text{after } n \text{ ceremonies} = 0.$$

The surface ledger is a positive, linearly growing count of symbols an adversary can physically capture: about 2.585 bits per keystroke, plus one verdict bit at the end. The effective ledger, the only one that matters for security, is identically zero and stays zero forever. Three independent mechanisms enforce the zero: (i) the **collapse** discards the ring/rotation arrangement before any decision is made; (ii) the **witness synonym** f, one of 720 unknown bijections, severs the bearing-language transcript from the color-language challenge; and (iii) the **one-time XOR pad** applied by the witnessing agent uniformizes the wire payload by Vernam's theorem. The pass/fail accumulator then releases a single bit that, for honest play, is a deterministic "yes" carrying no secret content.

So the challenge question that opened this document gets its honest two-part answer: *an observer can write down about two and a half bits per character, and can use exactly none of them.* For a security team, that is the line worth quoting in a threat model: **the bits an observer can capture and the bits an observer can use are not the same number, and in the Rosario interaction the second number is zero.**

---

10. A Formal Proof That the Usable Bits Are Zero

The preceding sections argued, mechanism by mechanism, that the observer learns nothing usable. This section assembles those arguments into a single self-contained proof for the formal reader: a statement of the theorem, the random variables it ranges over, four lemmas with explicit information-theoretic justification, and a conclusion in the language of Shannon's perfect secrecy. The proof is "descriptive" in the sense that every equation is paired with a plain-English reading and a full breakdown of its symbols, operators, and functions, so the argument can be followed without reconstructing the algebra from scratch.

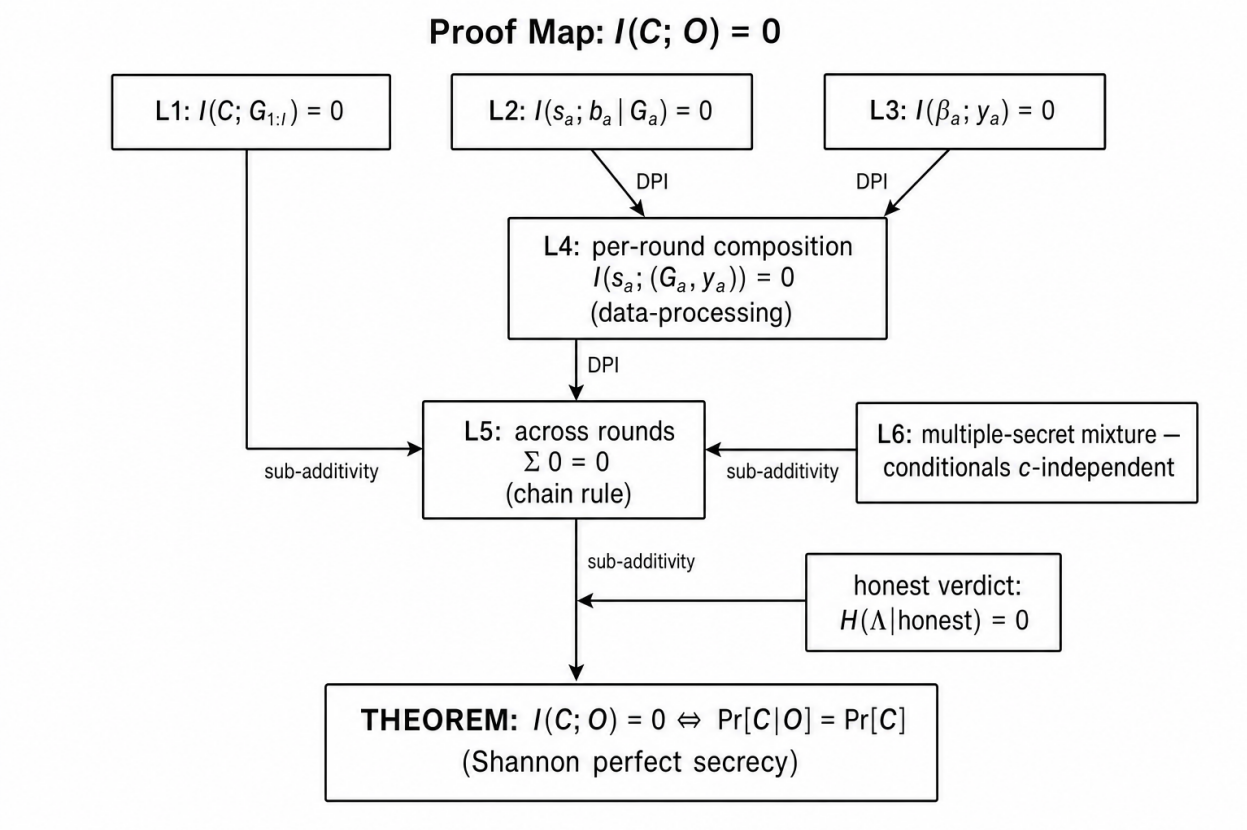

Figure 10: *Figure 10 — The six-lemma dependency graph assembling I(C;O) = 0*

Figure 10. The architecture of the proof. Six lemmas feed a single conclusion: geometry independence (L1), per-round bearing independence (L2), mask independence (L3) compose per round (L4), aggregate across rounds (L5), survive the multiple-secret mixture (L6), and combine with the honest-verdict term to force $I(C;O) = 0$ — Shannon's perfect-secrecy condition applied to the entire transcript.

10.1 The objects the proof ranges over

We first name every random variable precisely, because an information-theoretic proof is only as trustworthy as its sample space.

$$\underbrace{C}_{\text{secret}}\ ,\quad \underbrace{f}_{\text{private map}}\ ,\quad \underbrace{(E,\tau)}_{\text{session randomness}} \longrightarrow \underbrace{G_{1:L}}_{\text{geometries}} \longrightarrow \underbrace{b_{1:L}}_{\text{bearings}} \xrightarrow{\oplus k_{1:L}} \underbrace{y_{1:L}}_{\text{wire payloads}} \longrightarrow \underbrace{\Lambda}_{\text{verdict}}\ .$$

Reading left to right: the secret, the private map, and the fresh session randomness are the *inputs*; from the randomness the system derives the per-round screen geometries; the user combines geometry, secret, and map to choose a bearing each round; the witnessing agent XOR-masks each bearing into a wire payload; and at the very end the verifier emits one accept/reject bit. The observer is allowed to see everything to the right of the secret, the geometries, the wire payloads, and the verdict, but never the secret, the map, or the raw entropy.

Parameter / operator breakdown. - $C = (s_1, \dots, s_L)$, the secret commitment; a random variable over the space of possible secrets. s_r is the round-r target character. - f, the private color→bearing bijection, a random variable uniform over the 720 maps from the observer's standpoint. - (E, τ), the session entropy pool and microsecond timestamp; the source of all per-round freshness. - $G_{1:L} = (G_1, \dots, G_L)$, the sequence of round geometries; the shorthand $1{:}L$ means "indexed from 1 through L." - $b_{1:L}$, $y_{1:L}$, $k_{1:L}$, the per-round bearings, masked wire payloads, and one-time pads, respectively. - Λ, the single verdict bit. - $\oplus$, bitwise exclusive-or; the masking operator. - $\longrightarrow$, "deterministically produces" (a function arrow), so each stage is computed from the stage(s) before it.

We bundle everything the observer can capture into one symbol:

$$O \;=\; (\, G_{1:L},\; y_{1:L},\; \Lambda \,).$$

O is the observer's complete dossier after a full ceremony: every screen, every transmitted symbol, and the final verdict. The whole proof is about one quantity, the mutual information between this dossier O and the secret C.

Parameter / operator breakdown. - O, the observation; the *entire* record available to the adversary. - The comma-tuple $(\cdot, \cdot, \cdot)$, joint observation; the observer holds all three components simultaneously.

10.2 The theorem

> ***Theorem (zero usable bits).*** *For an honest ceremony and an observer who does not hold the private map f, the mutual information between the secret and the complete observation is zero:*
>
> $$I(C\,;\,O) \;=\; 0, \qquad \textit{equivalently} \qquad \Pr[\,C = c \mid O = o\,] \;=\; \Pr[\,C = c\,] \;\; \forall c, o.$$

Knowing the observer's entire dossier does not change the probability of any particular secret by even a hair: the attacker's beliefs about the secret *after* watching are identical to their beliefs *before* watching. The two formulations are equivalent, "zero mutual information" and "the posterior equals the prior" are two ways of saying the observation is irrelevant to the secret.

Parameter / operator breakdown. - $I(C;O)$, mutual information (in bits) between the secret and the observation; the average number of bits the observation reveals about the secret. Zero means statistical independence. - $\Pr[C = c \mid O = o]$, the *posterior* probability of secret value c given that the observation took the specific value o. - $\Pr[C = c]$, the *prior* probability of c (before any observation). - $\forall c, o$, "for all" secret values c and all observation values o; the equality is universal, not merely on average.

We prove the mutual-information form; the posterior form follows because $I(C;O) = 0$ is *definitionally* equivalent to $C \perp O$ (independence), which is the statement that posterior equals prior.

10.3 Lemma 1, The geometry is secret-independent

$$I(C\,;\,G_{1:L}) = 0.$$

The sequence of screens the user is shown is produced entirely from fresh session entropy and the clock, with no input from the secret. Therefore watching the screens, no matter how many rounds, tells you nothing about the secret.

Parameter / operator breakdown. - $G_{1:L}$, all L geometries. - The justification: each G_r is a deterministic function of (E, τ, r) only (Equation 6, the rotation operator). Since (E, τ) is drawn independently of C, the function output is independent of C, and a function of a variable independent of C is itself independent of C.

10.4 Lemma 2, Per round, the bearing is secret-independent given the geometry

$$\Pr[\, b_r = b \mid s_r,\, G_r \,] \;=\; \tfrac{1}{6} \;=\; \Pr[\, b_r = b \mid G_r \,] \qquad \Longrightarrow \qquad I(s_r\,;\, b_r \mid G_r) = 0.$$

Fix the screen. Whatever the secret character is, the bearing the user emits is equally likely to be any of the six directions, because the unknown private map spreads every possible "true color" uniformly across all six bearings. So the bearing, conditioned on the screen, carries no information about which character was the target.

Parameter / operator breakdown. - $\Pr[b_r = b \mid s_r, G_r]$, probability the round-$r$ bearing equals b, given the secret character and the geometry. - $\frac{1}{6}$, the uniform value; arises because exactly $5! = 120$ of the 720 maps send a fixed color to a fixed bearing, and $120/720 = 1/6$ (Equation 19). - $I(s_r; b_r \mid G_r)$, conditional mutual information between the secret character and the bearing, *given* the geometry. The implication arrow $\Longrightarrow$ holds because equality of the conditional and marginal distributions is the definition of conditional independence. - This lemma is the formal core; it is the witness-synonym mechanism of Section 4 expressed as a conditional-independence statement.

10.5 Lemma 3, The mask makes the wire payload bearing-independent

$$\Pr[\, y_r = y \mid b_r \,] = 2^{-3} \qquad \Longrightarrow \qquad I(b_r\,;\, y_r) = 0.$$

What actually travels on the wire is the bearing XORed with a fresh one-time pad. Because the pad is uniform and used once, the transmitted symbol is uniform regardless of the bearing, so even the wire payload reveals nothing about the bearing (and therefore nothing about the secret).

Parameter / operator breakdown. - $\Pr[y_r = y \mid b_r]$, probability the transmitted symbol is y, given the bearing. - $2^{-3} = 1/8$, uniform over the three pad bits (Equation 24); the encoding enc places each bearing in $\{0,1\}^3$. - $I(b_r; y_r)$, mutual information between bearing and wire payload; zero by Vernam's one-time-pad theorem.

10.6 Lemma 4, Per-round composition: the full round observation is secret-independent

Combining Lemmas 1–3 through the data-processing inequality, the round-r observation (G_r, y_r) is independent of the secret character:

$$I(s_r\,;\,(G_r, y_r)) \;\le\; I(s_r\,;\,(G_r, b_r)) \;=\; 0.$$

The wire payload is computed *from* the bearing by a process (XOR masking) that uses no extra knowledge of the secret. The data-processing inequality says you cannot manufacture information about the secret by post-processing something that already had none. Since the geometry-plus-bearing pair carried zero (Lemmas 1–2), the geometry-plus-payload pair carries zero too.

Parameter / operator breakdown. - (G_r, y_r), what the observer actually keeps for round r (screen plus transmitted symbol). - $\le$, the data-processing inequality: post-processing $b_r \to y_r$ can

only decrease or preserve mutual information, never increase it. - The right-hand = 0 comes from Lemma 2 (and Lemma 1 for the geometry component).

10.7 Lemma 5, Across rounds: independence aggregates rather than accumulates

Because each round draws fresh geometry and a fresh pad, the per-round observations are conditionally independent, and the chain rule for mutual information gives

$$I(C\,;\,G_{1:L}, y_{1:L}) \;=\; \sum_{r=1}^{L} I(s_r\,;\,(G_r, y_r)\,\big|\,\text{past rounds}) \;=\; \sum_{r=1}^{L} 0 \;=\; 0.$$

Watching more rounds is just adding up zeros. Each additional round is an independent, freshly randomized experiment that, by Lemma 4, leaks nothing, and there is no cross-round residue for the sum to accumulate. This is precisely why collecting a longer transcript (or many transcripts) never converts into knowledge: a sum of zeros is zero no matter how many terms.

Parameter / operator breakdown. - $\sum_{r=1}^{L}$, summation over all L rounds; the chain rule decomposes the joint mutual information into a sum of per-round contributions. - $I(\cdot \mid \text{past rounds})$, the information a round adds *given everything already observed*; the conditioning is what makes the chain rule exact rather than an inequality. - Each summand is 0 by Lemma 4 together with the per-round freshness (Lemma 1), which removes any dependence on earlier rounds.

10.8 Lemma 6, The multiple-secret mixture forecloses cross-session inference

A real deployment lets the user authenticate with any one of several configured secrets. An attacker pooling n sessions therefore observes samples from a *mixture* whose component weights are themselves unknown:

$$\Pr[\,O = o\,] \;=\; \sum_{c} \pi_c \,\Pr[\,O = o \mid C = c\,], \qquad \text{with } \Pr[O = o \mid C = c] \text{ independent of } c \text{ by Lemmas 1–5.}$$

Even if an attacker could somehow accumulate signal for a single fixed secret across thousands of logins, the user is not using a single fixed secret, they are drawing from a set of them with unknown frequencies. But the deeper point is that each component of the mixture is *already* non-informative (Lemmas 1–5), so the mixture is a blend of identical, secret-independent distributions. A weighted average of identical distributions is that same distribution, leaving nothing for frequency or correlation analysis to separate.

Parameter / operator breakdown. - π_c, the (unknown) mixture weight: how often the user authenticates with secret c. - $\Pr[O = o \mid C = c]$, the observation distribution conditioned on using secret c; by Lemmas 1–5 this does not depend on c. - $\sum_c$, summation over all secrets in the user's set. - Because every conditional term is the same function of o, the marginal $\Pr[O = o]$ equals that common term, so observing O cannot distinguish which c (or which weights π_c) generated it.

10.9 Handling the verdict bit

The observation O includes the verdict Λ. For an honest ceremony the verdict is the constant 1:

$$\Pr[\,\Lambda = 1 \mid \text{honest}\,] = 1 \qquad \Longrightarrow \qquad H(\Lambda \mid \text{honest}) = 0 \qquad \Longrightarrow \qquad I(C\,;\,\Lambda \mid \text{honest}) = 0.$$

When the legitimate user authenticates, the answer is always "accept." A variable that is always the same value carries no information, its entropy is zero, so it cannot reveal anything about the secret. (An *active* attacker who submits guesses can extract the microscopic $H_2(1/|\mathcal{H}|)$ of Section 8 per attempt, but that is brute force against the verifier, not inference from observation, and the ledger rate-limits it.)

Parameter / operator breakdown. - $H(\Lambda \mid \text{honest})$, entropy of the verdict for an honest session; zero because the outcome is certain. - $I(C;\Lambda \mid \text{honest})$, information the verdict carries about the secret in an honest session; zero, since a constant is independent of everything.

10.10 Conclusion (QED)

Assembling the lemmas: the geometry component contributes zero (Lemma 1), each round's payload contributes zero and the rounds aggregate to zero (Lemmas 2–5), the multiple-secret mixture cannot be unmixed (Lemma 6), and the verdict contributes zero for an honest ceremony (10.9). By sub-additivity of mutual information over the components of O,

$$I(C\,;\,O) \;\le\; I(C\,;\,G_{1:L}, y_{1:L}) \;+\; I(C\,;\,\Lambda \mid G_{1:L}, y_{1:L}) \;=\; 0 + 0 \;=\; 0.$$

Since mutual information is non-negative, $I(C;O) \ge 0$, and we have just shown $I(C;O) \le 0$; therefore

$$\boxed{I(C\,;\,O) = 0} \qquad \Longleftrightarrow \qquad \Pr[\,C \mid O\,] = \Pr[\,C\,].$$

The total information the observer's complete dossier carries about the secret is squeezed between zero (from below, because information cannot be negative) and zero (from above, because every component contributes nothing). It is therefore exactly zero. Equivalently, the attacker's posterior belief about the secret equals their prior, observation changed nothing. This is Shannon's perfect-secrecy condition, applied not to a ciphertext but to the entire observable footprint of the authentication ceremony.

Parameter / operator breakdown. - The inequality step uses *sub-additivity* / the chain rule: the information in the joint observation is at most the information in its parts taken in sequence. - $I(C;\Lambda \mid G_{1:L}, y_{1:L})$, the *extra* information the verdict adds once the geometries and payloads are already known; zero for honest play by 10.9. - ≥ 0, non-negativity of mutual information, a basic property; combined with the upper bound of 0 it forces equality. - $\Longleftrightarrow$, logical equivalence between the mutual-information statement and the posterior-equals-prior statement.

Operational restatement for the security reader. The dossier an adversary assembles from a perfect camera, a perfect screen recorder, a perfect network tap, and unlimited offline compute has *zero* mutual information with the credential. The usable-bit count is not "small," "computationally hard to extract," or "hard for now", it is **zero, structurally and permanently**, because the information is absent from the recording rather than merely obscured within it. Every observation-based attack, shoulder-surfing, replay, frequency analysis, correlation across sessions, machine-learning inference, future quantum analysis, therefore has the same success probability as a blind guess, and the only remaining adversarial action is rate-limited brute force against the verifier.

11. Relationship to Shannon Perfect Secrecy

The proof of Section 10 is, at bottom, a derivation that the authentication ceremony satisfies the oldest and strongest standard in cryptography. This section makes that connection explicit, because the synthesis is otherwise spread implicitly across the synonym (§4), the XOR (§5), and the formal proof (§10), and a reviewer deserves to see it stated in one place.

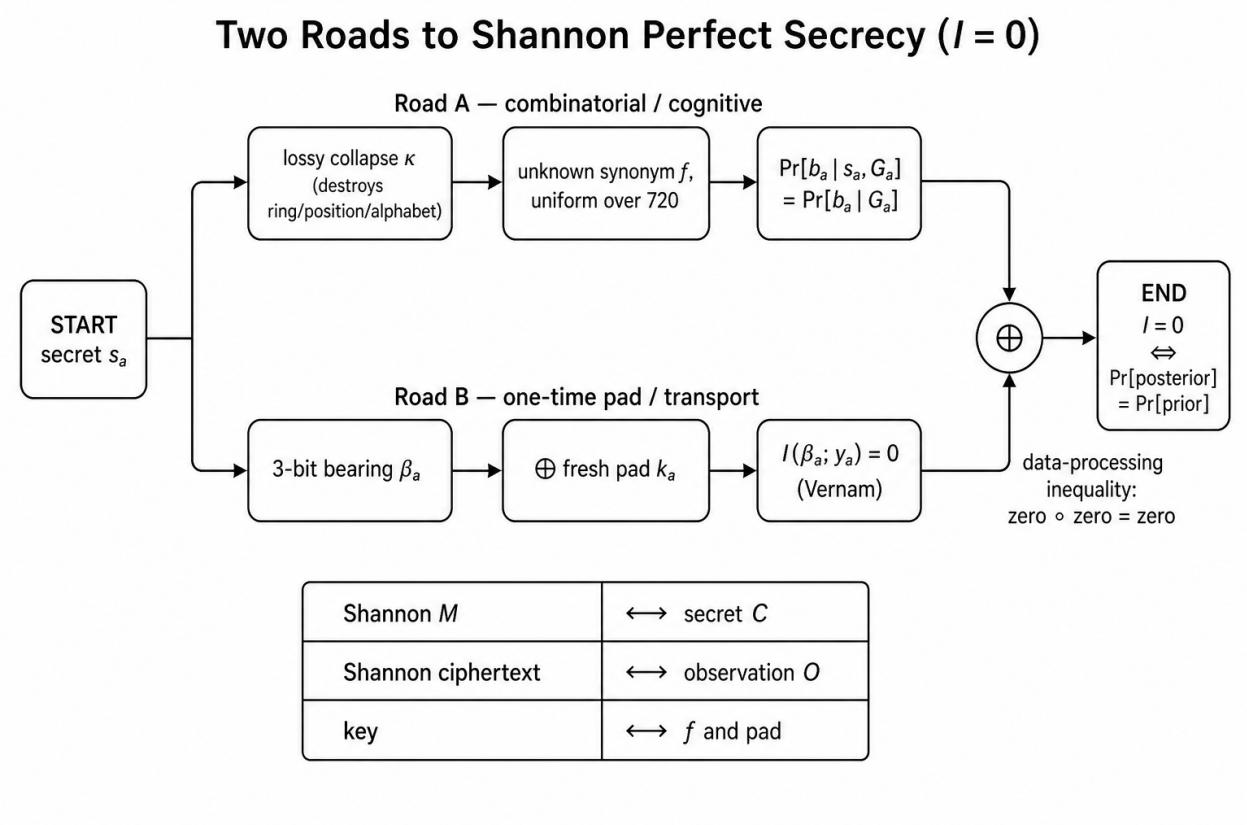

Figure 11: *Figure 11 — Two independent roads to $I = 0$, both satisfying Shannon's perfect-secrecy condition*

Figure 11. Two independent roads arrive at the same destination, $I = 0$. **Road A** (cognitive layer) reaches $\Pr[b_r \mid s_r, G_r] = \Pr[b_r \mid G_r]$ via the lossy collapse and the unknown synonym; **Road B** (transport layer) reaches $I(\beta_r; y_r) = 0$ via the Vernam one-time pad. Either alone suffices; composed by the data-processing inequality they yield end-to-end zero — the same condition Shannon defined in 1949, here applied to a whole interactive transcript rather than a single ciphertext.

11.1 Shannon's definition, restated

In his 1949 paper *Communication Theory of Secrecy Systems*, Claude Shannon defined a cipher to be **perfectly secret** when the ciphertext gives an adversary no information about the plaintext:

$$\Pr[\,M = m \mid C = c\,] \;=\; \Pr[\,M = m\,] \qquad \Longleftrightarrow \qquad I(M;C) = 0.$$

Plain English. A scheme is perfectly secret when seeing the ciphertext leaves the attacker's belief about the message exactly where it started. The two sides of the equivalence say the same thing in different dialects: "the posterior equals the prior" (probability language) and "the mutual information is zero" (information-theory language). Perfect secrecy is the gold standard precisely because it is *unconditional*, it makes no assumption that the attacker is computationally bounded.

Parameter / operator breakdown. - M, the plaintext message (the secret), a random variable. - C, the ciphertext, the signal the adversary observes. (Note: this is Shannon's C; the present document reuses the letter C for the *secret commitment*, so we map the names carefully in §11.2.) - $\Pr[M = m \mid C = c]$, the *posterior* probability of message m after observing ciphertext c. - $\Pr[M = m]$, the *prior* probability of m before any observation. - $I(M;C)$, mutual information in bits; zero denotes statistical independence. - $\iff$, logical equivalence between the two formulations.

11.2 The message / signal / key correspondence

Because the document reuses C for the secret, the mapping to Shannon's notation must be stated to avoid confusion:

Shannon (1949)	This document	What it is
Message M	Secret commitment $C = (s_1, \ldots, s_L)$	The thing that must stay hidden
Ciphertext / signal C	Observation $O = (G_{1:L},\, y_{1:L},\, \Lambda)$	Everything the observer can capture: screens, wire payloads, verdict
Key	Private synonym f together with the session pad material (E, τ, k_r)	The shared secret that renders the signal non-informative

So the phrase "zero mutual information between the message and the signal" is, in this document's symbols, exactly the Section 10 result $I(C;O) = 0$, with message = the secret and signal = the full observable footprint.

11.3 Two independent roads to $I = 0$

The construction reaches Shannon's condition by **two separate routes**, either of which alone would suffice, composed by the data-processing inequality (§10.6):

- **Road A, combinatorial perfect secrecy (cognitive layer).** Even with no encryption applied, the per-round channel already has zero capacity. The lossy collapse κ (§3) destroys the ring, position, and alphabet structure, and the unknown synonym f, uniform over $6! = 720$ dictionaries, smears every true color uniformly across the six bearings, giving $\Pr[b_r \mid s_r, G_r] = \frac{1}{6} = \Pr[b_r \mid G_r]$ (Lemma 2). This *is* Shannon's $\Pr[b \mid s] = \Pr[b]$, achieved by geometry and a private permutation rather than by a cipher.
- **Road B, the one-time pad (transport layer).** The bearing is XORed with a fresh, single-use, uniform pad, $y_r = \beta_r \oplus k_r$ (§5). This is the Vernam cipher, the canonical and unique realization of perfect secrecy, giving $I(\beta_r; y_r) = 0$ (Lemma 3).

Plain English. Road A hides the secret behind a dictionary the observer does not have; Road B hides the (already meaningless) bearing behind a pad the observer does not have. Because post-processing cannot manufacture information, stacking a zero on a zero yields zero, so the end-to-end leakage to any passive observer is zero. The pad is defense-in-depth on top of an already-perfect cognitive channel.

11.4 How the construction exceeds classical perfect secrecy

Three properties make the claim stronger than "we used a one-time pad," and they are worth stating for a cryptographer:

1. **The secret is never enciphered or transmitted.** Classical perfect secrecy protects a message that is *sent*. Here the secret is only ever a private *selector*; the signal O is a *witness to knowledge*, not an encryption of the secret. There is no ciphertext-of-the-secret to attack at all, only a proof that the membership relation $s_r \in W^r_{c^\star}$ holds.
2. **Perfect secrecy of an entire interactive transcript, not one ciphertext.** Shannon's theorem is usually stated for a single message under a single key. Section 10 extends $I = 0$ across all L rounds (the chain rule, a sum of zeros) and across all n sessions (the multiple-secret mixture). The property is therefore *time-extended and replay-immune*, which a static one-time pad does not provide on its own.
3. **No long-key burden.** The one-time pad's notorious cost is a key as long as the message, perfectly shared and never reused. Here the heavy lifting against the observer is done by Road A, whose "key" is a memorized 1-of-720 map plus session-fresh entropy; the pad is short and per-round.

11.5 The one honest boundary

Shannon's $I(M;C) = 0$ describes a **passive** observer of the signal. The single place anything leaks is the **verdict bit to an active attacker** who submits a guess and learns accept or reject; that channel is bounded by $H_2(1/|\mathcal{H}|)$ per attempt (about 10^{-11} bits for realistic parameters) and is rate-limited by the ledger (§8). For an honest ceremony the verdict is a constant and contributes zero (§10.9). So the perfect-secrecy claim concerns observation and interception; the residual against online guessing is deliberately metered brute force, not a leak in the Shannon sense.

---

12. Historical Context and Standing Against Modern and Post-Quantum Cryptography

A security architect evaluating this construction needs to know not only *that* it is perfectly secret but *where it stands* in the long arc of cryptography and against the schemes now being standardized for the quantum era. This section provides that context.

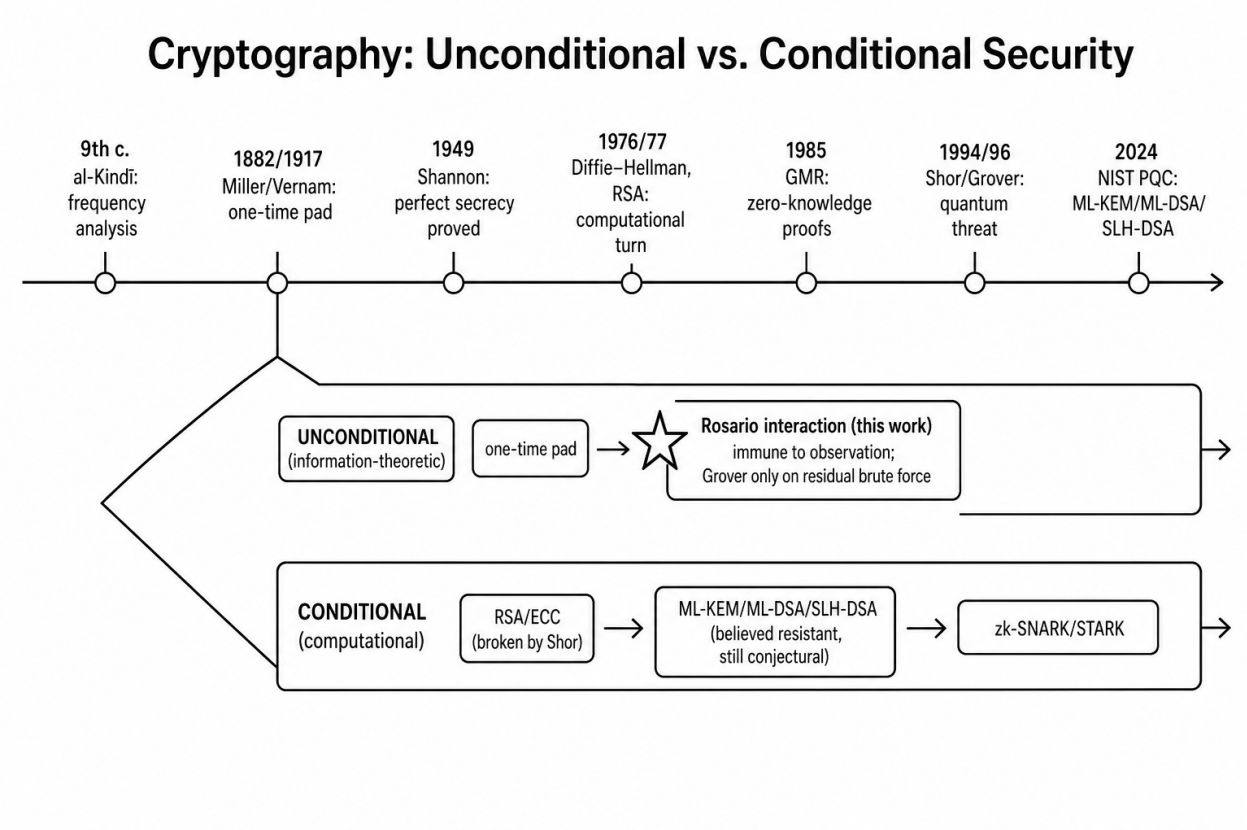

Figure 12: *Figure 12 — A timeline of cryptography and where the Rosario interaction sits on the unconditional branch*

Figure 12. The long arc of the message-versus-signal problem, and the fork between *unconditional* (information-theoretic) and *conditional* (computational) security. The construction analyzed here sits on the unconditional branch with the one-time pad — immune to observation-based and quantum attack on its transcript — while RSA/ECC and the NIST post-quantum standards (ML-KEM, ML-DSA, SLH-DSA) sit on the conditional branch.

12.1 A compressed history of the message-versus-signal problem

The discipline has oscillated between two kinds of guarantee for over a thousand years.

- **Frequency analysis and the fall of substitution (9th century).** The Arab polymath al-Kindi described, in the 9th century, how the non-uniform frequency of letters betrays a substitution cipher. Every classical cipher that maps symbols to symbols inherits this weakness: the signal carries the statistical fingerprint of the message. The collapse and synonym of this document are, in a direct sense, the deliberate destruction of that fingerprint (Sections 1, 3).
- **The one-time pad and unconditional secrecy (1882, 1917, 1949).** Frank Miller (1882) and Gilbert Vernam (1917) introduced the one-time pad; Shannon (1949) proved it perfectly secret and proved that perfect secrecy *requires* key material as plentiful as the message. This is the first and, until now, essentially the only branch of unconditional security, and it is the branch this document belongs to.
- **The computational turn (1976, 1977).** Diffie and Hellman (1976) and Rivest, Shamir,

and Adleman (1977) launched modern cryptography by trading *unconditional* security for *practical* security: keys became short, and safety rested on the conjectured hardness of problems like integer factoring and discrete logarithms. Nearly all deployed encryption, key exchange, and signatures live here.

- **Interactive and zero-knowledge proofs (1985).** Goldwasser, Micali, and Rackoff formalized interactive proofs and zero-knowledge, where a prover convinces a verifier of a fact while revealing nothing else. The Rosario ceremony is an interactive proof of knowledge in this lineage, but with an information-theoretically non-informative transcript rather than a merely computationally hiding one.
- **The quantum threat and the post-quantum response (1994 onward).** Shor's algorithm (1994) breaks factoring and discrete logarithms on a sufficiently large quantum computer, and Grover's algorithm (1996) gives a square-root speedup against unstructured search. This drove the NIST post-quantum standardization that concluded its first standards in 2024.

12.2 Where modern and post-quantum schemes draw their security

The crucial observation for this document is that **the post-quantum migration does not change the *kind* of guarantee; it only changes the hard problem.**

- **Classical (RSA, ECC):** security from the conjectured hardness of factoring and discrete logarithms. Broken in principle by Shor on a cryptographically relevant quantum computer.
- **Post-quantum NIST standards: ML-KEM** (the Kyber-derived key-encapsulation mechanism, FIPS 203), **ML-DSA** (the Dilithium-derived signature scheme, FIPS 204), and **SLH-DSA** (the SPHINCS+ hash-based signature scheme, FIPS 205). Their security rests on the conjectured hardness of structured lattice problems (Module-LWE and Module-SIS) or of hash-function properties. These are believed to resist quantum attack, but the guarantee remains **computational and conjectural**: it assumes the adversary cannot solve the underlying problem, and a future algorithmic advance could in principle weaken it.
- **Succinct proof systems (zk-SNARKs, zk-STARKs, Bulletproofs):** soundness and zero-knowledge that are typically computational, and sometimes dependent on a trusted setup. Powerful for verifiable computation, but not unconditional in the Shannon sense.

12.3 Where the Rosario interaction sits

Against this backdrop the construction analyzed here occupies the *unconditional* branch, and its quantum stance follows directly. Because $I(C;O) = 0$, the secret-relevant information is simply absent from the observation, so there is no structure for an algorithm to exploit. The most a quantum adversary can do against the residual brute-force channel is Grover search over the hypothesis space, a square-root speedup:

$$T_{\text{quantum}} \;\approx\; \sqrt{|\mathcal{H}|} \;=\; \sqrt{720 \cdot 6^{L} \cdot K},$$

which is still exponential in the round count L and is, in any case, throttled by the verifier's rate limit.

Plain English. Even a perfect quantum computer cannot read information that was never emitted; the best it can do is speed up blind guessing, and only by taking the square root of the search space, which leaves an astronomically large and rate-limited problem intact.

Parameter / operator breakdown. - T_{quantum}, the order of magnitude of guessing effort for a Grover-style search. - $|\mathcal{H}| = 720 \cdot 6^L \cdot K$, the hypothesis-space size from §7 (maps × bearing sequences × secret candidates). - $\sqrt{\cdot}$, the square-root (Grover) speedup, the *only* general quantum advantage against unstructured search. - L, the number of rounds; K, the number of candidate secrets.

A comparison table makes the standing concrete:

Scheme	Security basis	Stance vs. quantum	What it protects
One-time pad	Information-theoretic (unconditional)	Immune (no structure to attack)	Confidentiality of one message
RSA / ECC	Computational (factoring, discrete log)	Broken by Shor	Encryption, key exchange, signatures
ML-KEM / ML-DSA / SLH-DSA	Computational (lattice / hash, conjectural)	Believed resistant, still conditional	Key encapsulation, signatures
zk-SNARK / zk-STARK	Computational soundness (often + setup)	Varies; mostly conditional	Verifiable computation, proofs
Rosario interaction (this document)	Information-theoretic (unconditional)	Immune for observation; Grover only on residual brute force	Human-interactive proof of knowledge / authentication

12.4 An honest scoping note

It would be an overclaim to say this construction *replaces* Kyber or Dilithium, and a careful reader should resist that reading. The post-quantum standards solve problems this scheme does not address: encapsulating keys for bulk data encryption, and producing transferable digital signatures over arbitrary artifacts. The Rosario interaction solves a *different* problem, proving knowledge of a human-held secret across an interactive ceremony whose transcript leaks nothing, and it does so with an unconditional guarantee that the computational schemes cannot offer for their own problem. The correct framing is **complementary**: ML-KEM and ML-DSA continue to protect data at rest and in transit and to sign artifacts, while the construction here removes the credential itself, and the entire observation channel around it, from the attack surface. Defense in depth gains an unconditional layer exactly where the human meets the machine, which is historically the weakest link and the one least helped by stronger hard-problem assumptions.

---

13. Conclusion

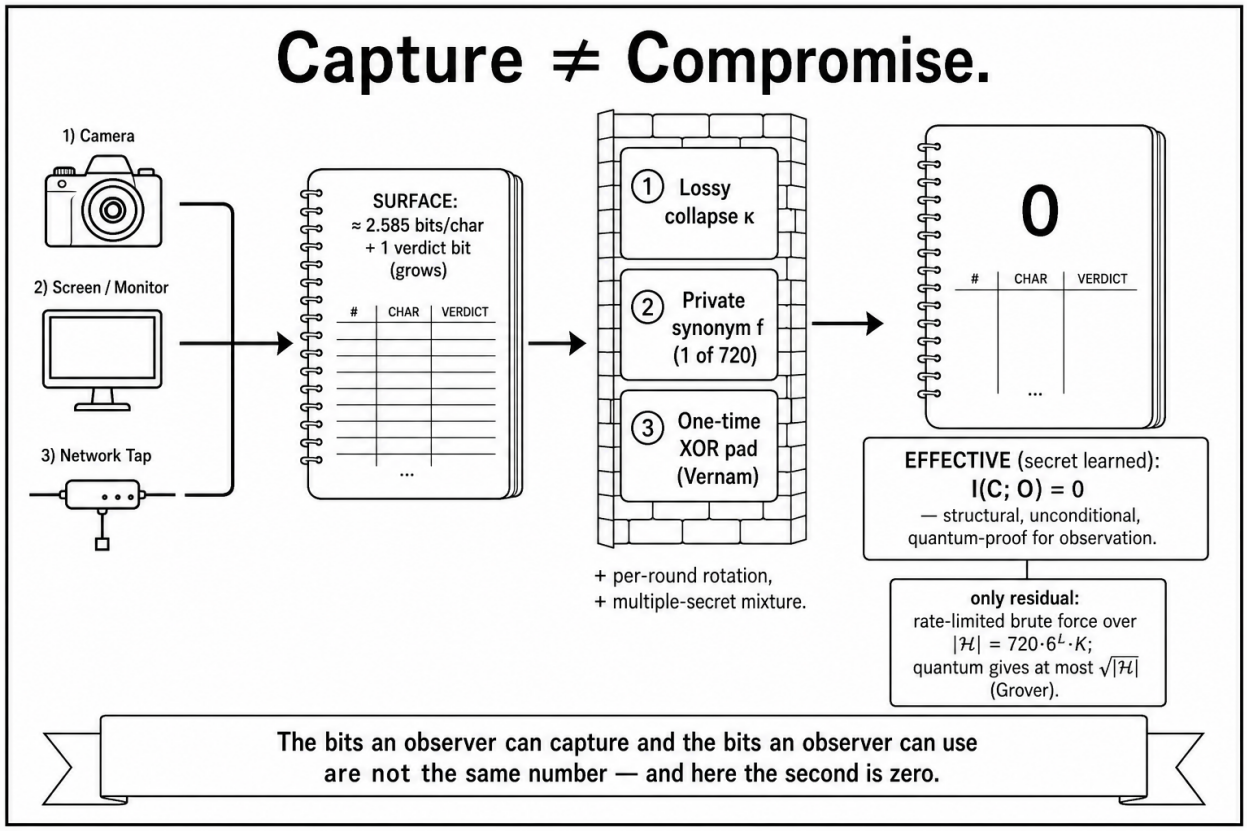

Figure 13: *Figure 13 — Capstone: capture is not compromise — three mechanisms hold the effective ledger at zero*

Figure 13. The one-page takeaway. An observer with a perfect camera, screen recorder, and network tap fills a surface ledger at ≈ 2.585 bits per character plus one verdict bit, yet the three composable mechanisms — lossy **collapse**, private **synonym** (1 of 720), and one-time **XOR pad** — hold the effective (secret-relevant) ledger at exactly 0, leaving only rate-limited brute force against $720 \cdot 6^L \cdot K$, with at most a square-root Grover speedup.

We have given a complete information-theoretic accounting of the observer's bit budget in a full Rosario cipher interaction and separated, once and for all, the two quantities that the word "bits" conflates. The *surface* ledger — what a perfect camera, screen recorder, and network tap can physically write down — is a strictly positive, linearly growing quantity of about $\log_2 6 \approx 2.585$ bits per character plus a single final verdict bit. The *effective* ledger — the secret-relevant mutual information that those recordings actually pin down — is identically zero, per character, per ceremony, and after any number of observed ceremonies. The zero is enforced by three independent and composable mechanisms: a lossy hyperplane collapse that destroys ring/position/alphabet structure before any decision is made; a private witness synonym drawn uniformly from $6! = 720$ bijections that severs the bearing transcript from the colored challenge; and a per-round one-time XOR pad that uniformizes the wire payload by Vernam's theorem. The six-lemma proof of Section 10 establishes $I(C;O) = 0$, which Section 11 identifies as Shannon's 1949 perfect-secrecy condition applied not to a single ciphertext but to an entire interactive, replay-immune authentication transcript. Section 12 situates the result on the *unconditional* branch of cryptography: because the secret-relevant in-

formation is absent rather than merely encrypted, no future algorithm and no quantum computer can extract it, and the only surviving adversarial move is rate-limited brute force over a hypothesis space of size $720 \cdot 6^L \cdot K$, against which the best quantum advantage is a square-root Grover speedup. The practical consequence for a security team is a single quotable line for any threat model: *the bits an observer can capture and the bits an observer can use are not the same number, and in the Rosario interaction the second number is zero.*

Glossary of Terms

Index of Equations

1. **Surface vs. effective bits (thesis).** $H_{\text{surface}}(\text{char}) = \log_2 6 \approx 2.585$; $I_{\text{effective}}(\text{secret}; \text{obs}) = 0.$, §0.
2. **Disjoint rings.** $A = \{A_1, \dots, A_M\}$, $A_i \cap A_j = \emptyset.$, §1.
3. **Secret commitment.** $C = (s_1, \dots, s_L)$, $s_r \in A_{j(r)}.$, §1.
4. **Color / bearing sets.** $|\mathcal{C}| = |\mathcal{B}| = 6.$, §1.
5. **Witness synonym.** $f : \mathcal{C} \to \mathcal{B}$ bijective., §1.
6. **Rotation operator.** $\delta_a^r = H(E \| \tau \| r \| a) \bmod n_a.$, §2.
7. **Foliation.** $\Pi_6(A_a, r) = \{\alpha_{a,0}^r, \dots, \alpha_{a,5}^r\}.$, §2.
8. **Zone collapse / witness set.** $W_z^r = \bigcup_{a=1}^{M} \alpha_{a,z}^r.$, §2, §3.
9. **True color.** \$c^ _r = \$ unique c with $s_r \in W_{\chi^{-1}(c)}^r.$, §2.
10. **Bearing response.** $b_r = f(c_r^\star).$, §2.
11. **Full per-character pipeline.** $(E, \tau, r) \to G_r \xrightarrow[+f]{+s_r} b_r.$, §2.
12. **Rotation configuration count.** $N_{\text{rot}} = \prod_a n_a = 6760.$, §3.
13. **Geometry entropy bound.** $H(G_r) \le \log_2 N_{\text{rot}} \approx 12.72.$, §3.
14. **Collapse map.** $\kappa(a, z) = \chi(z)$ (surjective)., §3.
15. **Decision entropy bound.** $H(c_r^\star) \le \log_2 6 \approx 2.585.$, §3.
16. **Collapse ratio.** $\eta = \log_2 N_{\text{rot}} / \log_2 6 \approx 4.92.$, §3.
17. **Map count.** $|\{f \text{ bijective}\}| = 6! = 720.$, §4.
18. **Uniform map prior.** $\Pr[f = \varphi] = 1/720$, $f \perp (s_r, G_r).$, §4.
19. **Conditional uniformity of bearing.** $\Pr[b_r = b \mid s_r, G_r] = 1/6.$, §4.
20. **Single-round zero leakage.** $I(s_r; (G_r, b_r)) = 0.$, §4.
21. **Bearing encoding.** $\beta_r = \text{enc}(b_r) \in \{0,1\}^3.$, §5.
22. **Keystream.** $k_r = \text{trunc}_3(H(E \| \tau \| \text{nonce} \| r)).$, §5.
23. **Masking XOR.** $y_r = \beta_r \oplus k_r.$, §5.
24. **OTP perfect secrecy.** $\Pr[y_r = y \mid \beta_r] = 2^{-3}$; $I(\beta_r; y_r) = 0.$, §5.
25. **Data-processing composition.** $I(s_r; y_r) \le \min(I(s_r; \beta_r), I(\beta_r; y_r)) = 0.$, §5.
26. **Surface bits per character.** $H_{\text{surface}}(O_r) = H(G_r) + \log_2 6.$, §6.
27. **Effective bits per character.** $I(s_r; O_r) = 0.$, §6.
28. **Surface transcript entropy.** $H_{\text{surface}}(T) = L \log_2 6 \approx 2.585L.$, §7.
29. **Aggregate zero leakage.** $I(C; T, G_{1:L}, y_{1:L}) = 0 \ \forall n \ge 0.$, §7.
30. **Hypothesis space.** $|\mathcal{H}| = 720 \cdot 6^L \cdot K.$, §7.
31. **Membership indicator.** $M_r = \mathbf{1}\{s_r \in W_{z_r}^r\}.$, §8.
32. **Accumulator.** $\Lambda = \bigwedge_{r=1}^{L} M_r.$, §8.

Term	Symbol	Definition
Ring (alphabet)	A_a	One of M disjoint alphabets (uppercase, lowercase, digits) loaded onto a rotating cylinder; size n_a.
Symbol universe	S	The union of all rings; $\|S\| = \sum_a n_a = 62$ in the reference design.
Secret commitment	C	The user's secret word $(s_1, \dots, s_L)$; never transmitted, only proven.
Secret character	s_r	The target character for round r.
Rounds / secret length	L	Number of interactive rounds, one per secret character.
Color set	$\mathcal{C}$	Red, Green, Blue, Yellow, Black, White; the language of the challenge; $\|\mathcal{C}\| = 6$.
Bearing set	$\mathcal{B}$	Up, Down, Left, Right, Forward, Back ; the language of the response; $\|\mathcal{B}\| = 6$.
Witness synonym / map	f	Private bijection $\mathcal{C} \to \mathcal{B}$; the cognitive key; one of $6! = 720$.
Session entropy	E	Fresh cryptographic randomness (256–512 bits) seeding the session geometry.
Time coefficient	τ	Microsecond timestamp mixed with E to make each session unique.
Round index	r	The current round, $1 \le r \le L$.
Rotation offset	δ_a^r	Hash-derived spin applied to ring a in round r.
Foliation operator	Π_6	Slices each rotated ring into six contiguous arcs (leaves).
Ring-slice	$\alpha_{a,z}^r$	The characters of ring a in slice z during round r.
Collapse map	κ	Many-to-one map sending each (ring, slice) cell to a color; destroys ring/position info.
Coloring	χ	Assignment of the six colors to the six zone indices.
Witness set	W_c^r	The characters shown inside the color-c zone in round r (union of same-index ring-slices).
Challenge geometry	G_r	The full round-r layout $\{(c, W_c^r)\}$; what the observer sees on screen.
True color	$c_r^\star$	The color of the zone that actually contains s_r.
Bearing response	b_r	The user's emitted direction, $b_r = f(c_r^\star)$.
Bearing codeword	β_r	3-bit encoding of b_r.
Keystream	k_r	One-time 3-bit pad derived from $H(E\|\tau\|\text{nonce}\|r)$.
Masked payload	y_r	Transmitted value $y_r = \beta_r \oplus k_r$ (Vernam-masked).
Transcript	T	The sequence of bearings $(b_1, \dots, b_L)$ observable across the ceremony.
Membership indicator	M_r	$\mathbf{1}\{s_r \in W_{z_r}^r\}$; round-$r$ pass/fail check.
Accumulator / verdict	Λ	$\bigwedge_r M_r$; the single released pass/fail bit.
Surface bits	H_{surface}	Shannon entropy of what the observer can physically record.
Effective bits	$I_{\text{effective}}$	Mutual information between secret and observation; the leakage.
Hypothesis space	$\mathcal{H}$	The set the attacker must blindly search; $\|\mathcal{H}\| = 720 \cdot 6^L \cdot K$.
Collapse ratio	η	$\log_2 N_{\text{rot}} / \log_2 \|\mathcal{C}\|$; arrangement bits discarded per decision bit retained.
Mutual information	$I(\cdot;\cdot)$	Bits one variable reveals about another; zero means independence.
Binary entropy	$H_2(p)$	$-p \log_2 p - (1-p)\log_2(1-p)$; entropy of a biased coin.
Bulk	(geometry)	The high-dimensional challenge configuration (rings, rotations, slices, alphabet, position); ≈ 12.72 bits.
Boundary	b_r, Λ	The low-dimensional encoding of the bulk: the emitted bearing per round, and the single verdict bit per ceremony.
Holographic map / dictionary	$\mathcal{H}$, f	The bulk-to-boundary encoding $\mathcal{H} = f \circ \text{zone}$; reconstructable only with the private synonym f.
Grover bound	T_{quantum}	$\sqrt{\|\mathcal{H}\|}$; the square-root quantum speedup against the residual brute-force search.

33. **Honest verdict entropy.** $H(\Lambda \mid \text{honest}) = 0.$, §8.
34. **Verdict leakage bound.** $I(C;\Lambda) \le H_2(1/|\mathcal{H}|).$, §8.
35. **Net release.** $I(C;T,G,y) + I(C;\Lambda \mid \text{honest}) = 0.$, §8.
36. **Two ledgers.** Surface $= L\log_2 6 + 1$; Effective $= 0.$, §9.
37. **Observation bundle.** $O = (G_{1:L}, y_{1:L}, \Lambda).$, §10.1.
38. **Zero-usable-bits theorem.** $I(C;O) = 0 \Leftrightarrow \Pr[C \mid O] = \Pr[C].$, §10.2.
39. **Lemma 1, geometry independence.** $I(C;G_{1:L}) = 0.$, §10.3.
40. **Lemma 2, per-round bearing independence.** $\Pr[b_r = b \mid s_r, G_r] = \frac{1}{6} \Rightarrow I(s_r; b_r \mid G_r) = 0.$, §10.4.
41. **Lemma 3, mask independence.** $\Pr[y_r = y \mid b_r] = 2^{-3} \Rightarrow I(b_r; y_r) = 0.$, §10.5.
42. **Lemma 4, per-round composition.** $I(s_r;(G_r,y_r)) \le I(s_r;(G_r,b_r)) = 0.$, §10.6.
43. **Lemma 5, cross-round aggregation.** $I(C;G_{1:L},y_{1:L}) = \sum_r 0 = 0.$, §10.7.
44. **Lemma 6, secret mixture.** $\Pr[O = o] = \sum_c \pi_c \Pr[O = o \mid C = c]$, conditionals c-independent., §10.8.
45. **Verdict independence (honest).** $H(\Lambda \mid \text{honest}) = 0 \Rightarrow I(C;\Lambda \mid \text{honest}) = 0.$, §10.9.
46. **Conclusion.** $0 \le I(C;O) \le 0 \Rightarrow I(C;O) = 0.$, §10.10.
47. **Holographic boundary encoding.** $\mathcal{H} : (s_r, G_r) \mapsto b_r = (f \circ \text{zone})(s_r, G_r).$, §3.1.
48. **Shannon perfect secrecy.** $\Pr[M = m \mid C = c] = \Pr[M = m] \Leftrightarrow I(M;C) = 0.$, §11.1.
49. **Two roads to zero.** Road A: $\Pr[b_r \mid s_r, G_r] = \Pr[b_r \mid G_r]$; Road B: $I(\beta_r; y_r) = 0.$, §11.3.
50. **Grover bound (quantum).** $T_{\text{quantum}} \approx \sqrt{|\mathcal{H}|} = \sqrt{720 \cdot 6^L \cdot K}.$, §12.3.

References

[1] C. E. Shannon, "Communication theory of secrecy systems," *Bell System Technical Journal*, vol. 28, no. 4, pp. 656–715, 1949.

[2] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423 and no. 4, pp. 623–656, 1948.

[3] G. S. Vernam, "Cipher printing telegraph systems for secret wire and radio telegraphic communications," *Journal of the AIEE*, vol. 45, no. 2, pp. 109–115, 1926. (One-time pad, introduced 1917.)

[4] S. M. Bellovin, "Frank Miller: Inventor of the one-time pad," *Cryptologia*, vol. 35, no. 3, pp. 203–222, 2011. (Documents Miller's 1882 anticipation of the one-time pad.)

[5] al-Kindī, *Risāla fī Istikhrāj al-Muᶜammā* (A Manuscript on Deciphering Cryptographic Messages), 9th century. (Earliest known description of frequency analysis.)

[6] W. Diffie and M. E. Hellman, "New directions in cryptography," *IEEE Transactions on Information Theory*, vol. IT-22, no. 6, pp. 644–654, 1976.

[7] R. L. Rivest, A. Shamir, and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems," *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, 1978.

[8] S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof systems," in *Proc. 17th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 291–304, 1985. (Journal version: *SIAM J. Comput.*, vol. 18, no. 1, pp. 186–208, 1989.)

[9] P. W. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," in *Proc. 35th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 124–134, 1994.

[10] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proc. 28th*

Annual ACM Symposium on Theory of Computing (STOC), pp. 212–219, 1996.

[11] National Institute of Standards and Technology, *Module-Lattice-Based Key-Encapsulation Mechanism Standard*, FIPS 203, 2024. (ML-KEM; derived from CRYSTALS-Kyber.)

[12] National Institute of Standards and Technology, *Module-Lattice-Based Digital Signature Standard*, FIPS 204, 2024. (ML-DSA; derived from CRYSTALS-Dilithium.)

[13] National Institute of Standards and Technology, *Stateless Hash-Based Digital Signature Standard*, FIPS 205, 2024. (SLH-DSA; derived from SPHINCS+.)

[14] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Hoboken, NJ: Wiley-Interscience, 2006. (Mutual information, chain rule, data-processing inequality.)

[15] G. 't Hooft, "Dimensional reduction in quantum gravity," in *Salamfestschrift*, World Scientific, 1993, arXiv:gr-qc/9310026. (Holographic principle.)

[16] L. Susskind, "The world as a hologram," *Journal of Mathematical Physics*, vol. 36, no. 11, pp. 6377–6396, 1995.

[17] J. Katz and Y. Lindell, *Introduction to Modern Cryptography*, 3rd ed. Boca Raton, FL: CRC Press, 2020. (Perfect secrecy, one-time pad, computational vs. unconditional security.)