

New Dimensions in Cryptography

Structured-Ambiguity Commitment and Sealed-Verifier Proof of Knowledge

INVITED PAPER

Frank Dylan Rosario[†] and Dr. Lin (Grant) Wang[‡]

[†]Corresponding author. Independent Researcher, Cryptographic Systems. [‡]Co-author.

Transactions on Information-Theoretic Security, Vol. I, No. 1, 2026

May 2026

Abstract—Two kinds of contemporary developments in cryptography are examined. The widening reach of quantum computation and ever-faster classical hardware have given rise to a need for cryptographic systems whose security does not rest on the difficulty of any mathematical problem and which leave no recoverable trace of the secret in the observable record. This paper suggests a way to meet these needs. It introduces the **structured-ambiguity commitment** and the **sealed-verifier proof of knowledge**, and it argues that secrecy can be relocated from the computational difficulty of inverting a function to the information-theoretic absence of a signal in the channel. It also discusses how Shannon’s theory of communication, rather than the theory of computational complexity, is beginning to provide the tools to solve cryptographic problems whose solutions have until now depended on unproven assumptions.

Index Terms—Information-theoretic security, commitment schemes, zero-knowledge proof of knowledge, structured ambiguity, post-quantum cryptography, Shannon secrecy, transcript silence, carrier-independent protocols.

Manuscript prepared May 2026. This paper deliberately mirrors the structure, language, and argumentative arc of W.,Diffie and M.,E.,Hellman, *New Directions in Cryptography* (IEEE Trans. Inform. Theory, vol. IT-22, no. 6, 1976), recasting the public-key thesis of that work in terms of an information-theoretic, channel-shaping foundation. The underlying construction is the subject of a United States patent specification by F.,D.,Rosario.

I. Introduction

We stand today on the brink of a second revolution in cryptography. The development of large-scale quantum hardware threatens to undo the design assumptions of the public-key era and to bring the cost of inverting the mathematical problems on which that era was built down to where the secrets they protect can be recovered in practice. In turn, this prospect creates a need for a new kind of cryptographic system, one which minimizes the dependence on any unproven hardness assumption and which leaves in the observable record nothing for a future cryptanalyst to seize upon. At the same time, theoretical results in information theory, settled three quarters of a century ago but never before brought to bear on the problems of commitment and proof of knowledge, show promise of providing unconditionally secure cryptosystems, changing this hardness-bound art back into a science.

The development of cryptanalytic machinery, both classical and quantum, promises effortless recovery of secrets that today seem safely buried behind the factorization and discrete-logarithm problems. For many applications, a commitment made today must remain secure for decades; a national identity anchor, an archival legal obligation, a cryptographic seal intended to outlive its creator. At present, however, every such commitment is a wager on the durability of a mathematical assumption. Contemporary cryptography is unable to meet the requirement of permanence, in that the security it offers is mortal: it weakens precisely as the adversary’s clock speeds up, and it must eventually be replaced by a system whose hardness assumption has not yet been undermined.

The best-known cryptographic guarantee is that of computational secrecy: the secret is present, in mathematically determinate form, in the public data, and the only thing standing between the adversary and the secret is the cost of the inversion. In order to use such a system safely, however, it is currently necessary to assume that no one will ever discover a fast algorithm for the underlying problem, and that no machine will ever arise that renders the problem easy. This assumption is reasonable today and may remain reasonable for some years, but it is an assumption about the adversary’s clock, and a clock can be made to run faster. A private commitment that must survive a half-century cannot rest on a guess about which problems will still be hard in fifty years’ time.

Section III proposes an approach to committing to a secret in public without placing the secret in the public record at all. In a **structured-ambiguity commitment**, the public artifact is not a one-way summary that seals a single truth, but a scaffold that displays many equally plausible truths, E and D being replaced by a public table $\mathcal{F}$ and a private selector path σ , such that recovering σ from $\mathcal{F}$ is not merely computationally infeasible but information-theoretically impossible: the data required to single out σ is not present in $\mathcal{F}$. The public scaffold can thus be published without compromising the secret, because the secret was never written into it. Each user of the system can place a scaffold in a public directory. This enables any user to commit, to prove, and to be verified without ever transmitting the quantity that would let an observer impersonate them.

We propose techniques for developing such commitments,

and we show that the construction is not confined to one substrate.

The **sealed-verifier proof of knowledge** offers a corresponding approach to demonstrating possession of a secret without revealing it. In such a system, a prover responds to challenges over a public channel, and a verifier becomes convinced, while a third party recording the entire exchange must find it not merely difficult but impossible to recover the secret from what was recorded. A possible solution to this problem is given in Section IV, and it relies, as the commitment does, on the structural absence of information rather than on the difficulty of a computation.

A second problem, amenable to the same treatment, which stands in the way of trusting cryptographic systems over the long horizon, is authentication that does not betray its hidden state. The Enigma machine of the last century was broken not because its key space was small but because the structure of its output betrayed its hidden state, and once the betrayal was understood the rest unravelled. Every system in deployment today contains the same vulnerability in a milder form: the secret is in the public data. In order to have a purely digital instrument that does not betray its state, each user must be able to produce a transcript whose validity can be checked by the verifier, but from which no observer, not even one granted unlimited time and a working quantum computer, can extract the hidden state. Current authentication techniques cannot meet this need.

Section IV discusses the problem of providing a true, transcript-silent proof of knowledge. For reasons brought out there, we refer to this as the **sealed-verifier problem**. A solution is given, and it is shown how any structured-ambiguity commitment can be transformed into a sealed-verifier authentication system.

Section V will consider the interrelation of these constructions and introduce the notion of the **private morphism**, the structural analogue of the trap door, whose possession distinguishes the legitimate prover from every observer.

At the same time that quantum computation has given rise to new cryptographic anxieties, the older theory of communication, Shannon's theory, has begun to supply the tools for their resolution. The search for unbreakable systems is one of the oldest themes of cryptographic research, and until the present construction all proposed systems for commitment and proof have ultimately rested on a hardness assumption. In the nineteen forties, however, the one-time pad was placed on a firm foundation by information theory and shown to be unconditionally secure. One-time pads require keys as long as the message and are therefore prohibitively expensive for most applications. The construction described here achieves an analogous unconditional guarantee in the domain of commitment and proof, without requiring a key as long as the message. Section VI explores this possibility.

Before proceeding to the construction, we introduce terminology and define the threat environment in the next section.

II. Conventional and Computational Cryptography

Cryptography is the study of systems for solving two kinds of security problems: confidentiality and authentication. A confidentiality system prevents the extraction of information by unauthorized parties from data transmitted or published over a public channel. An authentication system prevents the unauthorized fabrication of a valid transcript, assuring the verifier of the legitimacy of the prover.

Almost every system in widespread deployment today, the integer-factorization system, the discrete-logarithm key-exchange and signature systems, the elliptic-curve systems, and the lattice-based and hash-based systems now being standardized for the post-quantum era, shares a single architectural assumption. The assumption is that secrecy is achieved by hiding a secret behind a mathematical problem that is easy in one direction and prohibitively difficult in the reverse direction. We will call any such system **computationally secure**. A computationally secure system is one whose public artifact fully and deterministically encodes the secret, and which is safe to publish only because recovering the secret would, on the best available hardware, take longer than the adversary can afford. By contrast, a system which can resist any cryptanalytic attack, no matter how much computation or how advanced a machine is brought against it, we will call **unconditionally secure** or **information-theoretically secure**. Unconditionally secure systems belong to that portion of information theory, the Shannon theory, which is concerned with the limits of what any receiver can extract from a channel, irrespective of computation.

The distinction is fundamental and is worth stating sharply. In a computationally secure system, the secret *lives inside the public data*; the security guarantee is a statement about the adversary's clock. When the clock speeds up, through algorithmic advance, faster hardware, or the arrival of a quantum computer, the guarantee weakens. In an unconditionally secure system, the secret *is not present in the public data at all*; the security guarantee is a statement about the channel, and no acceleration of the adversary's clock can be brought to bear, because there is nothing in the record to compute upon.

Figure 1 illustrates the flow of information in a structured-ambiguity system used for confidentiality of commitment. There are three parties: a prover, a verifier, and an eavesdropper. The prover holds a secret selector path σ which is never transmitted. The prover publishes a scaffold $\mathcal{F}$, an open table that displays many possible truths, over an insecure channel. In order to prevent the eavesdropper from learning σ , the prover does not encrypt σ ; rather, the prover declines to place σ anywhere in the public record, relying on the combinatorial structure of $\mathcal{F}$ to make σ one of k^n equally plausible interpretations. The private morphism ϕ is shared only with the verifier, over a secure out-of-band channel indicated by the dashed path in Figure 1. The secure channel cannot be used to transmit the proof itself for reasons of generality: the entire purpose is that




Figure 1. Carrier-agnostic projection. The same mathematical manifold $\mathcal{M}$ projects onto physical carriers (cards, mechanical wheels, holographic grids, or digital silicon) while preserving the three invariants.

the verification can then proceed over the open channel without leaking the secret.

A structured-ambiguity system is a single-parameter family $\{\mathcal{F}_\sigma\}_{\sigma \in \Sigma_{\mathcal{F}}}$ of equally plausible interpretations

$$\mathcal{F} : I \times L \rightarrow V^d \quad (1)$$

from a set I of n positions, each offering k labeled candidates, to a space V^d of payload values. The committed secret is the selector path $\sigma : I \rightarrow L$, chosen from the set of paths, which is of size k^n . When the message spaces of paths and payloads are equal we will denote them both by $\{M\}$. When discussing an individual interpretation $\mathcal{F}_\sigma$, we will sometimes omit mention of the family and merely refer to the path σ .

The goal in designing the system $\{\mathcal{F}_\sigma\}$ is to make the operations of committing, proving, and verifying inexpensive, but to ensure that no cryptanalytic operation against the recorded data can succeed, regardless of cost. There are two approaches to making cryptanalysis fail. The first is to make it expensive: this is the computational approach, and it succumbs to an adversary with unlimited computation. The second is to make it impossible: this is the information-theoretic approach, and it is the one taken here. Unconditional secrecy results from the existence of multiple meaningful solutions consistent with the public record. For example, the scaffold offers at each of its n positions k labeled candidates, all displayed openly, and the recorded transcript is consistent with every one of the k^n possible selector paths. A computationally secure artifact, in

contrast, contains sufficient information to uniquely determine the secret; its security resides solely in the cost of computing it.

We will call a recovery task **information-theoretically impossible** if the data required to perform it is not present in the observable record, so that no algorithm, classical or quantum, given any amount of time, can perform it. This is a stronger statement than computational infeasibility, which asserts only that the task is finite but impossibly large. The structured-ambiguity construction targets impossibility, not infeasibility.

The first step in assessing the adequacy of such a system is to classify the threats to which it is to be subjected. The following threats may occur to a system employed for confidentiality or authentication. It is assumed throughout that the opponent knows the general system in use, since this information can be obtained by studying the device or the protocol; the security must therefore reside in the hidden selector path and the private morphism alone.

- A **transcript-recording attack** is an attack in which the cryptanalyst possesses a complete recording of the public channel, all challenges, all responses, the public scaffold, and the final verification bit, for one or more sessions. This is the baseline threat, and it is assumed to be always available.
- A **cross-session accumulation attack** is an attack in which the cryptanalyst possesses recordings of many sessions and attempts to aggregate small leakages into a usable inference




Figure 2. Flow of information in a structured-ambiguity system. The recorded transcript carries no information about the hidden selector path σ or the private morphism φ .

about the hidden state.

- **A chosen-nonce or replay attack** is an attack in which the cryptanalyst replays a recorded successful transcript to a verifier expecting a different session, or attempts to choose the session parameters so as to make a recorded transcript valid again.
- **An unbounded-adversary attack** is an attack in which the cryptanalyst is granted unlimited classical computation, or a fully working quantum computer, or both, and attempts to recover the hidden state from the recorded data.

In conventional cryptography, the most formidable of the analogous threats, the chosen-plaintext attack, is countered by making the inversion expensive. Here, the most formidable threat, the unbounded-adversary attack, cannot be countered by expense, because the adversary's expense is assumed to be unbounded. It is countered instead by ensuring that the recorded data does not carry the information the adversary would need, so that unbounded computation has nothing to act upon. A system which succumbs to the transcript-recording attack is considered totally insecure; a system which resists even the unbounded-adversary attack is unconditionally secure, and it is such systems that this paper is concerned to construct.

There is a further threat which conventional cryptography cannot treat and which motivates the present work. The Enigma machine was secure against a key-space search, its key space being large, and yet it fell, because the structure of its ciphertext betrayed its hidden rotor state through statistical reg-

ularities that no operational countermeasure could fully suppress. The deployed systems of today do not betray their state through statistical regularity, but they betray it in a stronger and more permanent sense: the secret is a deterministic function of the published key, awaiting only a fast enough inversion. We will call this the **structural betrayal** of the hidden state, and we will say that a system is **transcript-silent** if its recorded output is statistically independent of its hidden state, so that no betrayal, statistical or structural, is possible. The construction of this paper is transcript-silent by design.

III. Structured-Ambiguity Commitment

As shown in Figure 1, conventional cryptography has been a derivative security measure: once a hard problem exists behind which a secret can be hidden, the security can be extended to other operations by reducing them to that problem. The effect has been to tie the entire architecture of digital security to a small number of mathematical problems and to the assumption that those problems will remain hard.

In order to develop systems whose security survives the falsification of any such assumption, this must be changed. We propose that it is possible to develop systems of the type shown in Figure 2, in which a party commits to a secret over a public channel using only publicly known techniques, and in which the public artifact does not contain the secret in any form, deterministic or otherwise. We examine the construction in two

parts, the **commitment layer** and the **verification layer**, the first of which is treated in this section and the second in the next.

A structured-ambiguity commitment is a pair $(\mathcal{F}, D)$: a public scaffold $\mathcal{F}$ and a public binding tag D , together with a private selector path σ , such that:

1. For every position $i \in I$, the scaffold $\mathcal{F}$ displays all k labeled candidate values openly, so that no candidate is privileged from the public view.
2. Given σ and $\mathcal{F}$, the value selected at each position and the binding tag D are easy to compute.
3. Given $\mathcal{F}$ and D alone, it is information-theoretically impossible to determine σ : the min-entropy of the path conditioned on the public artifact is

$$H_\infty(\sigma \mid \mathcal{F}) = n \log_2 k \quad (2)$$

so that for $k = 4, n = 128$ there remain $4^{128} \approx 2^{256}$ equally plausible paths. 4. Given σ , it is feasible to compute the public pair $(\mathcal{F}, D)$, and given a later revelation of σ the verifier can confirm the binding tag, so that the committer cannot change which path was the true one.

Because of the third property, a committer's scaffold $\mathcal{F}$ can be made public without compromising the secret path σ . The system is therefore split into two parts, a public scaffold which supplies **hiding** and a binding tag which supplies **binding**, in such a way that, given the public artifact, it is impossible to find the hidden path. This decomposition is the central structural feature of the construction. In a conventional commitment, hiding and binding are fused into a single algebraic operation, and a break of the underlying hardness assumption compromises both at once. Here they rest on distinct foundations: hiding is combinatorial and information-theoretic, depending on no assumption whatever; binding is computational, depending on the collision-resistance of the hash function used for D . A break of the hash function would compromise only the binding tag, leaving the hiding property, supplied by the scaffold, fully intact. This is defense in depth applied at the level of the cryptographic primitive itself.

The fourth property guarantees that there is a feasible way of computing the public pair from any chosen secret path when no constraint is placed on what the path is to be. In practice, the equipment must contain a true random number generator for the session entropy from which the scaffold's internal structure is drawn, together with an algorithm for generating $(\mathcal{F}, D)$ from the path and the entropy.

Given a system of this kind, the problem of long-term confidentiality is vastly simplified. Each committer generates a scaffold and a binding tag and publishes them. The selector path must be kept secret, but need never be communicated on any channel. The scaffold can be placed in a public directory along with the committer's name and address. Anyone can then verify a later proof, but no one can recover the path from the published artifacts. Structured-ambiguity commitments can thus be regarded as the information-theoretic analogue of a pub-

lished public key, with the crucial difference that the published artifact does not determine the secret.

It is worth giving a suggestive, although deliberately simplified, example of how the hiding arises. Represent the scaffold as a table of n columns and k rows, every cell of which is filled with an openly displayed value. The secret is the choice, at each column, of one row. From the table alone, every one of the k^n row-choice patterns is equally consistent with what is displayed, because every cell is shown and no marking distinguishes the chosen cell from the others. There is no statistical regularity to exploit, in sharp contrast to the Enigma, whose displayed output did carry a regularity, that no letter ever mapped to itself, from which the hidden state could be unwound. The structured-ambiguity scaffold displays nothing that distinguishes the true path; the regularity that betrayed the Enigma is, by construction, absent.

A more powerful instantiation makes the hiding property serve a proof of knowledge directly, rather than merely a later revelation. For this we introduce the new technique that constitutes the heart of the construction. It makes use, not of the apparent difficulty of computing logarithms over a finite field, but of the structural impossibility of recovering a private one-to-one map from a transcript that is statistically independent of it.

Each prover establishes, with the verifier, an independent private morphism ϕ , a bijection chosen uniformly from the $H!$ bijections between the public set of H hyperplane labels $\mathcal{H}$ and a disjoint synonym vocabulary Σ . The morphism is kept secret and shared only between prover and verifier, over an out-of-band channel, much as the secret path is kept private. At the start of each session, an entropy source produces fresh randomness $\mathcal{E}$, from which a session geometry of H leaves is derived and the associated rings of values are rotated. When a prover wishes to demonstrate knowledge, the prover responds to each challenge with a synonym

$$v_r^t = \phi(h_{z_r^t}) \quad (3)$$

rather than with the public label itself. The synonyms are drawn from a vocabulary disjoint from the public labels and from the value rings,

$$\Sigma \cap \mathcal{H} = \emptyset, \quad \Sigma \cap A_j = \emptyset \quad \forall j \quad (4)$$

so that no recorded synonym can be aligned with any visible token. Under balanced foliation, the conditional distribution of the response is independent of the leaf the prover identified,

$$\Pr[v_r^t = v_k \mid h_{z_r^t} = h_i] = \Pr[v_r^t = v_k] = \frac{1}{H} \quad \forall i, k \quad (5)$$

so that the recorded transcript distribution is identical under every one of the $H!$ candidate morphisms. Frequency analysis returns nothing; the recorded responses look like uniform




Figure 3. The structured-ambiguity scaffold. All $n \times k$ values are public. The selector path σ is private. From $\mathcal{F}$ alone, all k^n paths are equiconsistent; the secret is information-theoretically hidden.

noise.

An observer attempting to recover the morphism from the transcript faces a coupled factorial barrier. The reconstruction problem decomposes into two coupled subproblems of size $H!$ each, the recovery of the response-to-label correspondence and the recovery of the per-round leaf assignments, giving a candidate space of size

$$W_{\text{recon}} = \Omega((H!)^2) \quad (6)$$

but the operative point is not the size of this space. The operative point is that, by equation (5), every one of these candidates is exactly as consistent with the recorded transcript as every other. We thus see that if the morphism could be read from the transcript the system could be broken; but the morphism cannot be read from the transcript, because the transcript is statistically independent of it. While we do not claim that no future construction could leak less, we do not see any way for an observer to compute φ from the recorded responses without first obtaining φ itself by some means outside the channel.

If H leaves and L trials per round over R rounds are used, then the probability that an observer without the morphism produces an accepting transcript by chance is bounded by

$$\Pr[\text{forge}] \leq H^{-LR} \quad (7)$$

and against a quantum adversary armed with Grover's search this becomes $H^{-LR/2}$. For $H = 6, L = 7, R = 10$, equation (7) gives 6^{-70} , equivalent to 181 bits of classical security, and 6^{-35} , equivalent to 90 bits of post-quantum security, comfortably exceeding modern thresholds. The cryptanalytic

effort therefore grows exponentially relative to the legitimate effort, exactly as in the discrete-logarithm construction of the previous era, with the decisive difference that here the exponential barrier is reinforced by an information-theoretic one: even an adversary willing to traverse the entire $(H!)^2$ space gains nothing, because the recorded data does not rank its elements.

Worked Example.

A concrete scaffold and morphism. To make the construction definite, take a small scaffold with $n = 4$ positions and $k = 4$ labels $\{A, B, C, D\}$ per position; the public table $\mathcal{F}$ therefore displays $n \times k = 16$ values in the open. Let the committer's private path be $\sigma = (B, D, A, C)$, selecting one label at each position. From $\mathcal{F}$ alone an observer sees all sixteen values but cannot tell which label is "true" at any position; the number of equiconsistent paths is $k^n = 4^4 = 256$, and the unconditional min-entropy of equation (2) is $H_\infty(\sigma | \mathcal{F}) = n \log_2 k = 4 \cdot 2 = 8$ bits. At deployment scale ($n = 128, k = 4$) the same expression yields 256 bits, and the equiconsistent count $4^{128} \approx 10^{77}$ is of the order of the number of atoms in the observable universe.

For the proof of knowledge, fix $H = 6$ hyperplane leaves and draw a private morphism φ uniformly from the $H! = 720$ bijections onto a disjoint synonym vocabulary $\Sigma = \{v_1, \dots, v_6\}$ satisfying equation (4). With $L = 7$ trials per round over $R = 10$ rounds, equation (7) bounds an outsider's forgery probability at $\Pr[\text{forge}] \leq 6^{-70} \approx 2^{-181}$, while equation (5) makes every one of the 720 candidate morphisms produce an identical transcript distribution. An eavesdropper who records the entire session is left with the coupled barrier of equation (6), $W_{\text{recon}} = \Omega((6!)^2) \approx 5.2 \times 10^5$ candidate pairs, *every one of*




Figure 4. The coupled factorial reconstruction barrier. The morphism space and leaf assignments are coupled, creating a flat-amplitude search space with no Grover signal to amplify.

which is equally consistent with what was recorded. The toy parameters already deliver 181 bits of classical and 90 bits of post-quantum security from a layout small enough to be realized with physical cards on a table.

IV. The Sealed-Verifier Proof of Knowledge

The problem of authentication that does not betray its hidden state is perhaps an even more serious barrier to long-horizon cryptographic trust than the problem of confidentiality. Authentication is at the heart of any system involving contracts, credentials, and access. Current authentication systems cannot meet the need for a proof of knowledge whose recorded transcript is silent about the secret. They provide protection against forgery by an outsider, but the transcripts they emit expose per-round features, challenge-response pairs, individual proof terms, partial scores, which statistical and side-channel analyses can exploit across many sessions, and from which, in the limit of an unbounded adversary, the hidden state can be approached.

In order to develop a system whose transcript betrays nothing, we must discover a verification procedure with the same property as the structured-ambiguity scaffold: that the observable output is statistically independent of the hidden state. It must be easy for the verifier to become convinced, but impossible for any observer to recover the secret from what was recorded. We will call any such procedure a **sealed-verifier proof of knowledge**. Since any digital transcript can be copied

precisely, a true transcript-silent proof must be one whose recording is useless to the copier.

Consider the analogue of the login problem in a multiuser system. When a user authenticates, the standard practice is to send a credential, or a function of a credential, which the system compares against a stored value. Even when the credential is protected by a one-way function, the system still receives, per round, a value that an observer can record, and the security rests on the difficulty of inverting the one-way function. This leads to the apparently impossible requirement for a verification procedure capable of judging the validity of a proof without receiving, or emitting, any value from which the proof could be reconstructed. While appearing to be a logical impossibility, this requirement is satisfied by the following construction.

The verifier is implemented as a sealed unit that performs all per-round membership checks internally and emits, at the end of the entire session, a single bit: the value 1 if the session was internally consistent, the value 0 if it was not. The verifier does not emit per-round results, mismatch indices, partial scores, or any other intermediate information; by design it cannot be queried for finer-grained outputs. The accumulator that produces this bit is an indivisible algebraic operation,

$$a_0 = H(C), \quad a_i = H(a_{i-1} \parallel x_i), \quad Y = \text{eq}_{\text{const-time}}(a_{LR}, A^*) \quad (8)$$

in which the per-round membership results x_i are internal and never exposed, and the final comparison is performed in constant time so that a mismatch at the first position and a mis-




Figure 5. The ENIGMA Paper Print Challenge Interface. A detailed mockup of the physical/printed "Paper Challenge Interface" sheet and transparent overlay card, demonstrating operation in a zero-connectivity, physical environment.

match at the last position produce indistinguishable outputs. This denies the observer the per-round oracle that virtually every cryptanalytic attack on an interactive protocol of the past half-century has relied upon. There is no incremental query to narrow the search; there is only one bit, at the end, and that bit is the same whether the proof failed in one place or in many.

We assume that the verification procedure is public information, so that it is not ignorance of the procedure which protects the secret. Such procedures, whose observable output is statistically independent of the hidden input, are the information-theoretic counterpart of the one-way function. More precisely, a verification channel is **transcript-silent** if, for the hidden state φ and the recorded transcript $\mathbf{W}$, the mutual information satisfies

$$I(\varphi; \mathbf{W}) = 0 \quad (\text{idealized}), \quad I(\varphi; \mathbf{W}) \leq \varepsilon_\varphi \quad (\text{practical}) \quad (9)$$

It is important to note that this is a different kind of non-invertibility from the kind normally encountered in cryptography. A function is normally called one-way when its inverse is hard to compute though fully determined by the output; the secret is present in the output and merely buried. Here the secret is not present in the output at all. By Shannon's bound, $I(\varphi; \mathbf{W})$ is an upper bound on the rate at which any receiver can extract information about φ from $\mathbf{W}$; when this mutual information is zero, the bound is zero, and no algorithm, however clever, can extract from its input information that the input does

not contain. It is therefore necessary, and it is the central design obligation of the construction, that the channel be shaped so that this mutual information vanishes; a small residual ε_φ is tolerable, and, as discussed in Section VI, is the practical analogue of the idealized zero.

Each round of the protocol draws on a per-round hidden geometric quantity that depends jointly on the session entropy, a microsecond-scale timestamp, a session-specific salt, the round number, and a set of structural parameters,

$$G_r = f(\mathcal{E}, \tau, \rho, r, \Theta) \quad (10)$$

Because the timestamp τ is a closed-out historical event, a transcript captured in one session cannot be replayed in another: the geometry that would make it valid no longer exists, and the probability that a transcript valid for nonce τ is accepted under a different nonce τ' is bounded by $2^{-\lambda}$. The recorded transcript is therefore an artifact of a closed-out event, not a reusable credential. This is a stronger replay defense than a timestamp comparison, which can be defeated by clock-skew manipulation; here the geometry itself shifts with time, so that tampering produces an algorithmic mismatch rather than a mere timing rejection.

A structured-ambiguity commitment can be used to produce a sealed-verifier proof of knowledge, and conversely the sealed verifier can be used with an alternative commitment primitive, so that the two layers compose freely. If a prover wishes to demonstrate knowledge of the path committed in a




Figure 6. The sealed verifier as an information-theoretic one-way channel. Per-round membership results x_i are accumulated internally; the only observable output is one event-validity bit Y . No per-round oracle is exposed, so the recorded exchange is statistically independent of the hidden morphism φ .

scaffold, the prover engages the sealed verifier under the private morphism, responding in synonyms; the verifier accumulates the membership predicates internally and emits the single validity bit. Anyone may record the exchange, but, by equations (5) and (9), the recording does not distinguish the true morphism from any of the $H!$ alternatives, and the proof of knowledge follows immediately from the impossibility of that distinction.

There is also a one-way authentication scheme in the manner of the prior literature, in which the prover demonstrates, round by round, a familiarity that no observer can forge. The receiver is incapable of fabricating even a single accepting round without the morphism, and yet learns nothing from the rounds it observes. The combined effect of the sealed accumulator, the disjoint synonym vocabulary, the uniform marginals, the per-session geometry reset, and the nonce-bound time coefficient is that the recorded transcript is statistically identical under every one of the $H!$ possible morphisms. The observer is not searching a haystack for a needle; the observer is staring at a haystack from which the needle has been structurally removed.

V. Problem Interrelations and the Private Morphism

In this section we show that the constructions presented thus far can be reduced to one another, thereby defining a loose ordering according to strength, and we introduce the structural role of the **private morphism**, which is to this construction what the trap door was to public-key cryptography.

In Section III we showed that a structured-ambiguity scaffold provides information-theoretic hiding. In Section IV we showed that a sealed verifier provides a transcript-silent proof. Each of these can be used to build the other, and both can be used to build further cryptographic objects.

A structured-ambiguity commitment can be used to produce a sealed-verifier proof of knowledge. Take the scaffold $\mathcal{F}$ with its private path σ , establish a private morphism φ between the public labels and a disjoint synonym vocabulary, and let the prover answer challenges about σ in synonyms accumulated by a sealed verifier. The resulting proof is transcript-silent because the morphism is private and the responses are statistically independent of it. Public knowledge of the scaffold and of the verification procedure is now equivalent to public knowledge of the system without the secret, exactly the situation we require.

The converse relationship also holds in a useful direction. **A sealed verifier can be used as the proof backend for an alternative commitment primitive**, for instance a conventional algebraic commitment, so that the prover demonstrates knowledge of the committed value through the transcript-silent channel rather than through an algebraic opening. This shows that the verification layer is not parasitic on the scaffold; it is a general-purpose primitive that can be attached to other commitments to give them transcript silence they would not otherwise possess.

The private morphism is the structural analogue of the trap door. In public-key cryptography, the trap door is a piece of private information, the factorization, the discrete logarithm,

the random bit string that produced the key pair, whose possession lets the legitimate party perform an inversion that is infeasible for everyone else. The morphism plays the same role, but with a crucial difference in kind. The trap door makes an infeasible computation feasible; the holder and the observer differ in their computational reach. The morphism makes an *impossible* recovery *unnecessary*; the holder already knows the renaming and so reads the transcript directly, while the observer cannot recover the renaming from the transcript at all, not because the observer's computation is too slow, but because the transcript does not contain it. The situation is precisely analogous to a combination lock, as the public-key literature observed of the trap door: anyone who knows the combination opens it in seconds, and one who has forgotten it is no better off than a stranger. But where the trap-door lock can in principle be picked, given enough time, the morphism lock cannot be picked at all, because the recorded scratches on its face are independent of its combination.

This distinction induces an ordering. A construction whose security rests on the impossibility of extracting absent information is strictly stronger than one whose security rests on the infeasibility of a computation, because the former survives any advance that defeats the latter. We may therefore say that the structured-ambiguity commitment and the sealed-verifier proof occupy a category above the computationally secure systems: they can simulate the security guarantees of those systems, in the domains of commitment and proof, while the computationally secure systems cannot simulate the unconditional guarantee of these.

A subtle point concerns composition. In computational systems, a small leak per session often composes: many transcripts aggregate into a usable inference, which is the cross-session accumulation attack of Section II. Here, because each session draws fresh entropy and the per-session transcript is statistically independent of the morphism, statistical independence does not aggregate into dependence. An observer who records a thousand sessions holds a thousand independent samples of uniform noise, and a thousand samples of noise carry no more information about the morphism than one does, namely none. The construction is therefore secure not only per session but under unbounded accumulation, a property that the computationally secure systems, whose every transcript leaks a determinate function of the key, cannot claim.

VI. Information-Theoretic Security

Cryptography differs from all other fields of endeavor in the ease with which its requirements may appear to be satisfied. A construction that looks secure, that scrambles its input beyond casual recognition, may yet betray its secret to a patient analyst, and the history of the subject is a history of systems that appeared sound and were not. In consequence, judging the worth of a new construction has always been the central concern of the cryptographer. For most of the modern era, that judgment has been made by reference to computational com-

plexity: a system was deemed secure if breaking it could be reduced to a problem believed to be hard.

It is worth recalling how the appeal to complexity arose, because the present construction represents a return to an older and firmer standard. During the era of public-key cryptography, the security of a system was argued by showing that an attack would require solving a problem in a class believed to be intractable. A function was said to belong to the class **P** if it could be computed in time bounded by a polynomial in the length of its input, and to the class **NP** if a proposed solution could be verified in polynomial time. The great open question, whether **P** equals **NP**, remains unsettled, and the security of the deployed systems rests, ultimately, on the *belief* that certain problems lie outside **P**. The general cryptanalytic problem was observed to be no harder than **NP**, since a key can be guessed nondeterministically and verified in polynomial time, and certain cryptographic problems were shown to be **NP**-complete. This was a powerful framework, but it purchased its guarantees on credit: every such guarantee is contingent on an unproven separation of complexity classes and on the continued absence of a fast algorithm.

The present construction does not draw on the theory of computational complexity at all. It draws instead on the theory of communication, and specifically on Shannon's result that the mutual information $I(X; Y)$ between a source X and an observation Y is an upper bound on the rate at which any receiver can learn about X from Y . This bound is **non-relativizing**: it holds for classical receivers, for quantum receivers, for receivers with oracle access to functions we believe to be uncomputable, and for receivers we have not yet imagined. It does not assume a polynomial bound on the adversary, or any resource bound whatever. When a channel is shaped so that

$$I(\varphi; \mathbf{W}) = 0 \quad (11)$$

the bound on extractable information is zero, and the impossibility of recovery is not a conjecture about complexity classes but a theorem about channels. We note that, just as Shannon showed the one-time pad to offer perfect secrecy in the domain of encryption, the construction of this paper offers the analogous perfect guarantee in the domains of commitment and proof, and it does so without the operational impracticality, a key as long as the message, that has kept the one-time pad from general use.

We can now identify the position of the general transcript-recovery problem among all problems. In a computationally secure system, the recovery problem is well defined and merely hard: the secret is a determinate function of the public data, and the adversary's task is to compute the inverse. In the present construction, the recovery problem is **ill defined**: there is no function from the recorded transcript to the secret, because many secrets are equally consistent with the transcript, and no amount of computation can choose among equiconsistent alternatives. The adversary in a computational system is fighting *time*; the adversary here is fighting *information*, and

an adversary fighting information loses regardless of time or machine.

This reframing has a direct bearing on the quantum question, which is the proximate motivation for the work. Two quantum algorithms threaten the deployed systems. Shor’s algorithm efficiently solves the integer-factorization and discrete-logarithm problems, and so breaks the factorization, discrete-logarithm, and elliptic-curve systems by attacking the algebraic group inversions on which they rest. The present construction performs no algebraic group inversion; there is no modulus to factor, no logarithm to compute, no curve point to invert. Shor’s algorithm is therefore not resisted but rendered *categorically inapplicable*: it has no structure in the construction to attach to. Grover’s algorithm provides a quadratic speed-up for search problems equipped with an oracle that distinguishes a target from a non-target. For the task of forging an accepting transcript, Grover halves the security exponent, from H^{-LR} to $H^{-LR/2}$, and this is compensated simply by doubling the round count. For the task of recovering the morphism, Grover obtains *no advantage at all*, because the oracle it would amplify is flat: by equation (5), every candidate morphism produces the same observable statistics, so the oracle returns the same value for all $H!$ candidates, the amplitude rotation is zero, and the quantum search collapses to classical behaviour over a space whose elements are in any case equicon-sistent. Grover requires a signal to amplify; the construction denies it one.

The contrast with the post-quantum standardization effort is therefore not one of degree but of kind. That effort selects new hardness assumptions, lattice problems, hash-inversion problems, believed to resist quantum attack, and the security of the resulting systems remains contingent on those beliefs. The present construction selects *no* hardness assumption, and so has none that a future algorithm could falsify. The standardization story asks which hardness assumption the security argument should rest upon; the present story asks whether the security argument should rest upon a hardness assumption at all, and answers that, in the domains of commitment and proof, it need not.

For a planner whose system must remain secure for decades, this distinction is decisive. A hardness-based system, however carefully chosen, is a wager on the durability of a mathematical belief. The present construction is not a wager; it is a structural argument that does not become more or less likely to hold as the years pass. The bits in the recorded transcript are what they are, Shannon’s bound is what it is, and the construction inhabits the space those two facts define.

VII. Comparison and Security Parameters

The categorical difference between the present construction and the deployed primitives is most clearly seen by placing them side by side. Table 1 compares each major primitive on the axes that matter: the basis of its security, whether the secret is present in the public data, its standing against the two

principal quantum algorithms, and whether it is locked to a particular computational substrate.

The security level of the proof of knowledge is set by three engineering parameters: the number of hyperplane leaves H , the number of trials per round L , and the number of rounds R . By equation (7), the classical security in bits is $LR \log_2 H$, and the post-quantum security under Grover’s bound is half that. Table 2 gives representative configurations; the same bounds are achievable in a purely physical embodiment.

VIII. Historical Perspective

While at first the structured-ambiguity commitment and the sealed-verifier proof may appear unportended by past cryptographic developments, it is possible to view them as the natural outgrowth of trends stretching back through the whole history of the subject, and in particular as the resolution of the lesson taught by the Enigma.

Secrecy is at the heart of cryptography, but there has always been a confusion about *where* the secrecy should reside. The earliest cryptosystems, the Caesar cipher and its kind, kept the entire method secret and were undone the moment the method was learned. The invention of the distinction between a general system and a specific key, codified by Kerckhoffs in the last century in the principle that the compromise of the system should inconvenience no one, moved the secret from the method to the key. The rotor machines of the twentieth century, of which the Enigma was the most famous, embodied this principle: the wiring was, in effect, public, and the security was meant to reside in the daily key. And yet the Enigma fell, not because its key space was small, but because the *structure of its output betrayed its hidden state*; the cipher never mapped a letter to itself, and that single regularity, joined to the predictable habits of its operators, gave the codebreakers at Bletchley Park a foothold from which the whole system could be unwound.

The lesson of the Enigma was learned only in part by the cryptography that followed. The public-key revolution of the nineteen seventies moved the secret yet further from view, allowing the enciphering key itself to be published and resting the security on the difficulty of inverting a one-way function. This was a great advance, and it decreased once more the portion of the system that had to be hidden. But it did not heed the deeper lesson, for it left the secret *present in the public data*, a determinate function of the published key, betrayed in the same structural sense as the Enigma’s hidden rotor state, and awaiting only a fast enough inversion to be recovered. Each development in the history of the subject decreased the portion of the system that had to be protected from public knowledge; the present construction is the natural continuation of that trend to its limit, the point at which the secret is removed from the public data altogether and the betrayal that undid the Enigma becomes, for the first time, structurally impossible.

There is a second historical thread worth drawing out. The failure of repeated attempts to demonstrate the soundness of cryptographic systems by mathematical argument led, after the

Table 1. Comparison with deployed cryptographic primitives. Every computationally secure primitive keeps the secret in the public data and weakens as the adversary’s clock speeds up; the structured-ambiguity construction removes the secret from the record entirely and so is invariant to that acceleration.

Primitive	Security basis	Secret in public data?	Shor	Grover	Substrate
RSA	Integer factorization (computational)	Yes (determines key)	Broken	Mild	Digital only
Elliptic-curve	EC discrete log (computational)	Yes (determines key)	Broken	Mild	Digital only
Conventional ZK (SNARK/STARK)	Algebraic hardness + setup	Yes (in transcript terms)	Often broken	Mild	Digital only
Lattice / hash PQC	Short-vector / preimage (computational)	Yes (determines key)	Resisted	Mild	Digital only
One-time pad	Information-theoretic	No	Inapplicable	Inapplicable	Digital only
This work	Information-theoretic (channel shape)	No (absent by design)	Inapplicable	No target to amplify	Card, wheel, hologram, acoustic, digital

Table 2. Representative security parameterizations. The configuration $H = 6, L = 7, R = 10$ delivers 181-bit classical and 90-bit post-quantum security, exceeding the 128-bit classical and 96-bit post-quantum thresholds set for present-day systems, and does so in the simplest physical embodiment.

H	L	R	Classical bits ($LR \log_2 H$)	Post-quantum bits	Representative carrier
4	8	2	32	16	demonstration only
6	6	8	124	62	cards / digital
8	6	8	144	72	wheel / digital
10	5	8	133	66	digital
6	7	10	181	90	cards on a glass table

sixteenth and seventeenth centuries, to a long period in which the only accepted certification of a system was its survival under cryptanalytic assault. The arrival of computational complexity theory in the nineteen seventies seemed to promise a return to mathematical certification, but the promise was only partial, for the proofs it offered were conditional on unproven assumptions about complexity classes. With the application of Shannon’s information theory to commitment and proof, the position of mathematical certification comes full circle and is reestablished on an unconditional footing: the guarantee offered here is not that breaking the system reduces to a problem believed to be hard, but that the information required to break it is not present in the record, which is a theorem and not a belief.

The last characteristic we note is the medium-independence of the construction, which has no parallel in the prior history. Cryptography has always been bound to its substrate: the rotor machine to its gears, the public-key system to its arithmetic. The present construction depends only on three projection invariants, that the public labels remain readable, that the value rings remain detectable, and that the private morphism remains concealed, and any physical medium that can satisfy those three invariants is a valid carrier. The same protocol, with the same parameters and the same security argument, can

be carried by a deck of cards on a glass table, by a wheel turning under a pointer, by a hologram, by an acoustic signal, or by software on silicon. A commitment may thus be archived in physical form, a printed card sealed in a safe, and verified a hundred years hence by the same protocol on whatever substrate then exists, long after the computational cryptography of today has been broken.

We hope this property will inspire others to work in a domain in which cryptographic security is, for the first time, the shape of the channel rather than the difficulty of a sum.

IX. Conclusion

The conventional cryptographic primitives of the past half-century share a single architectural assumption: the secret is kept in the public data, and security rests on the intractability of recovering it. This assumption has served the field with distinction, and it will serve it for years yet, but in the era of quantum computation and ever-faster classical hardware it has become a structural liability, for the guarantee it offers weakens precisely as the adversary’s clock speeds up.

This paper has proposed an alternative built on two composable primitives. The **structured-ambiguity commitment** publishes a scaffold of k^n equally plausible truths and with-




Figure 7. The Multi-Verifier Quorum & Threshold Accumulator. System topology for a k -of- n threshold verifier quorum aggregating independent binary residues under a binomial threshold.

holds the selector path, supplying hiding that is information-theoretic and binding that is computational, each on an independent foundation. The **sealed-verifier proof of knowledge** lets the holder of a private morphism demonstrate possession over an open channel while emitting only a single event-validity bit, so that the recorded transcript is statistically independent of the hidden state. Together they relocate secrecy from the difficulty of inverting a function to the structural absence of a recoverable signal in the channel, a guarantee that by Shannon’s bound no advance in algorithms, hardware, or quantum computation can erode.

Three consequences follow. The security argument is independent of any computational model; the public artifact does not become a liability when assumptions change; and the deployment surface is decoupled from any particular substrate, so that the same protocol and the same security argument apply on cards, wheels, holograms, and silicon alike. The construction is not a faster algorithm for an existing problem but a reframing of the problem, such that the adversary’s category of attack is no longer well defined. It belongs in the same tradition as Shannon’s information-theoretic secrecy and the Diffie–Hellman re-engineering of the cryptographic channel: the tradition of changing the operation itself so that the question of difficulty no longer arises.

Acknowledgments

The authors thank the cryptographic-engineering community whose scrutiny of transcript-level and side-channel attacks

shaped the threat model adopted here, and acknowledge the foundational influence of C. E. Shannon, whose information-theoretic treatment of secrecy is the basis on which the present construction rests.

References

1. C. E. Shannon, “Communication Theory of Secrecy Systems,” *Bell System Technical Journal*, vol. 28, no. 4, pp. 656–715, Oct. 1949.
2. C. E. Shannon, “A Mathematical Theory of Communication,” *Bell System Technical Journal*, vol. 27, pp. 379–423 and 623–656, Jul. and Oct. 1948.
3. W. Diffie and M. E. Hellman, “New Directions in Cryptography,” *IEEE Trans. Inform. Theory*, vol. IT-22, no. 6, pp. 644–654, Nov. 1976.
4. F. D. Rosario, *Systems and Methods for Structured Ambiguity Commitments and Entropic Hyperplane Proof-of-Knowledge Verification*, U.S. patent specification, PATENT2026/GRAND.ROSARIO.PATENT.2026.001.md.
5. F. D. Rosario, *Why an Eavesdropper Cannot Use Recorded Bits*, technical companion, PATENT2026/CORE.md.
6. R. L. Rivest, A. Shamir, and L. M. Adleman, “A Method for Obtaining Digital Signatures and Public-Key Cryptosystems,” *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, Feb. 1978.
7. S. Goldwasser, S. Micali, and C. Rackoff, “The Knowledge




Figure 8. The Mechanical Wheel/Rotor Verifier. A mechanical engineering drawing of a rotor-based verification device with concentric rotating wheels, a pointer, and a flywheel-driven physical entropy source.

- Complexity of Interactive Proof Systems,” *SIAM J. Comput.*, vol. 18, no. 1, pp. 186–208, Feb. 1989.
8. T. P. Pedersen, “Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing,” *Advances in Cryptology — CRYPTO ’91*, pp. 129–140, 1991.
 9. P. W. Shor, “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer,” *SIAM J. Comput.*, vol. 26, no. 5, pp. 1484–1509, Oct. 1997.
 10. L. K. Grover, “A Fast Quantum Mechanical Algorithm for Database Search,” *Proc. 28th Annual ACM Symp. on Theory of Computing (STOC ’96)*, pp. 212–219, May 1996.
 11. A. Scherbius, “Cipher Machine,” U.S. Patent 1,657,411, Jan. 24, 1928.
 12. D. Kahn, *The Codebreakers: The Story of Secret Writing*. New York: Macmillan, 1967.
 13. R. M. Karp, “Reducibility Among Combinatorial Problems,” in *Complexity of Computer Computations*, R. E. Miller and J. W. Thatcher, Eds. New York: Plenum, 1972, pp. 85–104.
 14. National Institute of Standards and Technology, *FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard*, Aug. 2024.
 15. National Institute of Standards and Technology, *FIPS 205: Stateless Hash-Based Digital Signature Standard*, Aug. 2024.

Appendix A: Notation

Author Biographies

Frank Dylan Rosario is the inventor of the structured-ambiguity commitment and the sealed-verifier proof of knowledge described in this paper, and is the author of the underlying United States patent specification. His work concerns information-theoretic foundations for cryptographic confidentiality and carrier-independent verification.

Dr. Lin (Grant) Wang contributed to the information-theoretic analysis and the comparative framing of the construction against deployed public-key, zero-knowledge, and post-quantum primitives.

Suggested Citation

F. D. Rosario and L. Wang, “New Dimensions in Cryptography: Structured-Ambiguity Commitment and Sealed-Verifier Proof of Knowledge,” *Transactions on Information-Theoretic Security*, vol. I, no. 1, 2026.

This paper deliberately mirrors the structure, language, and argumentative arc of W. Diffie and M. E. Hellman, “New Directions in Cryptography” (IEEE Trans. Inform. Theory, 1976). For the formal engineering disclosure, parameter tables, and claim language, the reader is referred to the patent specification

Table 3. Notation used throughout this paper.

Symbol	Meaning
$\mathcal{F}$	Public structured-ambiguity scaffold (table of $n \times k$ openly displayed values)
σ	Private selector path, $\sigma : I \rightarrow L$; the committed secret
n, k	Number of positions and labels per position; k^n equiconsistent paths
D	Public binding tag $D = H(\text{enc}(\sigma) \parallel \text{enc}(s_\sigma))$
$H_\infty(\sigma \mid \mathcal{F})$	Conditional min-entropy of the path given the public artifact, $= n \log_2 k$
φ	Private bijective morphism from public labels $\mathcal{H}$ to synonym vocabulary Σ
$\mathcal{H}, \Sigma$	Public hyperplane-label set and disjoint synonym vocabulary; $\Sigma \cap \mathcal{H} = \emptyset$
H	Number of hyperplane leaves; morphism space has $H!$ members
L, R	Trials per round and number of rounds
v_r^t	Synonym response emitted at trial r , round t
G_r	Per-round hidden geometry $G_r = f(\mathcal{E}, \tau, \rho, r, \Theta)$
$\mathcal{E}, \tau, \rho, \Theta$	Session entropy, microsecond timestamp/nonce, salt, structural parameters
W	Recorded transcript (the full observable exchange)
$I(\varphi; \mathbf{W})$	Mutual information between morphism and transcript; $= 0$ idealized
Y	Single event-validity bit emitted by the sealed verifier, $Y \in \{0, 1\}$
λ	Security parameter for replay/nonce binding

PATENT2026/GRAND.ROSARIO.PATENT.2026.001.md; for the transcript-recovery analysis, PATENT2026/CORE.md; and for the platform narrative, PATENT2026/ROSARIO-CIPHER-PLATFORM-NARRATIVE.md.