

CNVS

Closed Native Verification Systems: Formal Mathematical Synthesis

Massimo Comitato
Independent Researcher, Italy
Email: massimocomitato@gmail.com
ORCID: 0009-0002-8301-3589

Abstract

This paper introduces Closed Native Verification Systems (CNVS), a formal class of verification systems in which raw structural membership in a candidate state is primitive, whereas admissibility, accepted state validity, and state evolution are verification-dependent. Thus, an element may structurally belong to a candidate configuration independently of its current verification status; however, only candidate states satisfying local admissibility, relational consistency, and global invariant constraints are admitted as valid CNVS states.

The theory combines a typed structural universe, a context-preserving and type-compatible recursive decomposition, random assignment of non-uniform fragments, limited verifier knowledge, and global invariant checking, first enforcing data convergence in a local verification environment.

External verifiers are not provided with the internal declared payload $d \in D_o$; they receive only a task specification and independently measure the assigned attribute, the result of which is internally compared to the value maintained by the system.

CNVS therefore does not perform syntactic slicing of the payload. The decomposition generates atomic evaluation primitives (which have the local semantics to be computable via a local verification function) and injects them into a randomized routing topology.

The fragment is not the payload, but a measurement grid that is randomly distributed along the graph and that requires a convergent observation with respect to the unknown payload.

A dedicated lemma formalizes bounded topological leakage: the routing network is not required to reveal zero information about the payload, but any semantic information leaked through topology and metadata must remain quantitatively bounded and insufficient for deterministic reconstruction.

A conditional probabilistic theorem limits unauthorized reconstruction through weighted adversarial information and exponential concentration inequalities. An asymptotic scaling theorem shows that security can increase as the system expands in the presence of explicit assumptions.

The technical difficulty of reconstruction outside the system arises from the strictly local knowledge available to each verifier.

Keywords: Closed Native Verification Systems; topological leakage; emergent security; conditional min-entropy; distributed systems.

1. Primitive Domains and Typed Structural Universe

Let $B = \{0, 1\}$ and let B^* be the set of all finite binary strings.

Let $\mathcal{T}$ be a strictly finite set of structural types.

To guarantee semantic consistency and type preservation, we define the primitive domains not as flat sets, but as families of sets indexed by their structural type $\sigma \in \mathcal{T}$.

- $\text{Id}_\sigma \subseteq B^*$: the space of identifiers for type σ .
- $\text{D}_\sigma \subseteq B^*$: the space of declared data payloads structured according to type σ . An element $d \in \text{D}_\sigma$ represents a candidate assertion, digital archive, structured datum, object specification, or cryptographic record submitted by one or more users, entities, or system agents to the CNVS layer for verification.
- $\text{At}_\sigma \subseteq B^*$: the space of attribute classes for type σ .
- $\text{Obs}_\sigma \subseteq B^*$: the space of externally reported values for type σ .

Let P be a finite set of logical properties, and let $\mathcal{P}(P)$ denote the set of all its strictly finite subsets.

To model the absence of data without introducing domain collisions or ambiguity with the empty string $\varepsilon \in B^*$, we introduce a distinguished element $\perp \notin B^*$. We define the observational domain as a disjoint union:

$$\text{Obs}_\sigma^\perp \triangleq \text{Obs}_\sigma \uplus \{\perp\}$$

where the $\perp$ element rigorously denotes the absence of an external observation at non-terminal structural levels.

We define the Structural Universe $\mathcal{S}$ as a dependent sum (Sigma-type) over the set of structural types $\mathcal{T}$:

$$\mathcal{S} \triangleq \sum_{\sigma \in \mathcal{T}} (\text{Id}_\sigma \times \text{D}_\sigma \times \text{At}_\sigma \times \text{Obs}_\sigma^\perp \times \mathcal{P}(P)) \quad (1)$$

Consequently, each structural state $s \in \mathcal{S}$ assumes the dependent form:

$$s = \langle \sigma, (\text{id}, d, a, \text{obs}, p) \rangle$$

This construction guarantees ontologically that the identifier, the payload, and the attributes are intrinsically bound to their specific structural type σ at instantiation, preventing type-mismatch vulnerabilities prior to any verification phase.

2. State Space and Verification Signatures

A system state at time t is formally defined as a tuple:

$$S(t) = \langle E(t), R(t), C(t) \rangle \quad (2)$$

where:

- $E(t) \in \mathcal{P}_{\text{fin}}(\mathcal{S})$ is a finite set of structural elements active in the system at time t .
- $R(t) \subseteq E(t) \times E(t)$ is the relational graph defining structural dependencies between elements.
- $C = \{c_1, \dots, c_m\}$ is a finite set of global system invariants (consistency constraints). Formally, let $\mathcal{C}$ be the universe of all definable structural predicates. Each invariant $c_i \in \mathcal{C}$ is a boolean function $c_i : \mathcal{P}_{\text{fin}}(\mathcal{S}) \times \mathcal{P}_{\text{fin}}(\mathcal{S} \times \mathcal{S}) \rightarrow \{0, 1\}$ that evaluates a specific topological or logical consistency rule over the elements and their relations.

The state space $\mathcal{S}$ is then the strictly bounded set of all tuples $S(t)$ defined in formula (2), formally expressed as:

$$\mathcal{S} \triangleq \{ \langle E(t), R(t), C \rangle \mid E \in \mathcal{P}_{\text{fin}}(\mathcal{S}), R \subseteq E \times E, C \subseteq \mathcal{C} \} \quad (3)$$

Terminal Elements (Ter)

Let $\mathcal{D} : \mathcal{S} \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{S})$ be the recursive decomposition operator. We formally define the subset of terminal fragments (leaf nodes) at time t as atomic:

$$\text{Ter}(t) \triangleq \{s \in E(t) \mid \mathcal{D}(s) = \emptyset\} \quad (4)$$

Hierarchical Cardinality (n) and Partial States

Because the decomposition operator $\mathcal{D}$ is recursively applied, we define the hierarchical cardinality of the system as a natural number $n \in \mathbb{N}_{>0}$, representing the maximum path length from any root element to its terminal leaves. Consequently, the state $S(t)$ can be stratified into $n + 1$ discrete topological levels. We denote $S^{(k)}(t)$ as the partial state associated with decomposition level k , where $k \in \{0, \dots, n\}$. By definition, $S^{(0)}(t)$ corresponds to the root or nodal level, while $S^{(n)}(t)$ contains the terminal fragments $\text{Ter}(t)$. The recursive reconstruction and global verification process proceeds in the opposite direction, ascending from $S^{(n)}(t)$ toward $S^{(0)}(t)$.

Definition of Relational Consistency (Cons_R)

To ensure that the state topology strictly adheres to the decomposition rules, the relational consistency predicate evaluates to 1 if and only if the edge set $R(t)$ perfectly maps to the recursive decomposition operator $\mathcal{D}$ over the active element set $E(t)$:

$$\text{Cons}_R(R(t), E(t)) = 1 \Leftrightarrow \forall s, s' \in E(t), ((s, s') \in R(t) \Leftrightarrow s' \in \mathcal{D}(s)). \quad (5)$$

Verification Signatures

The Local Verification Function V_L evaluates the formal admissibility and attribute convergence of a single element. It maps an element to a tuple containing a boolean validity flag and a finite set of verified logical properties P :

$$V_L : \mathcal{S} \rightarrow \{0, 1\} \times \mathcal{P}(P). \quad (6)$$

Let $\pi_1 : \{0, 1\} \times \mathcal{P}(P) \rightarrow \{0, 1\}$,

be the standard projection operator extracting the boolean verification result.

The Global Verification Function V_G acts as the general structural veto for the framework. It evaluates the entire system state $S(t) = \langle E(t), R(t), C \rangle$, mapping it to a boolean validity flag:

$$V_G : \mathcal{S} \rightarrow \{0, 1\}$$

Using the default hierarchical cardinality n , the next global validity condition evaluates the system recursively from the terminal leaves up to the root:

$$V_G(S(t)) = 1 \Leftrightarrow (\forall s \in \text{Ter}(t), \pi_1(V_L(s)) = 1) \wedge \text{Cons}_R(R(t), E(t)) \wedge \text{Inv}_C(S(t)).$$

This recursive schema is formalized and strictly bounded in Axiom III, where global verification is definitively decomposed into local admissibility, relational consistency, and invariant satisfaction.

3. Local Verification and Admissibility Signatures

To formalize the local evaluation of a structural element, we define the verification primitives directly over the type-indexed components of the Structural Universe $\mathcal{S}$.

Let the Local Verification Function V_L be formally declared with the following signature:

$$V_L : \mathcal{S} \rightarrow \{0, 1\} \times \mathcal{P}(P).$$

For each structural type $\sigma \in \mathcal{T}$, we define — with respect to the elements active at time t — two distinct boolean verification predicates operating strictly on their corresponding index-bound domains:

1. Formal Admissibility (Form_σ): Evaluates the structural integrity, formatting, and cryptographic bounds of the element's components.

$$\text{Form}_\sigma : \text{Id}_\sigma \times D_\sigma \times \text{At}_\sigma \times \text{Obs}_\sigma^\perp \rightarrow \{0, 1\}.$$

2. Attribute Convergence (Conv_σ): Evaluates the semantic coherence between the declared payload d , the governing attributes a , and the externally reported values obs . The Conv_σ function performs the task natively without exposing the payload:

$$\text{Conv}_\sigma : \text{At}_\sigma \times D_\sigma \times \text{Obs}_\sigma^\perp \rightarrow \{0, 1\}.$$

The marker $\perp$ denotes absence of external observation and is only a structural placeholder in $\text{Obs}_\sigma^\perp$. It is neither a successful nor a failed convergence value.

For terminal fragments, local admissibility requires an actual observation $\text{obs} \in \text{Obs}_\sigma$.

If $\text{obs} = \perp$, the local verification instance remains unresolved, and the system cannot assert $\pi_1(V_L(s)) = 1$.

External observation is performed on the specific task for an element belonging to At_σ , while D_σ is internal to the system and is not accessible from the outside (**local epistemic restriction**).

Non-terminal elements carry $\text{obs} = \perp$, but their consistency is evaluated by V_G through relational consistency and invariant satisfaction.

Given an element $s = \langle \sigma, (\text{id}, d, a, \text{obs}, p) \rangle \in \mathcal{S}$, the operational definition of V_L is therefore constructed as:

$$V_L(s) \triangleq \langle \text{Form}_\sigma(\text{id}, d, a, \text{obs}) \wedge \text{Conv}_\sigma(a, d, \text{obs}), p \rangle \quad (7)$$

Let $\pi_1 : \{0, 1\} \times \mathcal{P}(P) \rightarrow \{0, 1\}$ be the standard projection to the boolean domain.

We define the local admissibility condition Adm_L for an element s as:

$$\text{Adm}_L(s) = 1 \Leftrightarrow \pi_1(V_L(s)) = 1$$

This architecture ensures that verification is inherently type-preserving and free of domain collisions, satisfying the prerequisites for the global state evaluation.

Remark 3a. (Polymorphism of Convergence and Tolerance Bounds)

The Conv_σ activity for obs on d must not be arbitrary.

Therefore:

$\forall$ type $\sigma \in \mathcal{T}$

and $\forall a \in \text{At}_\sigma$

we will have two categories:

1. Quantitative Convergence (Metric Spaces):

σ type imposes a continuous metric space (scalar or vector)

$$\text{Conv}_\sigma(a, d, \text{obs}) = 1 \Leftrightarrow \delta_\sigma(d, \text{obs}) \leq \varepsilon_a \quad (8)$$

where δ_σ is the function that calculates the gap between obs and d and ε_a defines a tolerance threshold.

2. Qualitative Convergence (Categorical and Semantic Spaces):

σ type imposes a discrete space (e.g., colors, classification labels, bit-strings).

The tolerance becomes a semantic neighborhood or a map of equivalence classes:

$$\Psi_a : D_\sigma \rightarrow \mathcal{P}(\text{Obs}_\sigma)$$

where Ψ_a maps the expected data.

$$\text{Conv}_\sigma(a, d, \text{obs}) = 1 \Leftrightarrow \text{obs} \in \Psi_a(d)$$

Convergence is 1 if obs falls within the valid neighborhood

(e.g., mapping "red" to a cluster of acceptable hex-color subsets or semantic neighbors).

This architecture securely delegates specific digitization encodings (e.g., continuous-to-discrete quantization, thresholding algorithms, clustering) to the system implementation

level, while maintaining mathematically rigorous global verification for candidate state acceptance.

Missing Observation at the Terminal Level

At the terminal level, the local verification function requires the production of an external observation $\text{obs} \in \text{Obs}_\sigma$. If the assigned external verification agent does not return this value within the time window defined by the protocol, the observational component of the corresponding terminal structured element is represented as $\text{obs} = \perp$, with $\perp \in \text{Obs}_\sigma^\perp$.

This case should not be confused with a failed local verification, in which the system states:

$$\pi_1(V_L(s^{(n)})) = 0.$$

The local verification instance remains operationally undefined.

Formally:

$$\text{Status}_L(s^{(n)}) = \perp,$$

where $s^{(n)}$ is a terminal structured element whose observational component is $\text{obs} = \perp$.

Since

$$\pi_1(V_L(s^{(n)})) \neq 1,$$

with

$$\text{Adm}_L(s) \neq 1,$$

the candidate state cannot transition to the set of allowed states for the next partial global verification.

The practical management of this unresolved condition, including timeout policies, task reassignment, suspension of the inactive verifier, recovery procedures, or verification updates, is the responsibility of the implementation layer of the network architecture.

These mechanisms preserve system continuity without altering the formal constraints of the CNVS.

4. Recursive Decomposition and Reconstruction

Let $\text{type} : \mathcal{S} \rightarrow \mathcal{T}$ be the fundamental projection function that assigns each structural element its specific type.

Let $\mathcal{D}$ be the primary decomposition operator such that:

$$\mathcal{D} : \mathcal{S} \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{S}).$$

Since the operator's output is constrained $\forall s \in \mathcal{S}$:

$$|\mathcal{D}(s)| = 0 \vee |\mathcal{D}(s)| \geq 2.$$

if $|\mathcal{D}(s)| = 0 \Leftrightarrow \mathcal{D}(s) = \emptyset, \forall s \text{ atomic.}$

if $|\mathcal{D}(s)| \geq 2$, the element is recursively decomposable in a finite way $\{f_1, f_2, \dots, f_k\}$.

The typological structure of the universe $\mathcal{S}$ allows the generation of elements heterogeneous by type σ because $\forall \sigma \in \mathcal{T}$ is associated with the finite feasible evaluative

spectrum:

$$\Omega(\sigma) \subseteq \mathcal{T},$$

where $\Omega(\sigma)$ is given by the child structural elements whose type is feasible starting from the type of the parent σ ,

such that type compatibility is preserved, not necessarily its identity,

formally:

$$\forall s \in \mathcal{S}, \forall f_i \in \mathcal{D}(s), f_i \in \mathcal{S} \wedge \text{type}(f_i) \in \Omega(\text{type}(s)), \quad (9)$$

which corresponds to

$$\Omega(\sigma) = \{\sigma\},$$

if the type retains its identity.

Thus, recursive decomposition remains compatible with the typed structural universe while allowing heterogeneous terminal fragments when the parent type admits heterogeneous local evaluative validation.

5. Non-Uniform Fragmentation

In CNVS, $\forall s \in \mathcal{S}$, the total information can be decomposed into two strictly positive functions, such that

$$I_{\text{total}}(s) = I_{\text{payload}}(s) + I_{\text{metadata}}(s) \quad (10)$$

where

- $I_{\text{payload}} : \mathcal{S} \rightarrow \mathbb{N}$ defining the semantic informational content inherited from the original root element.
- $I_{\text{metadata}} : \mathcal{S} \rightarrow \mathbb{N}$ defining the systemic structural information required for fragment identification, assignment consistency, relational reconstruction, and verification coordination.

Since the metadata and payloads are not accessible to the external verifier, which is only responsible for measuring them, unauthorized reconstruction is computationally hindered by the difficulty of inferring the topology of the structural graph of relations, the invariants, and the entropy threshold of the payload.

$\forall s \in \mathcal{S}$, the recursive decomposition is non-uniform, such that there are at least two child fragments of the same parent node that carry non-equal semantic payload measures:

1. $\forall s \in \mathcal{S}$ with $\mathcal{D}(s) \neq \emptyset$, $I_{\text{payload}}(s) = \sum_{f_i \in \mathcal{D}(s)} I_{\text{payload}}(f_i)$,
2. $\forall s \in \mathcal{S}$ with $\mathcal{D}(s) \neq \emptyset$, $\exists f_a, f_b \in \mathcal{D}(s)$ such that $I_{\text{payload}}(f_a) \neq I_{\text{payload}}(f_b)$

The system imposes finite decompression to avoid the computational overhead of metadata at the expense of the payload, which would result in an unsustainable monotonic increase in the structural cost.

6. Knowledge Restriction and Asymptotic Negligibility

Let $\mathcal{V}$ be a finite set of external verifiers. We introduce the knowledge measure function $\mathcal{K} : \mathcal{V} \times \mathcal{S} \rightarrow \mathbb{N}$, which quantifies the exact structural and semantic information (in bits) exposed to a specific verifier $v \in \mathcal{V}$. For each specific attribute $a \in \text{At}_\sigma$ of a terminal fragment $f \in \text{Ter}(t)$, let $I_{\min}(a) \in \mathbb{N}_{>0}$ denote the minimum complexity required to perform the given task without semantic loss.

It must necessarily hold that $\mathcal{K}(v, f) \geq 1$ since a total loss of information (knowledge convergent to zero) would make the verification undecidable.

We iteratively define the domains of total information:

- $I_{\text{total}}(f)$: the total information encoded in $f^{(n)}$, where n is the cardinality of maximum recursive decomposition.
- $I_{\text{total}}(s)$: the total information encoded in the reconstructed parent node $s = \text{Rec}(\mathfrak{D}(s))$.
- $I_{\text{global}}(S(t))$: the total information of the entire state of the system at time t .

The information subordination chain is rigorously constrained as follows:

$$I_{\min}(a) \leq \mathcal{K}(v, f) < I_{\text{total}}(f) \ll I_{\text{total}}(s) \ll I_{\text{global}}(S(t))$$

Implementation Note.

At the implementation level, the payload and structural metadata are compatible with ordinary cryptography. Each verifier $v \in \mathcal{V}$ accesses with a unique cryptographic key only the specific task assigned to it, but not that of the others, while the payload and metadata remain further protected. This application is not a replacement for the CNVS model, which is based on probabilistic assumptions, entropy, and limitation of the leakage of observable metadata, as formalized in Lemma 10.

Remark 6.

In general, the verifier's accessible knowledge $\mathcal{K}(v, f)$ does not coincide with the fragmental information $I_{\text{total}}(f)$ itself. The verifier receives only a minimal semantic projection required to independently report an external observation $\text{obs} \in \text{Obs}_\sigma^\perp$.

The strict inequality $\mathcal{K}(v, f) < I_{\text{total}}(f)$ guarantees that no single verifier holds the complete semantic state of a fragment.

Furthermore, the operator $\ll$ denotes asymptotic informational negligibility.

To formalize this without violating the finite cardinality bounds of $\mathcal{P}_{\text{fin}}(\mathcal{S})$, we define a sequence of system configurations using a global structural fragmentation scaling

parameter $m_{\text{Smin}} \in \mathbb{N}$, which denotes the lower bound of the system-level decomposition:
 $m_{\text{Smin}} = \min_{s \notin \text{Ter}(t)} |\mathcal{D}(s)|$,

where $\min$ is the operator for extracting the smallest fragmentation degree (or partition size) of $\mathcal{D}$.

By the principle of informational conservation, the parent node information scales with its partition complexity: $I_{\text{total}}(s) \geq I_{\text{payload}}(s) = \sum_{f_i \in \mathcal{D}(s)} I_{\text{payload}}(f_i)$.

Consequently, for any strictly bounded local verifier knowledge $\mathcal{K}(v, f)$, the informational ratio converges to zero as the macroscopic system scale expands:

$$\lim_{m_{\text{Smin}} \rightarrow \infty} \sup_{v \in \mathcal{V}, s \notin \text{Ter}(t), f \in \mathcal{D}(s)} (\mathcal{K}(v, f) / I_{\text{total}}(s)) = 0$$

This provides an asymptotic epistemic restriction condition supporting the subsequent security analysis.

Internal Payload and External Measurement-Grid Convention

A terminal fragment in CNVS must be understood as a closed native evaluative unit with two distinct informational layers.

First, each terminal fragment f_i possesses an internal declared payload d_i . This payload belongs to the native CNVS state and remains inside the system boundary during ordinary operation.

Second, each terminal fragment exposes to the external verifier only a measurement grid or task specification derived from its attribute structure. The verifier receives the local task, performs an independent observation obs_i , and returns that observation to the system. The verifier does not receive d_i .

Thus, the terminal fragment is not a payload share in the external cryptographic sense. It may contain internal semantic payload, but that payload is not externally disclosed to the verifier.

Accordingly, $I_{\text{payload}}(f_i)$ measures the internal semantic payload associated with f_i , not the information exposed to the verifier. The verifier-accessible knowledge is instead measured by $\mathcal{K}(v, f_i)$, and in the operational CNVS model:

$$\mathcal{K}(v, f_i) < I_{\text{total}}(f_i), \text{ because } d_i \text{ remains hidden.}$$

The probabilistic security model later adopts a strictly pessimistic convention: whenever f_i is assigned to a colluding verifier, the adversary is temporarily assumed to obtain full semantic access to d_i . This is not the ordinary CNVS exposure model, but a worst-case semantic exposure assumption used to upper-bound reconstruction risk.

Remark 6.a Definition: Structural Nature of Terminal Fragments

A terminal fragment $f \in \text{Ter}(t)$ is not a mere syntactic fraction (e.g., a data "shard" or a cryptographic sub-string) of the global payload, but rather constitutes an autonomous evaluative primitive. The decomposition function $\mathcal{D}$ does not geometrically slice the payload, but extracts localized logical predicates from the global state, encapsulating them into independent, non-uniform, discrete, and executable measurement tasks. Each task can be traced back to a terminal fragment that preserves a minimum local semantic coherence, strictly respecting the constraint: $I_{\text{payload}}(f) \geq I_{\text{min}}$.

These fragments are finally distributed by the system along a hidden graph, dispersing its connections in order to preclude the reconstruction of the global information by external verifiers.

7. Random Assignment

To prevent deterministic collusion and isolate adversarial clusters, terminal fragments are dynamically assigned to external verifiers through a randomized topological mapping.

Let $\mathcal{V}$ be the finite set of external verifiers, and let:

$$Q_v \triangleq |\mathcal{V}|$$

be the size of the verifier pool.

Let Ξ denote a cryptographic probability space acting as the source of secure randomness.

Let:

$$\text{Ter}(t) = \{f_1, \dots, f_k\}$$

be the set of terminal fragments active at time t , with:

$$k \triangleq |\text{Ter}(t)|. \tag{11}$$

The assignment mapping at time t is formalized as a probabilistic function:

$$\mathcal{A}_t : \text{Ter}(t) \times \Xi \rightarrow \mathcal{V}. \tag{12}$$

To satisfy this injective mapping constraint, the system structurally assumes $k \leq Q_v$ at any given assignment cycle t .

The CNVS framework imposes a one-fragment-per-verifier constraint at each assignment cycle. Therefore, for any random seed $\xi \in \Xi$ and for any two distinct terminal fragments $f_i, f_j \in \text{Ter}(t)$:

$$f_i \neq f_j \Rightarrow \mathcal{A}_t(f_i, \xi) \neq \mathcal{A}_t(f_j, \xi).$$

Thus, for each $\xi \in \Xi$, the map $\mathcal{A}_t(\cdot, \xi)$ is an injective matching from $\text{Ter}(t)$ into $\mathcal{V}$.

If an active verifier set $\mathcal{V}_t \subseteq \mathcal{V}$ with $|\mathcal{V}_t| = k$ is selected at time t , then $\mathcal{A}_t$ induces a bijection:

$$\mathcal{A}_t(\cdot, \xi) : \text{Ter}(t) \rightarrow \mathcal{V}_t. \quad (13)$$

Accordingly, the hidden assignment is not a set of independent verifier choices with replacement. It is a randomized injective matching, or equivalently a randomized permutation over the active verifier slots.

Let $\mathfrak{A}_{\text{inj}}(t)$ denote the set of all admissible injective assignments:

$$\mathfrak{A}_{\text{inj}}(t) \triangleq \{\alpha : \text{Ter}(t) \rightarrow \mathcal{V} \mid \alpha \text{ is injective}\}.$$

The general CNVS assignment law is a probability distribution over admissible injective assignments:

$$\mu_t : \mathfrak{A}_{\text{inj}}(t) \rightarrow [0, 1],$$

with:

$$\sum_{\alpha \in \mathfrak{A}_{\text{inj}}(t)} \mu_t(\alpha) = 1.$$

This formulation allows non-uniform randomized matching, including risk-weighted verifier selection, criticality-weighted assignment, or redundancy-aware assignment, provided that the realized assignment remains cryptographically randomized and not externally predictable.

For a terminal fragment f_i and a verifier $v \in \mathcal{V}$, the marginal assignment probability is defined as:

$$\mu_t(v, f_i) \triangleq \mathbb{P}(\mathcal{A}_t(f_i, \xi) = v)$$

or equivalently:

$$\mu_t(v, f_i) = \sum_{\alpha \in \mathfrak{A}_{\text{inj}}(t), \alpha(f_i) = v} \mu_t(\alpha). \quad (14)$$

For every $f_i \in \text{Ter}(t)$, the marginal probabilities satisfy:

$$\sum_{v \in \mathcal{V}} \mu_t(v, f_i) = 1.$$

In the uniform bijective baseline, where $|\mathcal{V}_t| = k$ and all bijections between $\text{Ter}(t)$ and $\mathcal{V}_t$ are equiprobable, the number of admissible assignments is:

$$k!,$$

and the corresponding assignment entropy is:

$$H_{\text{assign}}(k) = \log_2(k!). \quad (15)$$

In that case, for every $v \in \mathcal{V}_t$:

$$\mu_t(v, f_i) = 1 / k.$$

If the active verifier set is not fixed in advance and the assignment is uniformly sampled among all injective maps from $\text{Ter}(t)$ into the full verifier pool $\mathcal{V}$, with $Q_v = |\mathcal{V}| \geq k$, then the number of admissible injective assignments is:

$$Q_v! / (Q_v - k)!,$$

and the corresponding assignment entropy is:

$$H_{\text{assign}}(Q_v, k) = \log_2(Q_v! / (Q_v - k)!).$$

In this full-pool uniform injective case, the marginal assignment probability satisfies:

$$\mu_t(v, f_i) = 1 / |\mathcal{V}| = 1 / Q_v,$$

for every $f_i \in \text{Ter}(t)$ and every $v \in \mathcal{V}$.

However, this marginal uniformity does not imply independence between different fragment assignments. Under injective matching, if one verifier is assigned to one fragment, the same verifier cannot simultaneously be assigned to another fragment in the same assignment cycle.

Therefore, Section 7 defines the operational CNVS assignment mechanism as a randomized injective matching. The adversarial acquisition variables, compromise probabilities, and concentration bounds are introduced separately in the probabilistic security model.

8. Axiom I: Verification-Independent Structural Existence and Verification-Dependent Admissibility

The structural existence of an element within a candidate state of the system is independent of its verification state, but its admissibility is constrained by preliminary local verification.

Let $\mathcal{S}_{\text{cand}}$ denote the candidate state space. For any candidate state

$$S_{\text{cand}}(t) = \langle E_{\text{cand}}(t), R_{\text{cand}}(t), C \rangle \in \mathfrak{S}_{\text{cand}}, \quad (16)$$

the relation $s \in E_{\text{cand}}(t)$ denotes raw structural membership in the candidate configuration. This membership relation is not created, destroyed, or modified by the verification maps V_L or V_G . Verification does not determine whether a candidate element exists; it determines whether the element, and ultimately the state containing it, is admissible.

For every terminal element $s \in \text{Ter}(t)$, local admissibility is defined by:

$$\text{Adm}_L(s) = 1 \Leftrightarrow \pi_1(V_L(s)) = 1.$$

Accordingly, the locally admissible terminal subset is:

$$\text{Ter}^+(t) \triangleq \{s \in \text{Ter}(t) \mid \text{Adm}_L(s) = 1\}.$$

The accepted CNVS state space is defined as:

$$\mathfrak{S}^+ \triangleq \{S \in \mathfrak{S}_{\text{cand}} \mid V_G(S) = 1\}. \quad (17)$$

Therefore:

$$s \in E_{\text{cand}}(t) \not\Rightarrow \text{Adm}_L(s) = 1,$$

and

$$(\forall s \in \text{Ter}(t), \text{Adm}_L(s) = 1) \not\Rightarrow V_G(S_{\text{cand}}(t)) = 1.$$

The first implication fails because structural membership precedes local verification, on which its admissibility depends (V_L). The second implication fails because global validity additionally requires relational consistency and invariant satisfaction (Axiom III).

Thus, verification-dependent existence in CNVS must be understood strictly as verification-dependent admissibility into the accepted state evolution, not as primitive creation of raw structural elements. The verification process partitions candidate states into accepted and rejected states without retroactively denying the structural existence of the rejected candidate configuration.

9. Axiom II: Randomized Non-Uniform Fragmentation

Every element $s \in \mathcal{S}$ such that $\mathcal{D}(s) \neq \emptyset$ is recursively decomposed into a non-uniform set of fragments. The assignment operator $\mathcal{A}_t$ defines a probabilistic mapping to the verifier set $\mathcal{V}$:

- $|\mathcal{D}(s)| \geq 2$
- $\exists f_a, f_b \in \mathcal{D}(s) : I_{\text{payload}}(f_a) \neq I_{\text{payload}}(f_b)$
- $\text{Rec}(\mathcal{D}(s)) = s$

The assignment operator is a randomized injective matching as defined in Section 7.
The independent-assignment model is used only later as an analytical baseline.

10. Axiom III: Non-Reducibility of Global Validity

Global validity is not an emergent property of terminal elements, but a systemic invariant. Let $\text{Cons}_R(R(t), E(t))$ be the relational consistency and $\text{Inv}_C(S(t))$ be the satisfaction of the global invariants C . The Global Verification Function is defined as the strict conjunction:

$$\begin{aligned}\text{Cons}_R(R(t), E(t)) &= 1 \Leftrightarrow \forall s, s' \in E(t), ((s, s') \in R(t) \Leftrightarrow s' \in \mathcal{D}(s)), \\ \text{Inv}_C(S(t)) &= 1 \Leftrightarrow \forall c_i \in C, c_i(E(t), R(t)) = 1.\end{aligned}$$

$$V_G(S(t)) = 1 \Leftrightarrow (\forall s \in \text{Ter}(t), \pi_1(V_L(s)) = 1) \wedge \text{Cons}_R(R(t), E(t)) \wedge \text{Inv}_C(S(t)) \quad (18)$$

Non-Reducibility Principle:

$$\exists S(t) \in \mathcal{S} \text{ such that } (\forall s \in \text{Ter}(t), \pi_1(V_L(s)) = 1) \wedge (V_G(S(t)) = 0) \quad (19)$$

This asserts the existence of "locally valid when admissible but globally inconsistent" states. This proves that the global system state cannot be reduced to the mere verification of its terminal fragments (or the aggregation of local verification), thereby mandating the **Global Veto**.

11. State Transition Law

The evolution of the system state from the current condition $S(t)$ towards a successor state S' is governed by the local admissibility condition on the candidate state S'_{cand} and subsequently by the effect of the Global Veto on the admitted candidate state:

$$S(t+1) = \begin{cases} S' & \text{if } V_G(S'_{\text{cand}}) = 1 \\ S(t) & \text{if } V_G(S'_{\text{cand}}) = 0 \end{cases} \quad (20)$$

Any candidate state S'_{cand} with $V_G(S'_{\text{cand}}) = 0$ is formally rejected.

This establishes that state evolution is strictly verification-dependent and constrained by the totality of structural invariants C and the relational consistency Cons_R . The system admits no "invalid" transitions; the state space is self-correcting via the Global Veto.

Remark 11.1 (Topological Refresh and Sybil/Fault-Injection Resistance)

To preclude cross-cycle structural and semantic inference, the CNVS framework enforces an absolute Topological and Semantic Refresh of the measurement metadata tied to the native payload upon any state rejection. An adversarial coalition (e.g., executing a Sybil attack) might theoretically attempt a Fault Injection strategy: intentionally submitting invalid external observations to deliberately trigger the Global Veto ($V_G(S') = 0$).

If the structural topology and the measurement grids were static, the adversary could iterate this process infinitely, accumulating semantic and topological metadata across successive failed cycles to eventually reconstruct the relational graph and infer the hidden native payload.

To mathematically neutralize this vector, the CNVS mandates that the recursive decomposition operator ($\mathfrak{D}$), the random assignment mapping ($\mathcal{A}_t$) and the assigned tasks are strictly ephemeral and bound to the specific transition attempt.

Therefore, whenever $V_G(S'_{\text{cand}}) = 0$ (whether due to a local admissibility failure, a relational inconsistency, or a global invariant violation):

1. The prior topological decomposition $\mathfrak{D}(s)$ is dissolved.
2. A new, statistically independent randomization of the structural cut is generated via the cryptographic space Ξ .
3. The terminal fragments are completely redefined, generating new independent measurement tasks.
4. The new tasks are reassigned randomly across the verifier set $\mathcal{V}$.

Consequently, any I_{metadata} acquired or I_{payload} inferred by the adversary during a failed state transaction does not carry over directly to the next randomized transition attempt. The combinatorial entropy of the system is regenerated under a new randomized configuration at every new cycle, thereby limiting iterative metadata accumulation.

Remark 11.2 (Epistemic Isolation of Global Invariants and Survival of Out-of-Band Collusions)

In the case of prolonged adversarial collusion, it must be assumed that identified malicious nodes can maintain out-of-band (external to the system) communication channels with still-active colluding nodes.

Although Topological and Semantic Refresh precludes internal structural inference, the adversary could theoretically accumulate metadata and payloads historically inferred in previous cycles in an external ledger. However, the ultimate security of the CNVS is guaranteed by the epistemic isolation of Global Invariants (C).

As stated by Axiom III (**Non-Reducibility**), the global invariance conditions $\text{Inv}_c(S(t))$ are applied exclusively during global state reconstruction and are never projected, fragmented, or encoded within the payloads of terminal nodes ($\text{Ter}(t)$) or their measurement grids.

Therefore, even in the theoretical singularity where a colluding network manages to infer all local payloads, the adversary would remain constitutionally blind to systemic validation rules (e.g., macroscopic constraints, sum bounds, global parity conditions).

The lack of traces of the invariants within the measurement task makes the endorsement of a malicious state undecidable for the adversary, relegating the success of an attack to a mere probability of random collision with the hidden invariants of the system.

12. Fundamental Lemmas

The robustness of the CNVS framework is supported by the following lemmas, derived from the axioms and implications defined in the previous sections.

- **Lemma 1 (Terminal Membership):** $\forall s \in \text{Ter}(t), s \in E(t) \subseteq \mathcal{S}$ (Existence is primitive).
- **Lemma 2 (Global Necessity):** $\exists S(t) \in \mathcal{S} : (\forall s \in \text{Ter}(t), \text{Adm}_t(s)=1) \wedge V_G(S(t))=0$ (Proves that local validity does not imply global consistency).
- **Lemma 3 (Type Compatible Decomposition):**
 $\forall s \in \mathcal{S}$ and $\forall f \in \mathcal{D}(s)$, the generated fragment remains by its decomposition a well-typed element of the structural universe and its native type belongs to the admissible evaluative spectrum of the parent type: $\forall s \in \mathcal{S}, \forall f \in \mathcal{D}(s), f \in \mathcal{S} \wedge \text{type}(f) \in \Omega(\text{type}(s))$.
Consequently, recursive decomposition does not erase type structure. It preserves formal well-typedness while allowing heterogeneous local evaluative fragments.
- **Lemma 4 (Internal Payload Accounting):**

For every decomposable structural element $s \in \mathcal{S}$ with $\mathcal{D}(s) \neq \emptyset$, the internal semantic payload associated with s is distributed through the recursive decomposition of the sub-terminal units subject to partial global verification, up to the terminal evaluative units subject to local verification.

Formally: $I_{\text{payload}}(s) = \sum_{f_i \in \mathcal{D}(s)} I_{\text{payload}}(f_i)$.

This equation is an internal accounting identity. It does not imply that external verifiers receive semantic payload shares. In ordinary CNVS operation, the verifier receives only the measurement grid or task specification associated with f_i and independently produces obs_i . The internal payload d_i remains inaccessible to the verifier.

Therefore, internal payload conservation is compatible with external knowledge restriction.

- **Lemma 5 (Metadata Overhead):** $\forall s \in \mathcal{S}$ such that $\mathcal{D}(s) \neq \emptyset$, $I_{\text{metadata}}(\text{Rec}(\mathcal{D}(s))) < \sum_{f_i \in \mathcal{D}(s)} I_{\text{metadata}}(f_i)$.
(Fragmentation increases systemic metadata costs).

- **Lemma 6 (Asymptotic Epistemic Negligibility):**

Let $\mathcal{K}(v, f)$ denote the verifier-accessible information associated with a terminal fragment f , and let $I_{\text{global}}(S(t))$ denote the total informational content of the system state.

Let $m(t)$ denote the dynamic structural fragmentation scale of the system at time t . If the verifier-accessible information remains locally bounded while the global informational structure increases with $m(t)$, then:

$$\lim_{m(t) \rightarrow \infty} \mathcal{K}(v, f) / I_{\text{global}}(S(t)) = 0.$$

- **Lemma 7 (Residual Structural Inference Bound):** see Section 17a for the full statement and proof.
- **Lemma 8 (Hidden-Invariant Residual Non-Inferability Bound):**

To preclude deterministic algebraic reconstruction, the CNVS framework does not assume that an adversarial coalition has access to the instantiated global validation constraints.

An adversary may know the general public class of admissible invariant forms, denoted by C_{pub} , but does not know the internal instantiated invariant parameters θ_c , the complete relational topology R_{int} , or the specific invariant binding that maps terminal fragments into global state validation equations.

Let

$$C_{\text{int}} \triangleq \langle \theta_c, R_{\text{int}} \rangle \tag{21}$$

where:

θ_c = hidden internal invariant parameters

R_{int} = hidden internal relational topology.

Let f_{miss} denote a critical terminal fragment not directly controlled by the adversarial coalition, and let d_{miss} be its true internal semantic payload.

Let W_k denote the complete worst-case semantic information acquired from compromised terminal fragments, including the local payloads pessimistically granted to colluding verifiers in the adversarial model of Section 15.

Let M_s denote the observable topological metadata, subject to the bounded topological leakage condition defined in Lemma 10.

For the purpose of this residual non-inferability bound, the relevant adversarial view may be represented by the tuple of components:

$$\text{View}_{\text{adv}} \triangleq (W_k, M_s, C_{\text{pub}}). \quad (22)$$

The CNVS residual non-inferability condition requires that the adversarial view does not collapse the conditional min-entropy of the missing payload. There must exist a strictly positive residual entropy margin $h_{\text{min}} > 0$ such that:

$$H_{\infty}(d_{\text{miss}} \mid \text{View}_{\text{adv}}) \geq h_{\text{min}}. \quad (23)$$

Equivalently, the optimal adversarial guessing probability is bounded by:

$$\mathbb{P}_{\text{guess}}(d_{\text{miss}} \mid W_k, M_s, C_{\text{pub}}) \leq 2^{-h_{\text{min}}}. \quad (24)$$

This condition means that the adversary may possess many compromised local payloads and may know an abstract schema-level class of admissible validation forms, but still lacks the hidden instantiated constraint structure C_{int} required to deterministically infer the missing critical payload (d_{miss}).

The rigorous security condition is therefore not that the adversary observes fewer equations than unknowns in a purely syntactic linear-algebraic sense. Rather, the condition is that the adversarial view (View_{adv}) lacks sufficient information to reduce the residual min-entropy of the missing critical payload to zero ($H_{\infty}(d_{\text{miss}} \mid \text{View}_{\text{adv}}) > 0$).

Since conditional min-entropy bounds the optimal guessing probability, the inference probability for the missing critical fragment satisfies:

$$p_{\text{inf}} \leq 2^{-h_{\text{min}}}. \quad (25)$$

Consequently, if $h_{\text{min}} > 0$, then:

$$p_{\text{inf}} < 1.$$

If h_{min} is large, then p_{inf} becomes exponentially small. This provides the information-theoretic basis for the dependent-collusion model, in which the composite compromise probability remains strictly sub-unitary:

$$p_{\text{comp}} = q + (1 - q) p_{\text{inf}} < 1,$$

provided $q < 1$.

- **Lemma 9 (Decidability Constraint of Global Validation):**

To prevent systemic undecidability and avoid non-terminating validation procedures, the evaluative functions governing the Global Veto are formally restricted to a bounded validation language.

Let L_V denote the formal validation language used to express the systemic structural invariants C and the relational consistency checks Cons_R evaluated by V_G . The language L_V admits only total computable predicates with polynomially bounded evaluation. In particular, L_V excludes unbounded recursion, non-terminating procedures, and validation loops whose maximum decomposition cardinality is not structurally bounded by the finite state representation or by the hierarchical the minimum decomposition cardinality n of the state.

For every candidate state $S' = \langle E', R', C \rangle \in \mathfrak{S}$, assume that:

1. $\text{Cons}_R(R', E')$ is expressed in L_V ;
2. every invariant predicate $c_i \in C$ is expressed in L_V ;
3. all quantifications range only over the finite sets E', R' , and C ;
4. every iterative validation loop is bounded by the finite cardinalities of the state components or by the hierarchical the minimum decomposition cardinality n .

Then the global verification function

$$V_G(S') = 1 \Leftrightarrow (\forall s \in \text{Ter}(S'), \pi_1(V_L(s)) = 1) \wedge \text{Cons}_R(R', E') \wedge (\forall c_i \in C, c_i(E', R') = 1)$$

is decidable. Moreover, its evaluation time is bounded by a polynomial in the encoded size of the candidate state S' . If the encoded size of S' is polynomially bounded by the hierarchical the minimum decomposition cardinality n , then there exists a constant exponent $\alpha \in \mathbb{N}$ such that:

$$\text{Time}(V_G(S')) \in O(n^\alpha).$$

Proof. The sets E', R' , and C are finite by definition of the CNVS state space. The evaluation of V_G is a finite conjunction of: local verification predicates over terminal elements, the relational consistency predicate Cons_R , and the invariant predicates $c_i \in C$. By assumption, each of these predicates is expressed in L_V and therefore is total and polynomial-time bounded. Since finite sums and finite compositions of polynomial-time computations remain polynomial-time, the complete evaluation of V_G halts and returns a binary value in polynomial time. Therefore, for every $S' \in \mathfrak{S}$:

$$V_G(S') \in \{0, 1\}.$$

Proof Consequence. Because V_G is decidable under the L_V restriction, the State Transition Law introduced in Section 11 is strictly decidable. Each candidate transition is therefore completed or rejected in polynomial-time.

- **Lemma 10 (Bounded Topological Leakage and Entropy Separation):**

To prevent the theoretical vulnerability of entropic overlap — whereby the structural shape of the routing topology may reveal information about the native semantic payload — the CNVS framework does not require perfect semantic independence between topology and payload. Instead, it imposes a bounded-leakage condition.

Let X_S be the random variable representing the native semantic payload of the system state S , and let M_S be the random variable representing the observable structural metadata of the CNVS configuration, including routing topology, fragment assignment structure, and coordination metadata.

The CNVS topology is said to satisfy γ -topological leakage if: $I(X_S ; M_S) \leq \gamma_{\text{top}}$,

where $\gamma_{\text{top}} \geq 0$ is a bounded leakage parameter. Equivalently, the residual uncertainty of the semantic payload after observing the topological metadata satisfies:

$$H(X_S | M_S) \geq H(X_S) - \gamma_{\text{top}} \geq h_{\min} > 0, \quad (26)$$

therefore:

$$\gamma_{\text{top}} \leq H(X_S) - h_{\min}.$$

Thus, the routing topology is not required to be perfectly independent from the semantic payload. It may reveal limited structural information. However, such information remains quantitatively bounded and does not suffice for deterministic reconstruction.

A CNVS configuration is semantically non-reconstructive if there exists a strictly positive residual entropy margin $h_{\min} > 0$ such that $H(X_S | M_S, W_k) \geq h_{\min}$, (27)

where W_k denotes the fragment-level information acquired by the adversary, as defined in the adversarial model of Section 15.

This condition means that even after observing the available topology and the compromised local information, the adversary still faces irreducible uncertainty about the native semantic payload.

Proof Consequence. Under γ -topological leakage, the structural map may leak limited metadata about the organization of the verification process, but it does not collapse the semantic uncertainty of the payload. Therefore, the adversary cannot deterministically reconstruct the global state unless the residual entropy margin $h_{\min}$ vanishes.

In the idealized limiting case $\gamma_{\text{top}} = 0$, the model reduces to perfect semantic-agnostic topology: $I(X_S ; M_S) = 0$.

However, the general CNVS security model only requires the weaker and more realistic bounded-leakage condition: $I(X_s; M_s) \leq \gamma_{\text{top}}$, together with positive residual uncertainty: $H(X_s | M_s, W_k) \geq h_{\text{min}} > 0$.

13. Principle of Knowledge Restriction

The CNVS framework does not assume Knowledge Restriction as a passive, static property, but enforces it as a structural consequence of recursive decomposition, local task restriction, randomized assignment, and hidden global binding.

In the present compact formal paper the security argument only requires the following abstract condition: as the system expands, the verifier-accessible knowledge remains negligible relative to the total informational structure of the global state.

Let $m(t)$ denote the active structural fragmentation scale of the system at time t . As $m(t)$ increases, the verifier-accessible knowledge becomes strictly negligible relative to the global informational structure. Thus, we define epistemic restriction through the following asymptotic condition:

$$\lim_{m(t) \rightarrow \infty} \mathcal{K}(v, f) / I_{\text{global}}(S(t)) = 0. \quad (28)$$

This condition states that the information available to any verifier about a local terminal fragment remains asymptotically negligible with respect to the total informational content of the system state.

To translate this structural limit into information-theoretic terms, let X_s be the random variable representing the true internal semantic payload of a node s , and let View_v be the random variable representing the external observable information actually available to the verifier v .

The verifier's information about the native semantic content is asymptotically bounded by a vanishing sequence $\varepsilon_m \rightarrow 0$ as $m(t) \rightarrow \infty$:

$$I(X_s; \text{View}_v) \leq \varepsilon_m. \quad (29)$$

Equivalently, the conditional entropy of the internal payload, given the verifier's knowledge, remains close to maximal uncertainty:

$$H(X_s | \text{View}_v) \geq H(X_s) - \varepsilon_m. \quad (30)$$

Therefore, knowledge restriction is not treated here as a consequence of a specific density-control formula, but as a formal asymptotic condition enforced by CNVS decomposition, assignment, and verification architecture.

As anticipated in the previous sections, the probabilistic security analysis developed in the following sections intentionally adopts a strictly pessimistic adversarial convention, where $\text{View}_v = d_i$.

Therefore, the weighted adversarial information W_k analyzed below is not the ordinary information exposed to verifiers (honest or dishonest) during normal operation. It is a semantic exposure variable under the worst-case assumption. Any security bound tested under this pessimistic assumption also represents an upper bound on the risk of the real operating model in which d_i remains hidden and only obs_i is externally available.

14. Probabilistic Security Model

We formalize the security bounds against an adversarial coalition attempting unauthorized state reconstruction under a pessimistic semantic exposure model and using the assignment framework of Section 7.

Independence is an analytical baseline, not the CNVS operating model.

Let:

- $\mathcal{V}$ be the finite set of total verifiers, with $Q_{\mathcal{V}} = |\mathcal{V}|$;
- $\mathcal{V}_{\text{adv}} \subset \mathcal{V}$ be the adversarial colluding coalition;
- $k = |\text{Ter}(t)|$ be the total number of active terminal fragments;
- $q = |\mathcal{V}_{\text{adv}}| / Q_{\mathcal{V}}$ denotes the common marginal compromise probability under the full-pool uniform injective baseline;
- d_i be the internal declared payload associated with terminal fragment f_i ;
- $I_{\text{payload}}(f_i)$ be the semantic information content of d_i .

Although d_i is not disclosed to the verifier in the operational CNVS protocol, the present probabilistic model assumes, pessimistically, that an adversarial verifier assigned to f_i obtains full access to d_i .

Let $I_{\text{ref}} \in \mathbb{R}_{>0}$ be the semantic capacity of reference, defined rigorously as the maximum value assumed in the set:

$$I_{\text{ref}} = \max_{f_i \in \text{Ter}(t)} I_{\text{payload}}(f_i).$$

Consequently, every terminal fragment satisfies the condition:

$$0 < I_{\text{payload}}(f_i) \leq I_{\text{ref}}.$$

We define the normalized local semantic weight of fragment f_i as:

$$\omega_i \triangleq I_{\text{payload}}(f_i) / I_{\text{ref}}. \quad (31)$$

Hence:

$$0 < \omega_i \leq 1.$$

Let $Y_i \in \{0, 1\}$ be the adversarial acquisition indicator:

$Y_i = 1$ if f_i is assigned to a colluding verifier (i.e., $\mathcal{A}_t(f_i, \xi) \in \mathcal{V}_{\text{adv}}$),

$Y_i = 0$ otherwise.

Let $q_i \triangleq \mathbb{P}(Y_i = 1)$

denote the corresponding marginal compromise probability.

Using the marginal assignment probabilities defined in Section 7, this becomes:

$$q_i = \sum_{v \in \mathcal{V}_{\text{adv}}} \mu_t(v, f_i).$$

In the full-pool uniform injective case, if $|\mathcal{V}| = Q_{\mathcal{V}}$ and $|\mathcal{V}_{\text{adv}}| = r$, then: $q_i = r / Q_{\mathcal{V}}$.

If q denotes this common marginal compromise probability, then:

$$q = |\mathcal{V}_{\text{adv}}| / |\mathcal{V}| = r / Q_{\mathcal{V}}, \quad (32)$$

Thus, q is only the common compromise probability in the uniform marginal case.

Under injective assignment, the variables $Y_1, \dots, Y_k$ are generally not mutually independent. Therefore, the independent-assignment Hoeffding model should be understood as a simplified analytical baseline, while the exact operational assignment mechanism is the randomized injective matching defined in Section 7.

Under the independent randomized assignment assumption, the variables Y_i are independent Bernoulli random variables with:

$$\mathbb{E}[Y_i] = q.$$

The cumulative worst-case adversarial semantic acquisition is therefore:

$$W_k \triangleq \sum_{i=1}^k Y_i \omega_i. \quad (33)$$

Let:

$$\bar{\omega}_k \triangleq (1/k) \sum_{i=1}^k \omega_i$$

be the average normalized semantic weight of the terminal fragments.

Then:

$$\mathbb{E}[W_k] = \sum_{i=1}^k \mathbb{E}[Y_i] \omega_i = q \sum_{i=1}^k \omega_i = q k \bar{\omega}_k. \quad (34)$$

Unauthorized systemic reconstruction is modeled as a threshold event over the extensive worst-case semantic acquisition variable W_k . Let $\tau \in [0, 1]$ denote the critical normalized semantic acquisition density required for reconstruction. The adversarial reconstruction event is:

$$W_k \geq \tau k.$$

Thus, the conditional security regime is:

$$q \bar{\omega}_k k < \tau k.$$

This formulation is dimensionally coherent because W_k and τk are both extensive quantities measured in units of normalized terminal semantic capacity. It also preserves the pessimistic nature of the model: the adversary is assumed to acquire d_i whenever it controls the assigned verifier, even though d_i is not externally disclosed in the operational CNVS protocol.

In Sections 15–17, we restrict the analysis to the uniform independent-assignment baseline, where $q_i = q$ for all terminal fragments.

The general dependent-collusion model is addressed in Section 17a.

15. Conditional Security Theorem

We establish the foundational condition under which adversarial reconstruction is probabilistically bounded in the independent-assignment baseline.

Let the absolute reconstruction threshold be defined as:

$$\mathbb{T} \triangleq \tau k,$$

where $\tau \in [0, 1]$ represents the critical normalized semantic acquisition density required to compromise the state.

Theorem 1. Conditional Worst-Case Semantic Security

Let W_k be the cumulative worst-case adversarial semantic acquisition:

$$W_k = \sum_{i=1}^k Y_i \omega_i,$$

where each $\omega_i \in [0,1]$ is the semantic weight normalized on I_{ref} of the internal payload d_i associated with the terminal fragment f_i .

Let

$$\bar{\omega}_k \triangleq (1/k) \sum_{i=1}^k \omega_i$$

be the average normalized semantic weight.

If:

$$q \bar{\omega}_k < \tau,$$

then the expected worst-case semantic information acquired by the colluding coalition remains strictly below the reconstruction threshold τk .

Proof

By linearity of expectation:

$$\begin{aligned} \mathbb{E}[W_k] &= \mathbb{E}[\sum_{i=1}^k Y_i \omega_i] \\ &= \sum_{i=1}^k \mathbb{E}[Y_i] \omega_i \\ &= q \sum_{i=1}^k \omega_i \\ &= q k \bar{\omega}_k. \end{aligned}$$

By the theorem hypothesis:

$$q \bar{\omega}_k < \tau.$$

Multiplying both sides by k gives:

$$q k \bar{\omega}_k < \tau k.$$

Therefore:

$$\mathbb{E}[W_k] < \tau k. \tag{35}$$

Hence, even under the pessimistic assumption that every colluding verifier obtains the full internal payload d_i of its assigned fragment, the expected adversarial semantic acquisition remains below the reconstruction threshold.

16. Exponential Concentration Bound

The following concentration estimate is a baseline result under the simplified independent-assignment model. It assumes that terminal assignments are independent and that adversarial acquisition occurs through direct assignment to colluding verifiers.

Theorem 2. Baseline Hoeffding Concentration Bound

Let:

$$W_k = \sum_{i=1}^k Y_i \omega_i,$$

where each Y_i is an independent Bernoulli random variable with parameter q , and each weighted acquisition variable $Y_i \omega_i$ is bounded in the interval $[0, \omega_i]$.

If the worst-case semantic acquisition density satisfies:

$$q \bar{\omega}_k < \tau,$$

then:

$$\mathbb{P}(W_k \geq \tau k) \leq \exp(-2 k^2 (\tau - q \bar{\omega}_k)^2 / \sum_{i=1}^k \omega_i^2).$$

In particular, since $\omega_i \leq 1$ for all i , we can maximize the denominator to obtain a conservative (worst-case) upper bound:

$$\sum_{i=1}^k \omega_i^2 \leq k,$$

and therefore substituting this limit in the exponent, we obtain the generalized estimate:

$$\mathbb{P}(W_k \geq \tau k) \leq \exp(-2k(\tau - q \bar{\omega}_k)^2). \quad (36)$$

This theorem is a baseline concentration result. It proves exponential suppression under independent randomized assignment, but it does not claim that all realistic CNVS compromise processes are independent.

Proof

We write:

$$\begin{aligned} & \mathbb{P}(W_k \geq \tau k) \\ &= \mathbb{P}(W_k - \mathbb{E}[W_k] \geq \tau k - q k \bar{\omega}_k). \end{aligned}$$

Since:

$$\tau k - q k \bar{\omega}_k = k(\tau - q \bar{\omega}_k),$$

and since $\tau - q \bar{\omega}_k > 0$ by hypothesis, Hoeffding's inequality yields:

$$\mathbb{P}(W_k \geq \tau k) \leq \exp(-2 k^2 (\tau - q \bar{\omega}_k)^2 / \sum_{i=1}^k \omega_i^2).$$

Because each $\omega_i \leq 1$, we have:

$$\sum_{i=1}^k \omega_i^2 \leq k.$$

Therefore:

$$\mathbb{P}(W_k \geq \tau k) \leq \exp(-2k (\tau - q \bar{\omega}_k)^2).$$

This proves exponential decay in k whenever the positive security margin $\tau - q \bar{\omega}_k$ is bounded away from zero.

17. Emergent Security Scaling Theorem

We now state the asymptotic consequence of the baseline concentration bound.

Theorem 3. Asymptotic Emergent Security under Worst-Case Payload Exposure

Let $k(\lambda)$ be the number of terminal fragments generated at fragmentation scale λ , with:

$$k(\lambda) \rightarrow \infty \text{ as } \lambda \rightarrow \infty.$$

Assume:

1. in the independent-assignment analytical baseline, the acquisition variables Y_i are independent;
2. each colluding verifier is pessimistically assumed to obtain the full internal payload d_i of each assigned terminal fragment f_i ;
3. the normalized semantic weights satisfy $0 < \omega_i \leq 1$;
4. there exists a fixed positive security margin $\delta > 0$ such that:

$$\tau - q \bar{\omega}_{k(\lambda)} \geq \delta$$

for all sufficiently large λ .

Then:

$$\lim_{\lambda \rightarrow \infty} \mathbb{P}(W_{k(\lambda)} \geq \tau k(\lambda)) = 0.$$

Moreover:

$$\mathbb{P}(W_{k(\lambda)} \geq \tau k(\lambda)) \leq \exp(-2\delta^2 k(\lambda)). \tag{37}$$

Proof

From Theorem 2:

$$\mathbb{P}(W_k \geq \tau k) \leq \exp(-2k (\tau - q \bar{\omega}_k)^2).$$

By assumption, for sufficiently large $k = k(\lambda)$:

$$\tau - q \bar{\omega}_k \geq \delta > 0.$$

Therefore:

$$\mathbb{P}(W_{k(\lambda)} \geq \tau k(\lambda)) \leq \exp(-2\delta^2 k(\lambda)).$$

Since $k(\lambda) \rightarrow \infty$ as $\lambda \rightarrow \infty$, the exponent diverges to $-\infty$:

$$-2\delta^2 k(\lambda) \rightarrow -\infty.$$

Consequently:

$$\lim_{\lambda \rightarrow \infty} \mathbb{P}(W_{k(\lambda)} \geq \tau k(\lambda)) = 0.$$

Thus, even under the pessimistic semantic exposure convention in which colluding verifiers are assumed to obtain the internal payload d_i of every assigned fragment, unauthorized reconstruction probability converges to zero under the independent-assignment analytical baseline and a fixed positive semantic security margin.

Transition to the Generalized Dependent-Collusion Model

The independent Hoeffding baseline captures the idealized case in which fragment assignments are independent and adversarial acquisition occurs only through direct assignment to colluding verifiers.

However, real adversarial coalitions may attempt to exploit structural inference, partial metadata leakage, repeated observations, adaptive coordination, or correlations between compromised fragments. In such cases, the compromise events are no longer adequately modeled as independent Bernoulli acquisitions.

For this reason, the CNVS security analysis does not rely exclusively on the Hoeffding baseline.

The independent model is generalized by introducing a composite compromise probability:

$$p_{\text{comp}} = q + (1 - q)p_{\text{inf}},$$

where q represents direct assignment to colluding verifiers and p_{inf} represents the bounded probability of inferring an uncompromised critical fragment from already acquired information.

The dependent-collusion model replaces independence with a conditional upper bound:

$$\mathbb{P}(c_{\text{coll},i} \mid c_{\text{coll},1}, \dots, c_{\text{coll},i-1}) \leq p_{\text{comp}} < 1,$$

where $c_{\text{coll},i}$ denotes the adversarial event that the i -th critical fragment is successfully compromised.

Thus, the essential CNVS security requirement is not strict independence of all compromise events, but the weaker and more general condition that each additional critical fragment remains non-deterministically inferable with probability bounded away from one.

17a. Generalized Emergent Security Scaling under Dependent Collusion

The independent-assignment concentration analysis (Sections 15-17) provides an idealized asymptotic bound. We now generalize the model to evaluate dependent compromise events, where an adversarial coalition attempts cryptanalytic inference across multiple fragments.

Let:

- $q \in [0, 1)$ denotes the probability that a critical fragment is directly assigned to colluding verifiers under the uniform marginal baseline defined in Section 14.
- View_{adv} denotes the complete adversarial view available to the colluding coalition, including directly compromised fragments, observable metadata, public structural information, and any information accumulated through previous compromise history.
- $p_{\text{inf}} \in [0, 1)$ denotes the probability of structural inference — the likelihood that a fragment not directly controlled can be compromised via cryptanalytic derivation from already compromised fragments.
- $m \in \mathbb{N}$ denotes the strictly positive number of critical low-inference fragments required for unauthorized systemic reconstruction.

We formally define the single-fragment composite compromise probability p_{comp} via the Law of Total Probability:

$$p_{\text{comp}} \triangleq q + (1 - q) p_{\text{inf}} \tag{38}$$

To formalize the conditional dependence between compromise events, we establish the following bounding lemma.

Lemma 7. Residual Structural Inference Bound:

A critical fragment may be compromised in two distinct ways:

1. by direct assignment to a colluding verifier;

2. by structural inference from previously compromised fragments, observable metadata, and abstract schema-level knowledge about the admissible validation framework.

In the uniform independent-assignment baseline, the probability of direct assignment to the adversarial coalition is q .

If the i -th critical fragment is not directly assigned to the adversarial coalition, then its compromise can occur only through inference. By Lemma 8, every non-controlled critical fragment retains a strictly positive residual conditional min-entropy margin $h_{\min} > 0$ under the adversarial view. Therefore, the optimal inference probability is bounded by:

$$p_{\text{inf}} \leq 2^{-h_{\min}}.$$

Consequently, the conditional probability of compromising the i -th critical fragment, given any previous compromise history, satisfies:

$$\mathbb{P}(c_{\text{coll}_i} \mid c_{\text{coll}_1}, \dots, c_{\text{coll}_{i-1}}) \leq q + (1 - q)p_{\text{inf}}.$$

Substituting the definition given in Equation (38), we get:

$$\mathbb{P}(c_{\text{coll}_i} \mid c_{\text{coll}_1}, \dots, c_{\text{coll}_{i-1}}) \leq p_{\text{comp}}.$$

Since $p_{\text{inf}} \leq 2^{-h_{\min}}$, we obtain:

$$p_{\text{comp}} \leq q + (1 - q)2^{-h_{\min}}.$$

Therefore, if $q < 1$ and $h_{\min} > 0$, then:

$$p_{\text{comp}} < 1.$$

This establishes the conditional upper bound:

$$\mathbb{P}(c_{\text{coll}_i} \mid c_{\text{coll}_1}, \dots, c_{\text{coll}_{i-1}}) \leq p_{\text{comp}} < 1.$$

Thus, the dependent-collusion model does not require independence between compromise events. It requires only that each additional critical fragment remains non-deterministically inferable with probability bounded strictly below one.

Therefore for each additional critical fragment, the residual min-entropy condition is required to hold after incorporating into the adversarial view all previously compromised fragments, observable metadata, and any abstract schema-level knowledge available about the admissible validation language. Thus, previous compromise history may reduce residual uncertainty, but it must not collapse the conditional min-entropy of the next uncontrolled critical fragment to zero. The instantiated invariant parameters, the internal relational topology, and the specific invariant binding remain hidden inside the CNVS system.

Interpretive Note.

In Theorem 4, p_{comp} is used as a uniform conditional upper bound over the adversarial regime under consideration. The parameter p_{inf} need not be interpreted as a static fixed

scalar; it may represent the worst-case upper bound of an inference probability depending on the adversarial view accumulated so far. Equivalently, one may read p_{comp} as satisfying

$$p_{\text{comp}} \geq \sup_i \mathbb{P}(c_{\text{coll}_i} \mid c_{\text{coll}_1}, \dots, c_{\text{coll}_{i-1}}),$$

with $p_{\text{comp}} < 1$.

The dynamic entropy-driven formulation of Remark 17a.2 generalizes this static upper-bound presentation by allowing p_{inf} and p_{comp} to vary with time or accumulated adversarial knowledge.

Theorem 4 (Generalized Emergent Security Scaling)

Under the constraints of the Residual Structural Inference Bound, the probability of an unauthorized full state reconstruction, denoted as Rec^* , satisfies the exponential decay inequality:

$$\mathbb{P}(\text{Rec}^*) \leq (p_{\text{comp}})^m \quad (39)$$

Proof

Unauthorized full reconstruction strictly requires the simultaneous compromise of all m critical low-inference fragments. In set-theoretic terms, the reconstruction event is a subset of the intersection of all compromise events:

$$\text{Rec}^* \subseteq \bigcap_{i=1}^m c_{\text{coll}_i}$$

Therefore, by the monotonicity of probability measures:

$$\mathbb{P}(\text{Rec}^*) \leq \mathbb{P}(\bigcap_{i=1}^m c_{\text{coll}_i})$$

Applying the general chain rule of probability:

$$\mathbb{P}(\bigcap_{i=1}^m c_{\text{coll}_i}) = \prod_{i=1}^m \mathbb{P}(c_{\text{coll}_i} \mid c_{\text{coll}_1}, \dots, c_{\text{coll}_{i-1}})$$

Substituting the conditional upper bound established in Lemma 7 into the product series:

$$\mathbb{P}(\text{Rec}^*) \leq \prod_{i=1}^m p_{\text{comp}} = p_{\text{comp}}^m \quad (40)$$

This demonstrates that systemic security scales exponentially against dependent collusion, provided the underlying composite probability remains strictly sub-unitary.

Corollary 4.1 (Asymptotic Dependent Security)

Because $0 \leq p_{\text{comp}} < 1$, the geometric sequence $(p_{\text{comp}})^m$ converges to zero. Therefore, as the system scale expands and the requisite critical fragments increase:

$$\lim_{m \rightarrow \infty} \mathbb{P}(\text{Rec}^*) = 0$$

Corollary 4.2 (Non-Inferable Limit)

In the idealized non-inferable case $p_{\text{inf}} = 0$, the composite compromise probability reduces to $p_{\text{comp}} = q$.

The generalized security bound then reduces to the idealized direct-assignment boundary:

$$\mathbb{P}(\text{Rec}^*) \leq q^m$$

Systemic Design Formula: Minimum Critical Fragmentation

To operationalize the security model, the framework must determine the minimum decomposition cardinality required to satisfy a specific security target. Given a maximum tolerable target attack probability $\eta \in (0, 1)$, we seek the minimum number of critical low-inference fragments m_{min} such that:

$$\mathbb{P}(\text{Rec}^*) \leq \eta \tag{41}$$

By substituting the exponential bound derived in Theorem 4:

$$p_{\text{comp}}^m \leq \eta$$

Taking the natural logarithm of both sides:

$$m \ln(p_{\text{comp}}) \leq \ln(\eta)$$

Since $\eta \in (0, 1)$ and $p_{\text{comp}} \in (0, 1)$, both logarithmic terms are strictly negative. Consequently, dividing both sides by $\ln(p_{\text{comp}})$ reverses the inequality direction:

$$m \geq \ln(\eta) / \ln(p_{\text{comp}})$$

Because the fragmentation cardinality m must be a discrete natural number ($m \in \mathbb{N}$), we define the exact threshold using the ceiling function:

$$m_{\text{min}} = \lceil \ln(\eta) / \ln(p_{\text{comp}}) \rceil \tag{42}$$

This formula deterministically translates any theoretical cryptographic security requirement (η) into a precise topological constraint (m_{min}).

Remark 17a.1 (Semantic lower bound and practical scaling)

The asymptotic expression $m \rightarrow \infty$ utilized in Corollary 4.1 represents an idealized mathematical limit. In practical CNVS deployments, fragmentation cannot proceed indefinitely.

As formalized in Section 6, terminal fragments must preserve a sufficient semantic payload to allow honest external verifiers to perform their assigned local evaluation tasks without logical undecidability.

Therefore, the effective fragmentation cardinality is strictly constrained by the semantic lower bound:

$$\forall f \in \text{Ter}(t), I_{\text{payload}}(f) \geq I_{\min} \quad (43)$$

where $I_{\min}$ denotes the absolute minimum informational quantum required for task comprehension.

Consequently, the practical system design problem is not to maximize fragmentation to infinity, but to solve a Bounded Optimization Problem:

Maximize epistemic restriction (by increasing m) subject to the constraint of local semantic verifiability.

Let $m_{\max}$ denote the theoretical maximum number of fragments derivable from a root element s before the payload of any resulting terminal fragment drops below $I_{\min}$.

A specific CNVS configuration is considered structurally and cryptographically feasible if and only if the maximum semantically permissible fragmentation equals or exceeds the minimum security requirement:

$$m_{\max} \geq m_{\min} \quad (44)$$

This feasibility condition formally guarantees that the system can achieve the target adversarial resistance η without sacrificing the operational computability of the global state.

Remark 17a.2 (Entropic Interpretation of Inference Probability)

The structural inference probability p_{inf} may increase over time as the adversarial coalition accumulates compromised fragment-level information, observable metadata, and partial structural correlations.

Let $H(S(t_0))$ denote the initial maximum Shannon entropy of the global state space prior to adversarial observation. Let $H(S(t) | W(t))$ denote the residual conditional entropy of the global state at time t , given the cumulative adversarial information $W(t)$.

Assume $H(S(t_0)) > 0$ and $0 \leq \mathbb{E}[H(S(t) | W(t))] \leq H(S(t_0))$. We define the normalized residual entropy ratio:

$$r(t) \triangleq \mathbb{E}[H(S(t) | W(t))] / H(S(t_0)).$$

Thus:

$$0 \leq r(t) \leq 1.$$

A simple entropy-driven inference model may define the dynamic inference probability as:

$$p_{\text{inf}}(t) \triangleq 1 - r(t).$$

Equivalently:

$$p_{\text{inf}}(t) \triangleq 1 - \mathbb{E}[H(S(t) \mid W(t))] / H(S(t_0)).$$

Under this model, the dynamic composite compromise probability becomes:

$$p_{\text{comp}}(t) \triangleq q + (1 - q) p_{\text{inf}}(t).$$

Substituting the entropy-driven expression for $p_{\text{inf}}(t)$, we obtain:

$$p_{\text{comp}}(t) = q + (1 - q)[1 - \mathbb{E}[H(S(t) \mid W(t))] / H(S(t_0))].$$

Equivalently:

$$p_{\text{comp}}(t) = 1 - (1 - q) \cdot \mathbb{E}[H(S(t) \mid W(t))] / H(S(t_0)).$$

The dependent-collusion bound therefore becomes time-dependent:

$$\mathbb{P}(\text{Rec}^*) \leq p_{\text{comp}}(t)^{m(t)}.$$

This does not imply unconditional dominance of fragmentation over entropy erosion. If $p_{\text{comp}}(t)$ approaches 1 too rapidly, the exponential bound may fail to converge to zero.

A sufficient asymptotic condition for security is:

$$m(t) (-\ln p_{\text{comp}}(t)) \rightarrow \infty.$$

Equivalently, if we define the residual compromise gap:

$$\delta_p(t) \triangleq 1 - p_{\text{comp}}(t),$$

then a sufficient small-gap condition is:

$$m(t) \delta_p(t) \rightarrow \infty.$$

In the entropy-driven model above, this becomes:

$$m(t) (1 - q) \cdot \mathbb{E}[H(S(t) \mid W(t))] / H(S(t_0)) \rightarrow \infty.$$

Therefore, fragmentation suppresses dependent reconstruction only when the growth of the critical-fragment cardinality $m(t)$ dominates the rate at which residual entropy is depleted by adversarial knowledge accumulation.

This formulation preserves the thermodynamic interpretation while making the asymptotic security claim conditional, rate-aware, and mathematically defensible.

Note $\delta_p(t)$ is distinct from the fixed Hoeffding security margin δ introduced in Theorem 3. While δ measures the static acquisition-margin $\tau - q \bar{\omega}_k$ in the independent baseline, $\delta_p(t)$ measures the dynamic residual gap $1 - p_{\text{comp}}(t)$ in the dependent entropy-driven regime.

Remark 17a.3 (Worst-case leakage resilience)

The emergent security scaling demonstrated in Theorems 3 and 4 remains applicable even under a worst-case local leakage scenario

Assume a boundary condition where the local convergence function (Conv_σ) fails to project a minimized semantic view, granting the colluding verifiers absolute local knowledge of their assigned fragments. Formally, the Knowledge Restriction bound introduced in Section 6 collapses locally, yielding:

$$\mathcal{K}(v, f) = I_{\text{total}}(f), \forall v \in \mathcal{V}_{\text{adv}}, \forall f \in \text{Ter}(t) \mid \mathcal{A}_t(f, \xi) = v. \quad (45)$$

Even under this worst-case assumption, the asymptotic bound of Corollary 4.1 remains applicable provided that the composite compromise probability remains strictly sub-unitary:

$$p_{\text{comp}} < 1.$$

In that case:

$$\lim_{m \rightarrow \infty} \mathbb{P}(\text{Rec}^*) = 0.$$

This shows that worst-case local payload exposure does not by itself invalidate the dependent-collusion bound, as long as the residual inference condition still prevents p_{comp} from reaching 1.

Thus, the local verification mechanism and the fragmentation topology act as distinct defensive layers. The asymptotic claim remains conditional on the preservation of a strictly sub-unitary composite compromise probability.

Theorem Implications

This result formally generalizes the idealized independent-assignment model. It shows that exponential security scaling against coordinated collusion follows under the stated assumptions, provided that the inference probability of any single fragment remains strictly bounded by a constant smaller than one.

Remark 17a.4: Why No Classical Covariance Matrix is Required

In threshold-share cryptosystems, fragments are algebraically correlated shares of a common secret. For that reason, collusion analysis may require explicit covariance or reconstruction equations between shares.

CNVS does not model terminal fragments as algebraic shares of a monolithic payload. Terminal fragments are localized evaluative units whose external exposure is limited to measurement grids, task specifications, and reported observations. Their global meaning depends on hidden invariant binding and relational topology.

Therefore, a classical covariance matrix between algebraic shares is not the appropriate object for the CNVS security model.

This does not require assuming perfect semantic independence between fragments. By Lemma 10, topology and metadata may leak bounded information. The required condition is only that this leakage remains insufficient to determine the missing semantic payload or collapse its residual conditional min-entropy.

Accordingly, inter-fragment dependence is modeled through the bounded inference parameter p_{inf} , not through algebraic covariance.

The dependent-collusion analysis remains valid as long as each missing critical fragment retains positive residual non-inferability under the adversarial view.

18. Binding Entropy and Systemic Scaling interpretation

The structural growth of the CNVS framework increases the adversarial uncertainty associated with fragment assignment, hidden invariant binding, and relational topology. These sources of uncertainty must be distinguished.

At each assignment cycle, CNVS imposes a one-fragment-per-verifier constraint. Therefore, no active verifier receives more than one terminal fragment at the same time, and no terminal fragment is assigned to more than one active verifier.

Recalling the random assignment model introduced in Section 7, let

$$\text{Ter}(t) = \{f_1, \dots, f_k\}$$

be the set of terminal fragments active at time t .

Let $\mathcal{V}$ be the finite verifier pool, with $Q = |\mathcal{V}|$, and let $\mathcal{V}_t \subseteq \mathcal{V}$ be the active verifier set selected at time t .

The assignment operator

$$\mathcal{A}_t : \text{Ter}(t) \times \Xi \rightarrow \mathcal{V}$$

is subject to the injectivity constraint:

$$\forall f_i, f_j \in \text{Ter}(t), f_i \neq f_j \Rightarrow \mathcal{A}_t(f_i, \xi) \neq \mathcal{A}_t(f_j, \xi).$$

Thus, for each random seed $\xi \in \Xi$, the map $\mathcal{A}_t(\cdot, \xi)$ is an injective matching from terminal fragments to verifiers.

If $|\mathcal{V}_t| = k$, then $\mathcal{A}_t$ induces a bijection:

$$\mathcal{A}_t(\cdot, \xi) : \text{Ter}(t) \rightarrow \mathcal{V}_t.$$

In the uniform bijective baseline, all bijections between $\text{Ter}(t)$ and $\mathcal{V}_t$ are equiprobable. The number of admissible assignments is therefore:

$$k!,$$

and the corresponding assignment entropy is:

$$H_{\text{assign}}(k) = \log_2(k!).$$

By Stirling's approximation:

$$\log_2(k!) = \Theta(k \log_2 k).$$

This factorial term arises from the injective one-fragment-per-verifier matching constraint. It does not arise from unrestricted independent assignment with replacement.

If the active verifier set $\mathcal{V}_t$ itself is selected from a larger verifier pool $\mathcal{V}$ with $|\mathcal{V}| = Q \geq k$, then the number of admissible injective assignments is:

$$Q! / (Q - k)!,$$

and the corresponding assignment entropy is:

$$H_{\text{assign}}(Q, k) = \log_2(Q! / (Q - k)!). \quad (46)$$

The CNVS framework may also use non-uniform randomized matching, for example risk-weighted, criticality-weighted, or redundancy-aware assignment. Let $\mathfrak{A}_{\text{inj}}(t)$ denote the set of admissible injective assignments at time t , and let $\mu_t(\mathbf{a})$ be the probability assigned to an injective assignment $\mathbf{a} \in \mathfrak{A}_{\text{inj}}(t)$. Then the general assignment entropy is:

$$H_{\text{assign}}(t) = - \sum_{\mathbf{a} \in \mathfrak{A}_{\text{inj}}(t)} \mu_t(\mathbf{a}) \log_2 \mu_t(\mathbf{a}). \quad (47)$$

The uniform bijective and injective cases are recovered when μ_t is uniform over the admissible assignment set.

CNVS also contains a second source of uncertainty: the hidden binding between terminal fragments and the internal invariant slots or relational positions used by global validation.

If the system internally binds k terminal fragments to k hidden validation slots through an unknown bijection, then the number of possible internal bindings is again:

$$k!,$$

with binding entropy:

$$H_{\text{bind}}(k) = \log_2(k!).$$

This binding entropy is distinct from verifier-assignment entropy. Assignment determines which verifier receives which terminal fragment. Binding determines how terminal fragments are internally connected to hidden validation slots, invariant parameters, and relational constraints.

Finally, the hidden relational graph may contribute additional topological entropy.

Let $\mathcal{G}_t$ denote the admissible family of hidden relational graphs or structural bindings generated at time t . Then:

$$H_{\text{top}}(t) = \log_2 |\mathcal{G}_t|. \quad (48)$$

The total structural uncertainty should therefore be understood as a joint entropy:

$$H_{\text{struct}}(t) = H(\mathcal{A}_t, B_t, G_t), \quad (49)$$

where $\mathcal{A}_t$ denotes the injective verifier assignment, B_t denotes the hidden invariant binding, and G_t denotes the hidden relational topology.

If these sources are generated independently, the joint entropy decomposes additively:

$$H_{\text{struct}}(t) = H_{\text{assign}}(t) + H_{\text{bind}}(t) + H_{\text{top}}(t). \quad (50)$$

Without independence, this expression must be interpreted only as an upper decomposition, and implementation-specific lower bounds must be stated explicitly.

Consequently, the correct asymptotic interpretation is not that every CNVS uncertainty source is automatically factorial. Rather, factorial growth arises from injective matching and hidden bijective binding. Additional graph-topological uncertainty may further increase the adversarial search space.

Taking the limit as the system expands, CNVS structural uncertainty diverges whenever the admissible injective assignment space, hidden binding space, or topological graph family diverges:

$$\lim_{m \rightarrow \infty} H_{\text{struct}}(t_m) = \infty. \quad (51)$$

This divergence increases the adversarial search space for unauthorized reconstruction, provided that hidden binding and topology remain unavailable or only boundedly leaked according to Lemma 10.

Therefore, combinatorial entropy supports the CNVS security model as a structural uncertainty amplifier. It does not replace the probabilistic security bounds of Sections 15–17a; rather, it supplies the combinatorial substrate that makes hidden assignment, hidden binding, and hidden topology increasingly difficult to reconstruct as the system scales.

18a. Corollary: Global Veto Property

If at least one critical fragment essential to the evaluation of a global invariant $c_i \in C$ remains outside adversarial control and retains positive residual conditional min-entropy under the adversarial view, deterministic reconstruction of a globally valid state is blocked.

Therefore, any incorrect adversarial estimation of the missing structural or relational contribution will force the invariant to fail, triggering the absolute Global Veto:

$$V_G(S') = 0$$

and resulting in the immediate topological rejection of the candidate state S' .

Operational Remark 18a. — Authenticated Verifier Identities and Non-Paralyzing Validation

In concrete CNVS implementations, verifiers are assumed to be protocol-recognized entities bound to authenticated cryptographic keys. The global validation layer does not accept arbitrary anonymous submissions. Invalid, unauthenticated, duplicated, expired, or unauthorized verification objects are rejected at the admissibility layer before reaching global validation.

Accordingly, the Global Veto Property applies only to well-formed candidate transitions submitted through authorized verifier identities and admissible validation channels. If a malicious or colluding verifier submits false, inconsistent, duplicated, or fraudulent verification objects, the authenticated identity associated with the corresponding cryptographic key can be deterministically identified, isolated from subsequent validation rounds, penalized, or excluded according to the implementation rules.

Malicious or false submissions therefore produce deterministic rejection, accountability, penalty, or verifier exclusion; they do not create an unbounded right to force repeated global validation.

This operational requirement does not replace the formal CNVS security theorem. It specifies the admissibility conditions under which the global validation function remains finite, authenticated, and non-paralyzing.

18b. Critical Fragment Principle

The security of the CNVS framework does not scale merely with the percentage of verifiers under adversarial control. It depends on whether the adversarial coalition can eliminate the residual uncertainty of every critical fragment required for global reconstruction.

Let $\mathcal{V}_{\text{adv}} \subset \mathcal{V}$ be the adversarial coalition, and let $F_{\text{comp}} \subset \text{Ter}(t)$ be the subset of terminal fragments assigned to colluding verifiers.

A terminal fragment $f^* \in \text{Ter}(t)$ is called critical if its internal payload, structural role, or invariant binding is necessary for reconstructing a globally admissible candidate state.

A critical fragment f^* is called residual non-inferable with respect to the adversarial view if its internal semantic payload d^* satisfies:

$$H_{\infty}(d^* \mid \text{View}_{\text{adv}}) \geq h^*, \quad (52)$$

for some $h^* > 0$.

Equivalently, the optimal adversarial guessing probability is bounded by:

$$\mathbb{P}_{\text{guess}}(d^* \mid \text{View}_{\text{adv}}) \leq 2^{-h^*}. \quad (53)$$

Therefore, if at least one critical fragment f^* remains outside adversarial control and satisfies residual non-inferability, then any unauthorized reconstruction that requires d^* is bounded by:

$$\mathbb{P}(\text{Rec}^*) \leq 2^{-h^*}.$$

More generally, let $F_{\text{miss}} = \{f_1, \dots, f_r\}$ be the set of missing critical fragments required for reconstruction, and let $D_{\text{miss}} = (d_1, \dots, d_r)$ denote their internal semantic payloads.

We define the residual critical min-entropy as:

$$H_{\text{crit}} \triangleq H_{\infty}(D_{\text{miss}} \mid \text{View}_{\text{adv}}). \quad (54)$$

If:

$$H_{\text{crit}} > 0,$$

then the optimal adversarial guessing probability for the missing critical payload vector satisfies:

$$\mathbb{P}_{\text{guess}}(D_{\text{miss}} \mid \text{View}_{\text{adv}}) \leq 2^{-H_{\text{crit}}}. \quad (55)$$

Accordingly, any unauthorized reconstruction event Rec^* requiring the correct recovery of D_{miss} is bounded by:

$$\mathbb{P}(\text{Rec}^*) \leq 2^{-H_{\text{crit}}}.$$

The relevant security quantity is not the total information size $I_{\text{total}}(f^*)$ of the missing fragment. The relevant quantity is the residual conditional min-entropy that remains after conditioning on everything available to the adversary:

$$\mathbb{P}(\text{Rec}^*) \leq 2^{-H_\infty(D_{\text{miss}} \mid \text{view}_{\text{adv}})}. \quad (56)$$

Consequently, extensive collusion does not imply deterministic reconstruction if at least one reconstruction-critical component remains outside adversarial control and retains strictly positive residual conditional min-entropy.

This principle is conditional: if the adversarial view collapses the residual min-entropy of every critical missing fragment to zero, then the Critical Fragment Principle no longer provides protection. CNVS security therefore requires that critical fragments remain not only uncompromised, but also non-inferable under the hidden-invariant and bounded-leakage conditions of Lemmas 8 and 10.

Remark 18.1 (Gold-Custody Invariant)

Suppose a CNVS state encodes a physical gold-custody verification process. A colluding coalition controls 99 terminal fragments and obtains, under the pessimistic semantic exposure model, 99 local payloads corresponding to 99 kg of gold.

The coalition therefore knows that the exposed local payloads sum to 99 kg.

However, the coalition does not know the internal instantiated invariant C_{int} . In particular, it does not know whether the global validation condition requires:

- total mass equal to 100 kg;
- total mass equal to 120 kg;
- minimum mass above a hidden threshold;
- a mass-purity relation;
- a vault-distribution constraint;
- a temporal custody invariant;
- a relational mapping between different physical deposits.

Therefore, from the adversarial view alone, the missing fragment is not determined.

If the hidden instantiated invariant were publicly known to be: $x_1 + \dots + x_{100} = 100 \text{ kg}$,

then the missing value would be algebraically inferable: $x_{100} = 100 \text{ kg} - 99 \text{ kg} = 1 \text{ kg}$.

But this is not the CNVS adversarial model.

In CNVS, the adversary may know the general class of conservation-type invariants (C_{pub}), but not the internal invariant instance (C_{int}), its hidden parameters (θ_C), or its relational binding (R_{int}).

Thus, possessing 99 kg does not imply knowing that the missing fragment must be 1 kg. The missing fragment remains non-inferable as long as:

$$H_{\infty}(D_{\text{miss}} \mid W_{\text{adv}}, \text{View}_{\text{adv}}, C_{\text{pub}}) \geq h_{\min} > 0.$$

Remark 18.2 (Reactor Example)

To clarify the ontology of a terminal fragment within the CNVS framework, consider the controlled verification of a classified reactor construction process.

In a traditional data-sharding model, fragmentation would correspond to cutting the reactor blueprint into several pieces and distributing those pieces to different workers. In that case, each fragment would be a syntactic or cryptographic share of the global semantic payload. If sufficiently many workers colluded, they could attempt to reconstruct the underlying blueprint.

CNVS does not use this model.

In CNVS, terminal fragments are closed native evaluative units. Each fragment may contain an internal payload d_i within the system boundary, but the external verifier does not receive d_i . The verifier receives only a local measurement grid or task specification, performs an independent observation obs_i , and reports obs_i back to the system.

For example, the verifier may be instructed to measure the length and diameter of a specific screw. The verifier can execute this local task, but does not receive the internal expected value d_i , the global reactor design, the hidden invariant C_{int} , or the relational binding that explains how this local measurement participates in global validation.

Thus, the local task is semantically meaningful at the operational level, but not globally reconstructive by itself.

The topology and task structure may leak limited information about the system. CNVS therefore does not require perfect zero mutual information between metadata and payload. Instead, by Lemma 10, such leakage is bounded:

$$I(X_s ; M_s) \leq \gamma_{\text{top}}.$$

Moreover, by Lemma 8, the adversarial view must not collapse the residual conditional min-entropy of any missing critical payload. For a missing critical payload d_{miss} , CNVS requires:

$$H_{\infty}(d_{\text{miss}} \mid \text{View}_{\text{adv}}) \geq h_{\min} > 0.$$

The reactor example should therefore be understood through three adversarial scenarios.

Scenario 1. The adversary controls many local payloads, but lacks hidden binding and instantiated invariants.

The adversary may obtain many local payloads under the pessimistic semantic exposure model. However, without the hidden invariant instance C_{int} , the internal parameters Θ_C , and the relational binding R_{int} , those local payloads do not necessarily determine the missing critical payload.

Security in this scenario follows if the missing critical component retains positive residual min-entropy:

$$H_\infty(d_{miss} \mid W_k, M_s, C_{pub}) \geq h_{min} > 0.$$

Scenario 2. The adversary observes topology or metadata, but does not control the relevant payloads.

Topology and metadata may reduce uncertainty, but only within the bounded-leakage model. Observing M_s may reveal limited information about the verification structure, but does not by itself imply deterministic reconstruction if:

$$I(X_s ; M_s) \leq \gamma_{top}$$

and residual uncertainty remains positive after conditioning on the adversarial view.

Scenario 3. The adversary controls many payloads and observes substantial topology. This is the strongest case. CNVS does not claim unconditional security here. The system remains protected only if at least one reconstruction-critical component remains both uncompromised and non-inferable under the adversarial view.

Formally, if reconstruction requires a missing critical payload d_{miss} and:

$$H_\infty(d_{miss} \mid \text{View}_{adv}) \geq h_{min} > 0,$$

then:

$$\mathbb{P}_{guess}(d_{miss} \mid \text{View}_{adv}) \leq 2^{-h_{min}}.$$

Accordingly, deterministic reconstruction is blocked only while the residual min-entropy condition holds.

If, instead, the adversary obtains the compromised payloads, the relevant topology, the instantiated hidden invariant C_{int} , and enough information to reduce:

$$H_\infty(d_{miss} \mid \text{View}_{adv}) = 0,$$

then this specific CNVS security condition fails.

Conversely, if the instantiated invariants C_{int} remain cryptographically and epistemically isolated, the system maintains an intrinsic structural resilience against deterministic reconstruction: possessing almost all topological and semantic data is useless if the

adversary is deprived of the validation condition to solve the state equation (Remark 11.2 and Axiom III).

Therefore, the reactor example does not establish unconditional secrecy. It illustrates the CNVS principle that local measurement tasks, hidden invariant binding, bounded metadata leakage, and residual min-entropy jointly prevent deterministic reconstruction under explicit assumptions.

19. Scientific Positioning

The CNVS architecture synthesizes principles from Shannon Information Theory, probabilistic combinatorics, structural graph theory, and distributed formal verification into a unified mathematical framework.

Crucially, CNVS differs structurally from traditional threshold cryptography (e.g., Shamir's Secret Sharing). In standard threshold models, the fragment (the "share") intrinsically contains a mathematical projection of the plaintext payload.

In the CNVS framework, the verifier is never provided with a data share containing the internal semantic payload. Instead, the verifier receives strictly bounded metadata dictating a localized measurement task, subsequently producing observational data used solely for convergence verification. This epistemic decoupling shifts the security analysis away from cryptographic key-splitting and toward bounded topological leakage, hidden binding, and residual non-inferability.

Demarcation from Threshold Cryptography (Shamir Secret Sharing)

Traditional decentralized confidentiality relies heavily on threshold schemes like Shamir Secret Sharing (SSS). In SSS, the native payload is mathematically divided into polynomial shares; thus, the fragment itself *is* a cryptographic projection of the semantic payload. If an adversary gathers a threshold τ of shares, algebraic reconstruction is deterministic.

CNVS differs from this paradigm. Terminal fragments in CNVS do not carry polynomial shares of the native payload ($d \in D_\sigma$); they carry autonomous measurement grids (I_{metadata}). The global state is validated through observational convergence rather than algebraic threshold reassembly. Therefore, CNVS is not primarily exposed to polynomial-interpolation reconstruction in the same way as threshold-share schemes.

Demarcation from Secure Multi-Party Computation (SMPC)

Secure Multi-Party Computation allows multiple parties to jointly evaluate a function over private inputs without revealing the inputs themselves. While it shares with CNVS the broad goal of privacy-preserving evaluation, SMPC operates within a cryptographic computation model and often requires non-trivial cryptographic machinery, such as garbled circuits,

secret sharing protocols, or homomorphic techniques, to compute a function securely. Moreover, standard SMPC protocols typically rely on interactive communication among participating nodes, which may introduce substantial coordination and network overhead. CNVS, by contrast, models verification through restricted local observation, randomized assignment, and hidden global binding. External verifiers do not communicate with one another during the local verification phase, do not jointly compute a global function, and do not participate in a consensus or interactive computation protocol. They act only as epistemically isolated observers of assigned external measurement grids. Security is therefore analyzed through hidden-invariant residual non-inferability, bounded topological leakage, and the Global Veto, rather than through secure joint computation. This places CNVS in a different design space from SMPC, especially when verification involves physical, empirical, or externally observed attributes.

Demarcation from Verifiable Computation (VC) and zk-SNARKs

In Verifiable Computation frameworks, such as zk-SNARKs, a prover mathematically demonstrates the correct execution of a function $f(x)$ without revealing the input x . While superficially similar in its privacy-preserving goals, CNVS operates under a different formal ontology. In CNVS, the external verifier is not a cryptographic prover generating a mathematical proof of a known computation; rather, it acts as a physical oracle executing a real-world measurement. The verifier in CNVS does not possess the hidden input to prove it. Consequently, CNVS security is not primarily modeled through proof-generation hardness, but through restricted verifier knowledge, hidden invariant binding, and residual non-inferability.

Demarcation from Probabilistically Checkable Proofs (PCP)

The PCP theorem allows the verification of global proof validity by means of a randomized verifier that queries only a constant number of locations in an encoded proof string. While PCP techniques use local proof queries to certify global properties, the queried positions directly belong to the proof object itself; they are not external empirical observers and are not semantically blind in the CNVS sense. In contrast, the CNVS framework strictly enforces the Knowledge Restriction bound introduced in Section 6. The terminal fragment provides a measurement grid whose exposed information remains bounded and non-reconstructive with respect to the global semantic payload, in accordance with Lemma 10.

Demarcation from Blind Quantum Computing (BQC)

Blind Quantum Computing is conceptually adjacent to CNVS, as it allows a client to delegate a computation to a server without revealing the delegated circuit or the underlying data.

However, the demarcation is structural, physical, and operational. First, BQC achieves blindness through the delegation of quantum states and rotational encryption requiring quantum hardware. CNVS operates within a classical verification setting and seeks restricted verifier knowledge through randomized assignment, structural non-uniformity, and hidden binding. Second, BQC servers act as passive computational executors with no localized verification; they compute functions blindly without checking the semantic alignment of their outputs. Conversely, CNVS nodes function as restricted local empirical observers: each terminal fragment enforces a local verification function (V_L) that requires a real-world empirical observation (obs) to satisfy the local convergence predicate defined by the measurement-grid constraints. Furthermore, CNVS natively integrates a Global Veto mechanism (V_G) over structural invariants (C), a holistic state-validation feature not inherently present in the purely computational delegation of BQC.

Demarcation from Classical Distributed Verification

Classic distributed verification, including traditional BFT consensus models, relies fundamentally on peer-to-peer communication, network broadcasting, quorum formation, and threshold agreement rules in order to agree on the validity of a state.

The CNVS framework diverges structurally by eliminating the very concept of systemic voting. Verifiers assigned to terminal fragments operate under strict semantic blindness: they do not communicate with each other, do not know the global implications of their assigned task, and do not express any systemic consensus. They act only as isolated measurement probes over external grids, returning local observations to the system. Consequently, in CNVS, validation does not emerge from distributed agreement, but from deterministic mathematical convergence inside the closed verification structure. Even if all terminal fragments are locally admissible, the candidate state is rejected if relational consistency or the satisfaction of hidden global invariants fails.

This is formalized by the Global Veto: $V_G(S) = 1$ only when local admissibility, relational consistency, and satisfaction of global invariants jointly converge unambiguously.

Summary

CNVS should therefore be understood as a candidate formal class of closed native verification systems characterized by:

1. internal payloads not ordinarily disclosed to verifiers;
2. external measurement-grid based observation;
3. randomized assignment of terminal evaluative units;
4. hidden invariant binding and relational topology;

5. bounded topological leakage;
6. residual non-inferability of critical fragments;
7. global admissibility enforced by systemic veto.

The framework does not claim unconditional superiority over existing cryptographic or distributed verification systems. Its contribution is the formalization of a different verification architecture whose security properties are conditional on explicit assumptions: bounded leakage, hidden instantiated invariants, residual min-entropy, randomized assignment, and sub-unitary composite compromise probability.

20. Open Problems

While this paper establishes the foundational formal proofs of the CNVS architecture, several advanced vectors require further formalization:

1. Cryptographic Reductions: Formal mapping of the local knowledge restriction bounds to standard hard computational problems (e.g., Learning With Errors (LWE) or Discrete Logarithm).
2. Information-Theoretic Lower Bounds: Derivation of the absolute optimal adversarial allocation strategies and the corresponding tightest bounds for I_{metadata} expansion.
3. Empirical Validation: Extensive stress-testing of the asymptotic bounds ($\mathbb{P}(\text{Rec}^*) \rightarrow 0$) via large-scale Monte Carlo systemic simulations under dynamically shifting adversarial assignment ($\mathcal{A}_i$).

Appendix A. Minimal Formal Definition

A structure

$$M = (\mathcal{S}, E, R, C, D, A, V_L, V_G, \text{Rec}, \mathcal{K}(v, f))$$

is a Closed Native Verification System if and only if:

1. $\mathcal{S}$ is a typed finite-binary structural universe.
2. Raw structural membership belongs to the candidate state layer: an element s may belong to E_{cand} independently of its current verification status.
3. Terminal admissibility requires local verification: $\text{Adm}_L(s) = 1 \Leftrightarrow \pi_1(V_L(s)) = 1$.
4. Accepted state validity requires global verification: $S \in \mathfrak{S}^+ \Leftrightarrow V_G(S) = 1$.
5. Global verification requires terminal admissibility, relational consistency, and invariant satisfaction.
6. Decomposition is finite, recursive, and reconstructible.
7. Verifier knowledge is strictly restricted.

8. Terminal assignments are randomized.
9. Invalid candidate states are rejected from accepted state evolution; their rejection does not imply that their candidate elements were structurally non-existent.
10. The symbols V_L and V_G denote respectively the Local Verification Function and the Global Verification Function defined in this paper.

AI-Assistance Disclosure

The Author acknowledges the use of AI tools for editorial refinement, language optimization, formatting support, consistency checking, and assistance in the formal correction of notation, domains, and mathematical expressions.

The conceptual architecture, core theoretical framework, axiomatic structure, original mathematical relations, intended formal meaning of the equations, and scientific claims of the CNVS framework were developed and determined by the author. The author assumes full scientific responsibility for the content of this manuscript.

References

- [1] C. E. Shannon, "A Mathematical Theory of Communication," *The Bell System Technical Journal*, vol. 27, no. 3, pp. 379-423, 1948.
- [2] A. Shamir, "How to share a secret," *Communications of the ACM*, vol. 22, no. 11, pp. 612-613, 1979.
- [3] G. R. Blakley, "Safeguarding cryptographic keys," in *Proceedings of the 1979 International Workshop on Managing Requirements Knowledge (AFIPS)*, pp. 313-317, 1979.
- [4] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *Journal of the American Statistical Association*, vol. 58, no. 301, pp. 13-30, 1963.
- [5] H. Chernoff, "A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations," *The Annals of Mathematical Statistics*, vol. 23, no. 4, pp. 493-507, 1952.
- [6] A. Rényi, "On measures of entropy and information," in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, Volume 1: Contributions to the Theory of Statistics, University of California Press, pp. 547-561, 1961.
- [7] S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof systems," *SIAM Journal on Computing*, vol. 18, no. 1, pp. 186-208, 1989.
- [8] A. C. Yao, "Protocols for secure computations," in *Proceedings of the 23rd Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 160-164, 1982.
- [9] U. Maurer, "Information-theoretic cryptography," in *Advances in Cryptology — CRYPTO '99*, Springer, pp. 47-64, 1999.
- [10] S. Dziembowski and K. Pietrzak, "Leakage-resilient cryptography," in *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 293-302, 2008.
- [11] L. Lamport, R. Shostak, and M. Pease, "The Byzantine Generals Problem," *ACM Transactions on Programming Languages and Systems*, vol. 4, no. 3, pp. 382-401, 1982.
- [12] P. Feldman, "A practical scheme for non-interactive verifiable secret sharing," in *Proceedings of the 28th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 427-438, 1987.
- [13] T. P. Pedersen, "Non-interactive and information-theoretic secure verifiable secret sharing," in *Advances in Cryptology — CRYPTO '91*, Springer, pp. 129-140, 1992.
- [14] N. Alon and J. H. Spencer, *The Probabilistic Method*, 4th ed., John Wiley & Sons, 2016.

- [15] P. Erdős and A. Rényi, "On random graphs I," *Publicationes Mathematicae Debrecen*, vol. 6, pp. 290-297, 1959.
- [16] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, "Zerocash: Decentralized anonymous payments from Bitcoin," in *2014 IEEE Symposium on Security and Privacy*, pp. 459-474, 2014.
- [17] E. Ben-Sasson et al., "Scalable, zero knowledge, and quantum secure computational integrity," in *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing (STOC)*, ACM, pp. 595-604, 2015.
- [18] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith, "Fuzzy extractors: How to generate strong keys from biometrics and other noisy data," *SIAM Journal on Computing*, vol. 38, no. 1, pp. 97-139, 2008.
- [19] O. Goldreich, S. Micali, and A. Wigderson, "How to play any mental game or A completeness theorem for protocols with honest majority," in *Proceedings of the 19th Annual ACM Symposium on Theory of Computing (STOC)*, ACM, pp. 218-229, 1987.
- [20] M. Ben-Or, S. Goldwasser, and A. Wigderson, "Completeness theorems for non-cryptographic fault-tolerant distributed computation," in *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC)*, ACM, pp. 1-10, 1988.
- [21] D. Chaum, C. Crépeau, and I. Damgård, "Multiparty unconditionally secure protocols," in *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC)*, ACM, pp. 11-19, 1988.
- [22] Y. Lindell, "Secure Multiparty Computation," *Communications of the ACM*, vol. 64, no. 1, pp. 86-96, 2021.
- [23] R. Gennaro, C. Gentry, B. Parno, and J. Raykova, "Quadratic Span Programs and Succinct NIZKs without PCPs," in *Advances in Cryptology – EUROCRYPT 2013*, Springer, pp. 626-645, 2013.
- [24] J. Groth, "On the size of pairing-based non-interactive arguments," in *Advances in Cryptology – EUROCRYPT 2016*, Springer, pp. 305-326, 2016.
- [25] B. Parno, J. Howell, C. Gentry, and M. Raykova, "Pinocchio: Nearly practical verifiable computation," in *2013 IEEE Symposium on Security and Privacy*, IEEE, pp. 238-252, 2013.
- [26] C. Dwork, N. Lynch, and L. Stockmeyer, "Consensus in the presence of partial synchrony," *Journal of the ACM*, vol. 35, no. 2, pp. 288-323, 1988.
- [27] M. Castro and B. Liskov, "Practical Byzantine fault tolerance," in *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI)*, USENIX, pp. 173-186, 1999.

- [28] R. Canetti, "Universally composable security: A new paradigm for cryptographic protocols," in *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 136-145, 2001.
- [29] Y. Ishai, A. Sahai, and D. Wagner, "Private circuits: Securing hardware against probing attacks," in *Advances in Cryptology – CRYPTO 2003*, Springer, pp. 463-481, 2003.
- [30] S. Micali, "Computationally sound proofs," *SIAM Journal on Computing*, vol. 30, no. 4, pp. 1253-1298, 2000.
- [31] D. R. Stinson, "An explication of secret sharing schemes," *Designs, Codes and Cryptography*, vol. 2, no. 4, pp. 357-390, 1992.
- [32] A. Smith, "Scrambling adversarial errors using few random bits, optimal information reconciliation, and better private codes," in *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 395-404, 2007.
- [33] C. McDiarmid, "On the method of bounded differences," in *Surveys in Combinatorics*, Cambridge University Press, pp. 148-188, 1989.
- [34] J. Katz and Y. Lindell, *Introduction to Modern Cryptography*, 3rd ed., CRC Press, 2020.
- [35] M. Mitzenmacher and E. Upfal, *Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis*, Cambridge University Press, 2005.
- [36] F. B. Schneider, "Implementing fault-tolerant services using the state machine approach: A tutorial," *ACM Computing Surveys*, vol. 22, no. 4, pp. 299-319, 1990.
- [37] E. M. Clarke and E. A. Emerson, "Design and synthesis of synchronization skeletons using branching time temporal logic," in *Logics of Programs*, Springer, pp. 52-71, 1981.
- [38] R. A. Kemmerer, "Integrating formal methods into the development process," *IEEE Software*, vol. 7, no. 5, pp. 37-50, 1990.
- [39] O. Goldreich, *Foundations of Cryptography: Volume 1, Basic Tools*, Cambridge University Press, 2001.
- [40] O. Goldreich, *Foundations of Cryptography: Volume 2, Basic Applications*, Cambridge University Press, 2004.
- [41] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed., MIT Press, 2009.
- [42] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell, "Bulletproofs: Short proofs for confidential transactions and more," in *2018 IEEE Symposium on Security and Privacy*, IEEE, pp. 315-334, 2018.
- [43] J. A. Halderman et al., "Lest we remember: cold-boot attacks on encryption keys," in *Proceedings of the 17th USENIX Security Symposium*, pp. 45-60, 2008.

- [44] R. Cramer, I. Damgård, and J. B. Nielsen, *Secure Multiparty Computation and Secret Sharing*, Cambridge University Press, 2015.
- [45] S. Halevi, "A plausible approach to computer-aided cryptographic proofs," *IACR Cryptology ePrint Archive*, vol. 2005, p. 181, 2005.
- [46] S. M. Bellare and M. Merritt, "Encrypted key exchange: Password-based protocols secure against dictionary attacks," in *Proceedings of the 1992 IEEE Symposium on Security and Privacy*, pp. 72-84, 1992.
- [47] A. Shamir, "IP = PSPACE," *Journal of the ACM*, vol. 39, no. 4, pp. 869-877, 1992.
- [48] M. Blum, P. Feldman, and S. Micali, "Non-interactive zero-knowledge and its applications," in *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC)*, ACM, pp. 103-112, 1988.
- [49] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.