PoISV: A Peer-to-Peer System for Independent Semantic Validation

Fatih Dinc
fatdinhero@gmail.com
Pforzheim, Germany

Version 1.0 – April 2026
DOI: 10.5281/zenodo.19642292
OTS: poisv.com/verification

Abstract

A purely peer-to-peer system for evaluating the quality of digital statements would allow information to be validated directly between independent observers without relying on a central authority. Existing consensus mechanisms such as Proof-of-Work or Proof-of-Stake order transactions but do not assess whether the content of the ordered information is semantically consistent or independently validated. We propose a solution based on a network of validators that computes a composite score for each statement—measuring its semantic consistency with a public knowledge corpus—and enforces a statistical non-collusion test derived from the correlation of validator errors on dummy claims. Accepted statements are permanently anchored in a directed acyclic graph, forming a globally available, tamper-evident knowledge base. The system is secure as long as honest validators collectively perform more independent semantic work than any coordinated group of attackers. The mathematical foundation of the non-collusion test rests on the Meta-Bell Theory [6], a measure-theoretic extension of Bell's inequalities, which provides a formal criterion for distinguishing genuinely independent validation from coordinated manipulation.

1 Introduction

Trust on the Internet today relies almost entirely on central intermediaries. Financial platforms agree on account balances through central ledgers, social platforms agree on permissible content through central moderators, and artificial intelligence systems produce outputs whose correctness users cannot independently verify. This model works for much of everyday traffic but suffers from the well-known weaknesses of the trust paradigm: costs rise with the importance of the decision, and a baseline rate of errors and manipulation is accepted as inevitable.

Bitcoin introduced a way to reach consensus on the order of transactions without a central party by requiring computational work to be expended [1]. Proof-of-Stake achieves a similar ordering by locking capital as collateral. Neither mechanism, however, makes any verifiable claim about whether the content of the ordered information is semantically consistent or independently assessed. For pure payments this suffices; for the information age, it does not.

What is needed is a consensus mechanism in which the scarce resource is not raw computation or capital, but *demonstrably independent semantic work*. We call such a system **Proof of Independent Semantic Validation (PoISV)**. This work is the third in a sequence: the Meta-Bell Theory [6] establishes the mathematical foundation, Proof of WiseWork [7] is the first application, and PoISV is the completed, operationally precise protocol.

2 The Problem of Semantic Consensus

Consider a network where participants submit *claims*—statements about the world. We want the network to attach a quality score to each claim that reflects its agreement with an objective, publicly verifiable knowledge corpus and the independence of the validators who assessed it.

A single claim evaluated by a single validator is unreliable. A validator may be biased, may have an outdated knowledge base, or may be part of a coordinated group trying to push a false narrative. Even many validators may all be using the same flawed model or manipulated data source. We need a way to ensure that the validators are *truly independent* in their assessment—and we need to be able to prove it mathematically.

3 The PoISV Protocol

3.1 Information Units

We define a claim $A = (d, \sigma)$, where d is the serialized statement data and σ is an optional digital signature of the submitter. The knowledge corpus $\mathcal{K}$ is a versioned dataset addressable via content hashing. Its hash $H(\mathcal{K})$ is fixed in the genesis block.

3.2 The Semantic Consistency Score

For a claim A, the function $S_{con}(A) \in [0, 1]$ is computed deterministically by extracting entities and relations from A, querying $\mathcal{K}$ for supporting or contradicting facts, and returning a weighted agreement value. The exact implementation is fixed by the protocol version; its hash is included in the block header, ensuring full reproducibility.

3.3 The Non-Collusion Statistic Ψ

The Meta-Bell Theory [6] provides the formal basis for distinguishing genuinely independent measurements from coordinated ones. In the PoISV context, each validator i must solve k canonical dummy claims $D_1, \ldots, D_k$ whose correct scores $S^*(D_j)$ are hardcoded in the protocol. This produces an error vector

$$\mathbf{e}_i = \big(|S_i(D_j) - S^*(D_j)|\big)_{j=1}^{k}.$$

For N validators contributing to a block, the non-collusion statistic is

$$\Psi = 1 - \frac{2}{N(N-1)} \sum_{1 \le i < j \le N} |\rho(\mathbf{e}_i, \mathbf{e}_j)|,$$

where ρ denotes the Pearson correlation coefficient. Independent validators produce uncorrelated error patterns, yielding $\Psi \approx 1$. Colluding validators—sharing a common model or data source—produce identical error patterns, yielding $\Psi \approx 0$.

This statistic is the operational realization of the Meta-Bell entanglement measure $\Psi(X, Y)$ defined in [6]: a positive value proves that the observed correlations cannot be explained by any local hidden-variable model, i.e., by any form of shared coordination.

3.4 Acceptance Rule

A block is accepted if and only if

$$\frac{1}{|A|} \sum_{A \in \text{Block}} S_{con}(A) \;\ge\; \theta_{min} \quad \text{and} \quad \Psi \;\ge\; \Psi_{min}.$$

3.5 Block Weight and DAG Ordering

PoISV uses the GHOSTDAG protocol [5] to allow parallel block production without discarding honest work. The weight of a block is

$$\text{Weight}(B) = \Psi_B \cdot \sum_{A \in B} S_{con}(A).$$

The canonical chain is the path through the DAG with the highest cumulative weight. This incentivizes validators to produce high-quality semantic work and to maintain genuine independence.

4 Network Operation

The steps of the network are as follows.

1. New claims are broadcast to all validators.
2. Each validator computes S_{con} for the claim and records its error vector on the dummy claim set.
3. A validator wishing to propose a block collects claims, computes aggregate scores, and creates a block.
4. The block is broadcast along with a zero-knowledge proof of correct computation, produced by a zkVM [4].
5. Other validators accept the block if the scores meet the thresholds and the proof verifies.
6. Accepted blocks are added to the DAG, permanently anchoring each validated claim on-chain.

Validators always consider the DAG path with the highest cumulative weight to be the canonical one. New claims do not need to reach all validators immediately; as long as they reach many, they will enter a block before long.

5 Security Analysis

We consider an attacker controlling a fraction $q < 0.5$ of the validators.

5.1 Attack on Semantic Consistency

To push $S_{con}(A)$ above θ_{min} for a false claim, the attacker must either forge the hash of $\mathcal{K}$—cryptographically infeasible—or convince a majority of validators to accept false paths in $\mathcal{K}$. Since $\mathcal{K}$ is public, any honest validator can produce a fraud proof invalidating the block.

5.2 Attack on Ψ

By Hoeffding's inequality for U-statistics, the probability that an attacker with fraction q of colluding validators can forge $\Psi \geq \Psi_{min}$ is bounded by

$$P(\text{Success}) \;\leq\; \exp\!\left(-2k \cdot \frac{(1-q)^2}{q^2} \cdot \left(\Psi_{min} - (1-q^2)\right)^2\right).$$

Running some results, we can see the probability drop off exponentially with k:

q	k	Ψ_{min}	P(Success)
0.10	32	0.7	$< 10^{-12}$
0.20	32	0.7	$< 10^{-8}$
0.30	64	0.7	$< 10^{-7}$
0.40	64	0.7	$< 10^{-4}$
0.49	128	0.7	$< 10^{-3}$

For practical parameters ($k = 64$, $q = 0.4$, $\Psi_{min} = 0.7$), the probability is below 10^{-4}. Increasing k to 128 dummy claims reduces this to below 10^{-8} across all realistic attacker fractions.

5.3 Incentive Compatibility

Validators are rewarded for proposing valid blocks with newly minted native currency and transaction fees. Manipulation requires not only controlling a significant fraction of validators but making them appear independent—which forces the attacker to run genuinely diverse, independent validation pipelines, defeating the purpose of coordination. A rational attacker finds it more profitable to follow the rules.

6 Privacy and Light Clients

The raw evidence sources and model weights used by validators remain private within the zkVM witness; only final scores and the Ψ statistic are published. Light clients can verify block acceptance by downloading only block headers and zero-knowledge proofs, without running the semantic engine. Light clients may also apply subjective filters—for instance, only displaying claims validated by a specific ethics committee—without imposing these filters on the network as a whole. The protocol itself remains value-neutral.

7 Reclaiming Storage Space

Older validator outputs can be pruned without breaking block hashes. Validator outputs are hashed in a Merkle tree; only the root is stored in the block header. A block header of approximately 200 bytes, a zero-knowledge proof of roughly 200 kilobytes, and a small index per day suffice. At one block per second, this amounts to approximately 6 gigabytes per year—well within the capabilities of modern storage.

8 Conclusion

We have proposed a system for decentralized semantic validation that does not rely on a trusted third party. Starting from digital signatures and a public knowledge corpus, which provide strong provenance guarantees but are incomplete without a mechanism to prevent coordinated manipulation, we built a peer-to-peer network that computes a semantic consistency score and enforces a non-collusion test grounded in the Meta-Bell Theory. Accepted claims are permanently anchored on a DAG, forming a public, tamper-evident knowledge base. The network is robust in its unstructured simplicity. Validators work largely independently, express their acceptance of valid blocks by extending the highest-weight path, and reject invalid blocks by withholding their work. Any needed rules and incentives can be enforced through this consensus mechanism.

References

[1] S. Nakamoto, *Bitcoin: A Peer-to-Peer Electronic Cash System*, 2008.

[2] J. S. Bell, *On the Einstein Podolsky Rosen Paradox*, Physics 1(3):195–200, 1964.

[3] J. F. Clauser, M. A. Horne, A. Shimony, R. A. Holt, *Proposed Experiment to Test Local Hidden-Variable Theories*, Physical Review Letters 23:880–884, 1969.

[4] Nexus Labs, *Nexus zkVM: Enabling Verifiable Computation*, 2024.

[5] Y. Sompolinsky, S. Wyborski, A. Zohar, *PHANTOM GHOSTDAG: A Scalable Generalization of Nakamoto Consensus*, AFT 2021.

[6] F. Dinc, *Meta-Bell-Theorie: Eine maßtheoretische Erweiterung der Bell-Ungleichungen*, Version 1.0, SHA-256: 062e290009f6b7339e9a8b522ce1d94d9021d109d8e4bc41210d1f3dda053a3b, 2026.

[7] F. Dinc, *Proof of WiseWork: Ein Peer-to-Peer-System zur Konsensbildung über Wahrheit*, Version 2.0, SHA-256: e57cae993701a1933a3317e28c7bb7141a01b51b31674adee97b6cf89472c2eb, 2026.

[8] C. E. Shannon, *A Mathematical Theory of Communication*, Bell System Technical Journal 27:379–423, 1948.