# TAC-QSW CONJECTURE

*Quantum Stochastic Walk as a Structural Inference Engine*

*for Decision Eligibility Assessment*

---


Technical White Paper — Version 1.0

TAC-3D Decision Infrastructure

tac3d.sycds.com

April 2026

# Abstract

We introduce the TAC-QSW Conjecture: that the structural eligibility of a multi-evidence decision can be formally assessed by simulating a Quantum Stochastic Walk (QSW) over the pairwise compatibility graph of its evidence layers, and extracting four diagnostic dimensions — Tension, Alignment, Convergence, and Clarity — from the resulting probability flow dynamics.

Classical decision frameworks treat evidence aggregation as a weighted summation problem. We argue this is structurally insufficient. Evidence layers do not merely add; they interact, conflict, reinforce, and cascade. The TAC-QSW framework models these interactions as a quantum-classical hybrid walk on a compatibility graph, revealing structural properties that scalar aggregation cannot capture.

This paper presents the mathematical formulation of the conjecture, the computational architecture of its implementation, the four-dimensional TAC diagnostic framework, and empirical observations from live deployment on the TAC-3D platform.


# 1. Introduction

## 1.1 The Problem with Classical Evidence Aggregation

Consider a decision with three evidence layers: financial projections indicating strong growth, clinical trial data showing mixed efficacy signals, and regulatory intelligence suggesting an uncertain approval timeline. A classical weighted average might yield a moderate composite score — obscuring the fundamental structural tension between layers.

The central insight of the TAC-QSW framework is that decision eligibility is not a property of individual evidence layers in isolation, but of the structural dynamics that emerge from their interactions. A decision with internally coherent but collectively contradictory evidence layers is not simply "medium quality" — it is structurally compromised in a qualitatively distinct way.

## 1.2 Quantum Walks as Structural Probes

Quantum walks have been established as superior to classical random walks for graph exploration, achieving quadratic speedups in hitting time and revealing structural properties invisible to diffusive processes. In the context of decision analysis, we repurpose this exploratory power: the QSW does not search for an answer, but probes the compatibility landscape of the evidence graph to reveal its structural geometry.

The steady-state distribution of the walk — where probability amplitude has settled after sufficient decoherence — encodes information about which evidence nodes are structurally dominant, which are marginalized, and how coherently the system has resolved its internal tensions.

# 2. Mathematical Formulation

## 2.1 The Compatibility Graph

Let D be a decision and E = $\{e_1, e_2, ..., e_n\}$ be its evidence layers. We define the compatibility matrix C ∈ [-1, 1]^(n×n) where:

$$C_{ij} \in [-1, 1], \quad C_{ii} = 1, \quad C_{ij} = C_{ji}$$

C_ij = 1.0 indicates that evidence layer i strongly supports and reinforces layer j. C_ij = 0.0 indicates structural independence. C_ij = -1.0 indicates fundamental contradiction. The matrix is symmetric and assessed by a language model evaluating semantic and logical compatibility between evidence pairs.

## 2.2 The Hamiltonian

The system Hamiltonian is constructed from the compatibility matrix with diagonal terms encoding node strength:

$$H_{ij} = C_{ij} \quad (i \neq j)$$
$$H_{ii} = s_i = \mathrm{mean}\{ \max(0, C_{ij}) : j \neq i \}$$

The diagonal term s_i represents the mean positive connectivity of node i — its structural support from the rest of the evidence graph. Nodes with high s_i are well-integrated; nodes with low s_i are structurally isolated or adversarial.

## 2.3 The Lindblad Evolution Generator

We model decoherence via a simplified Lindblad operator. The effective evolution generator L_eff captures both coherent transport and dissipative relaxation:

$$L_{eff}[i,j] = C_{ij} \cdot \max(0, s_i - s_j) \cdot 2 \quad (\text{for } C_{ij} > 0,\ i \neq j)$$
$$L_{eff}[i,j] = |C_{ij}| \cdot 0.3 \quad (\text{for } C_{ij} < 0,\ i \neq j)$$
$$L_{eff}[i,i] = -\Sigma_{j\neq i} L_{eff}[j,i] - \gamma_i$$

The decoherence rate $\gamma_i = 0.02 \cdot (1 - s_i + 0.1)$ couples node-level decoherence to structural support: weakly supported nodes decohere faster, losing quantum coherence and relaxing toward classical diffusion more rapidly.

## 2.4 The Combined Generator and Time Evolution

The combined generator G integrates coherent Hamiltonian evolution with dissipative Lindblad dynamics:

$$G = -0.1H + L_{eff}$$

Starting from a uniform initial distribution ρ(0) = (1/n, ..., 1/n), the system evolves via Euler integration:

$$\rho(t + dt) = \rho(t) + dt \cdot G\, \rho(t)$$
$$\rho(t) \rightarrow |\rho(t)| \,/\, \| |\rho(t)| \|_1$$

The renormalization at each step maintains ρ as a probability distribution while allowing the non-conservative Lindblad terms to redistribute amplitude. The simulation runs for t_max = 20 time units with 200 steps (dt = 0.1).

# 3. The TAC Diagnostic Framework

From the QSW trajectory {ρ(0), ρ(1), ..., ρ(T)}, we extract four structural diagnostics that together constitute the TAC assessment.

## 3.1 Tension (T)

Tension measures the degree of adversarial structure in the evidence graph — how much contradictory pressure exists between evidence layers:

$$\text{neg\_strength} = \text{mean}\{\ |C_{ij}| : C_{ij} < -0.1,\ i < j\ \}$$
$$T = \text{clip}(100 \cdot (1 - \text{neg\_strength}), 0, 100)$$

High Tension (T → 100) indicates that negative compatibility pairs are weak — the evidence graph has little internal contradiction. Low Tension (T → 0) indicates strong systematic contradiction between evidence layers.

## 3.2 Alignment (A)

Alignment measures whether the QSW dynamics respect the compatibility structure — whether probability amplitude flows in directions consistent with the evidence relationships:

$$changes_i = \rho_i(T) - \rho_i(0)$$

$$A = \text{clip}\left( 100 \cdot |\{(i,j): \text{sgn}(C_{ij}) \cdot changes_i \cdot changes_j > 0\}| \,/\, |pairs|,\ 0,\ 100 \right)$$

An aligned system (A → 100) is one where compatible evidence layers gain or lose amplitude together, and contradictory pairs diverge. This measures the self-consistency of the QSW resolution.

## 3.3 Convergence (C)

Convergence measures the overall structural cohesion of the compatibility graph, combining mean compatibility with the proportion of negative pairs:

$$mean\_compat = \text{mean}\{ C_{ij} : i < j \}$$

$$neg\_ratio = |\{(i,j): C_{ij} < -0.1,\ i < j\}| \,/\, |pairs|$$

$$C = \text{clip}\left( 100 \cdot (mean\_compat+1)/2 \cdot 0.6 + (1 - neg\_ratio) \cdot 0.4,\ 0,\ 100 \right)$$

Convergence is a graph-level measure of how structurally coherent the evidence assembly is as a whole. It penalizes both low mean compatibility and high proportions of contradictory pairs.

## 3.4 Clarity (K)

Clarity measures the differentiation of the steady-state distribution — whether the QSW has resolved into a clear hierarchy of evidence dominance or remained diffuse:

$$entropy = -\Sigma_i\ \rho_i(T) \cdot \ln(\rho_i(T))$$

$$K = \text{clip}\left( 100 \cdot (1 - entropy / \ln(n)),\ 0,\ 100 \right)$$

Maximum Clarity (K = 100) corresponds to a maximally concentrated steady state — one evidence node has absorbed all amplitude. Minimum Clarity (K = 0) corresponds to the uniform distribution — the QSW has not differentiated between nodes. Clarity captures whether the system has a structural resolution.

# 4. The Eligibility Index and Verdict Framework

## 4.1 Composite Eligibility Index

The four TAC dimensions are combined into a composite Eligibility Index E via a weighted linear combination calibrated to reflect the structural priority ordering of decision analysis:

$$E = 0.30 \cdot T + 0.35 \cdot A + 0.20 \cdot C + 0.15 \cdot K$$

The weighting reflects the following priority logic: Alignment carries the highest weight (0.35) because self-consistency of the QSW resolution is the most direct indicator of structural coherence. Tension (0.30) is weighted heavily because adversarial evidence is the most common source of decision failure. Convergence (0.20) captures global graph structure. Clarity (0.15) captures resolution definiteness.

## 4.2 Verdict Thresholds

| Verdict | E Range | Structural Interpretation |
|---|---|---|
| ELIGIBLE | $E \geq 72$ | Evidence graph is structurally coherent. Decision is supported for execution. |
| ELIGIBLE WITH CONDITIONS | $48 \leq E < 72$ | Moderate structural tensions. Addressable with targeted risk mitigation. |
| DEFER | $28 \leq E < 48$ | Significant structural incompatibilities. Decision requires substantial restructuring. |
| RESTRUCTURE | $E < 28$ | Fundamental evidence contradictions. Decision premise requires reconception. |

# 5. The TAC-QSW Conjecture

## 5.1 Formal Statement

The TAC-QSW Conjecture states:

> *For any decision D with evidence layers $E = \{e_1, \ldots, e_n\}$, the structural eligibility of D — defined as the probability that D will achieve its stated objectives given the evidence assembly — is monotonically correlated with the Eligibility Index E(D) derived from the QSW dynamics on the compatibility graph C(E).*

More precisely: decisions with $E \geq 72$ exhibit structurally coherent evidence assemblies where the QSW reaches a well-differentiated steady state consistent with the compatibility structure;

decisions with $E < 28$ exhibit evidence assemblies where the QSW cannot resolve due to fundamental contradictions encoded in the compatibility graph.

## 5.2 Why QSW and Not Classical Random Walk

The conjecture specifically requires a quantum-classical hybrid walk rather than a classical random walk for the following reasons:

- Interference effects: The Hamiltonian term in G allows destructive interference between contradictory evidence pathways, amplifying structural contradictions that classical diffusion would average away.
- Decoherence-dependent relaxation: The node-level decoherence rate $\gamma_i$ couples structural support to relaxation speed. Weakly supported nodes lose coherence faster, creating an asymmetry that classical walks cannot model.
- Steady-state geometry: The QSW steady state reflects the full spectral structure of the compatibility graph, not merely the stationary distribution of a Markov chain. Clarity — which is only meaningful in the QSW context — would be identically zero for any ergodic classical walk on a connected graph.

## 5.3 The v1/v2 Comparison Program

TAC-3D operates two parallel assessment engines: v1 (the primary Claude-based natural language analysis) and v2 (the QSW engine described in this paper). Every decision submitted to the platform triggers both engines, with results stored in parallel tables (screenings and qsw_screenings respectively).

This parallel deployment constitutes an ongoing empirical test of the conjecture. The TAC-QSW Conjecture predicts that verdicts from v2 (pure structural analysis) should correlate with v1 verdicts (semantic analysis) in the majority of cases, with systematic divergences occurring precisely when v1's language model introduces contextual reasoning that goes beyond the structural compatibility signal.

Cases of divergence are theoretically informative: they identify decisions where semantic framing and structural geometry are in conflict — a meta-diagnostic signal of its own.

# 6. Implementation Architecture

## 6.1 Compatibility Matrix Generation

The compatibility matrix C is generated via a structured prompt to a large language model (Claude Sonnet), which evaluates each pair of evidence layers for semantic and logical compatibility on the [-1, 1] scale. The prompt enforces symmetry and includes the decision context to ensure compatibility is assessed relative to the specific decision being evaluated, not in the abstract.

This is the only step in the QSW pipeline that involves a language model. All subsequent computation is purely numerical.

### 6.2 Numerical QSW Simulation

The QSW is implemented in NumPy using Euler integration with t_max = 20, steps = 200, dt = 0.1. The implementation is deterministic given C. The computational complexity is $O(n^2 \cdot \text{steps})$ where n is the number of evidence layers.

NaN values in the output (arising from division by zero in degenerate single-layer cases) are handled downstream by null substitution, as the QSW framework requires $n \geq 2$ evidence layers for meaningful dynamics.

### 6.3 Integration with TAC-3D Platform

The QSW engine is deployed as a subprocess called by the Node.js proxy server upon completion of each v1 analysis. The input is a JSON object containing the decision text and evidence layer contents; the output is the TAC four-vector and verdict. Results are stored in the qsw_screenings table, linked to the corresponding v1 screening via screening_id.

## 7. Empirical Observations

### 7.1 Current Dataset Status

As of the date of this white paper, the v1/v2 parallel deployment has been operational for less than 24 hours following the resolution of a compatibility module dependency issue. The qsw_screenings dataset is in its initial accumulation phase.

Early observations from manual testing indicate that the QSW engine produces structurally intuitive results: decisions with genuinely contradictory evidence layers (e.g., strong financial projections paired with adverse clinical signals) receive low Tension scores and RESTRUCTURE verdicts, consistent with the conjecture's predictions.

### 7.2 Planned Validation

The planned validation program for the TAC-QSW Conjecture proceeds in three phases:

- Phase 1 — Accumulation: Collect 50-100 parallel v1/v2 assessments across diverse decision types (investment, clinical, strategic).

- Phase 2 — Correlation Analysis: Compute Spearman rank correlation between v1 and v2 Eligibility Index scores. The conjecture predicts $\rho > 0.7$.
- Phase 3 — Divergence Analysis: Identify and manually review cases where v1 and v2 verdicts disagree. Classify divergences as semantic (v1 introduces context beyond C) vs. structural (v2 detects incompatibilities v1 missed).

# 8. Discussion and Limitations

## 8.1 The LLM Compatibility Assessment Dependency

The conjecture's validity depends on the quality of the compatibility matrix C generated by the language model. If C is systematically biased — for example, if the LLM conflates thematic similarity with structural compatibility — the QSW dynamics will reflect that bias. The conjecture is therefore conditional: it claims that structural eligibility is captured by QSW dynamics on a correct compatibility graph.

This is not a weakness unique to this framework; any structural inference method depends on the quality of its input representation. The LLM-based compatibility assessment is a practical approximation of the theoretically ideal semantic compatibility function.

## 8.2 The Euler Approximation

The current implementation uses first-order Euler integration. For compatibility graphs with high spectral radius (large eigenvalues in G), the Euler method may introduce numerical instability. In practice, the renormalization step at each iteration provides a stabilizing correction. Future versions may implement Runge-Kutta or matrix exponential methods for improved numerical fidelity.

## 8.3 The Weight Calibration Question

The weights (0.30, 0.35, 0.20, 0.15) in the Eligibility Index are currently set by theoretical reasoning rather than empirical calibration. As the v1/v2 comparison dataset accumulates, these weights should be recalibrated via regression against v1 verdicts (treated as a noisy but informative ground truth) or against real-world decision outcomes where available.

# 9. Conclusion

The TAC-QSW Conjecture proposes a novel approach to decision eligibility assessment grounded in the structural dynamics of evidence compatibility graphs. By modeling evidence interactions as a quantum-classical hybrid walk, the framework captures structural properties — interference between contradictory pathways, decoherence-dependent relaxation, steady-state differentiation — that are invisible to classical scalar aggregation methods.

The framework is currently in live deployment as the v2 engine of the TAC-3D platform, running in parallel with v1 semantic analysis. The accumulating v1/v2 comparison dataset constitutes an ongoing empirical test of the conjecture.

We conjecture that structural eligibility is a real and measurable property of decision-evidence assemblies, that QSW dynamics provide a natural and principled framework for its measurement, and that the TAC four-vector (Tension, Alignment, Convergence, Clarity) constitutes a complete and interpretable structural diagnostic.

If the conjecture holds, it implies that a purely structural analysis — requiring no understanding of the domain, no access to ground truth outcomes, and no semantic interpretation — can reliably distinguish between decisions that are structurally ready for execution and those that are not.

---

# Appendix A: QSW Engine Implementation

The full implementation of the QSW engine is available at the TAC-3D platform. Core computational parameters:

| Parameter | Value | Description |
|---|---|---|
| `t_max` | `20` | Maximum simulation time |
| `steps` | `200` | Euler integration steps |
| `dt` | `0.1` | Time step size |
| `γ_base` | `0.02` | Base decoherence rate |
| `H_scale` | `0.1` | Hamiltonian coupling strength |
| `Negative rate` | `0.3` | Lindblad rate for $C_{ij} < 0$ |

# References

[1] Kempe, J. (2003). Quantum random walks: An introductory overview. Contemporary Physics, 44(4), 307-327.

[2] Lindblad, G. (1976). On the generators of quantum dynamical semigroups. Communications in Mathematical Physics, 48(2), 119-130.

[3] Childs, A. M. (2010). On the relationship between continuous- and discrete-time quantum walk. Communications in Mathematical Physics, 294(2), 581-603.

[4] Whitfield, J. D., Rodríguez-Rosario, C. A., & Aspuru-Guzik, A. (2010). Quantum stochastic walks: A generalization of random walks and quantum walks. Physical Review A, 81(2), 022323.

[5] TAC-3D Platform. (2026). Decision Eligibility Screener. tac3d.sycds.com

---