

# Ambient Structure Discovery via Stigmergic Mesh: A Self-Organizing Architecture for Surfacing Second-Order Organizational Ignorance

Jeremy McEntire  
jeremy@cageandmirror.com

February 2026

## Abstract

Organizations do not fail because they lack information. They fail because they cannot perceive what they do not know they are missing. Armour’s taxonomy of ignorance identifies five orders, of which the second—the absence of awareness that relevant knowledge exists—produces the costliest failures. IT project cost overruns follow a power-law distribution with exponent  $\alpha \approx 1.0$ , implying infinite expected loss. The cause is consistently traced to hidden interdependencies that no one thought to ask about. The mechanism driving this blindness—dysmemic pressure, the compound selection force that emerges when organizations compress information to coordinate at scale—is substrate-independent: the same compression-selection-equilibrium pattern appears in neural prediction, corporate hierarchy, machine learning, and academic publishing.

Existing approaches to organizational knowledge management share a structural deficiency: they require a query. Retrieval-augmented generation, search engines, dashboards, and multi-agent frameworks all assume someone knows what to look for. This assumption is precisely what second-order ignorance violates.

This paper introduces Ambient Structure Discovery, a paradigm in which a self-organizing computational mesh continuously ingests work artifacts, develops emergent topological specialization through competitive routing, and surfaces structural patterns that no query could have retrieved because no one knew to ask. The architecture draws formal grounding from Adaptive Resonance Theory (stability-plasticity guarantees via vigilance-based category creation), Crawford–Sobel strategic communication theory (quantified information degradation as a function of incentive divergence), Beer’s Viable System Model (cybernetic completeness of the five-system regulatory structure), and rate-distortion theory (empirically demonstrated hierarchical information loss with structural breaks in soft-information sensitivity at specific organizational seams).

Beyond the Discovery quadrant—where hidden interdependencies are surfaced—the paper formalizes a four-quadrant model that includes Normalized Deviance detection, where known risks are discussed but not acted upon, and the decay path by which surfaced findings migrate from discovery to normalized deviance through dysmemic absorption. A formal scoring function prioritizes findings by topological bridge distance, entity confidence, risk coupling, and communicability. The mesh is shown to satisfy the three properties of a Mirror—insulation, access, and authority—required for observation outside local selection pressure. A reference implementation with comprehensive test coverage demonstrates the architecture’s feasibility. The system operates as a prosthetic for organizational perception: it sees what the organization cannot see about itself, because operationally closed systems cannot observe the conditions of their own observation.

**Keywords:** stigmergy, self-organizing systems, second-order ignorance, organizational cybernetics, adaptive resonance theory, spectral anomaly detection, normalized deviance, dysmemic pressure

# 1 The Problem: What You Cannot Ask About

In 2000, Phillip Armour proposed a taxonomy of ignorance that remains underappreciated in software engineering and organizational theory [Armour, 2000]. The taxonomy distinguishes five orders. Zero-order ignorance is the possession of knowledge: you know something. First-order ignorance is the awareness of a gap: you know you do not know something, and you can formulate a question. Second-order ignorance is the absence of that awareness: you do not know that you do not know, and consequently no question exists to ask. Third-order ignorance is the lack of a process for discovering second-order ignorance. Fourth-order ignorance is meta-ignorance about the taxonomy itself.

The distinction between first and second order is the one that matters for engineering organizations. First-order ignorance is manageable. A team that knows it lacks expertise in distributed consensus can hire for it, read about it, or architect around it. The gap is visible and the remediation path is clear. Second-order ignorance is categorically different. A team that does not realize its booking system shares implicit state with its reservation system through a database view that neither team documented cannot take corrective action, because no one has identified a problem requiring correction. The interdependency exists. It will manifest. No one is looking for it.

Flyvbjerg et al.’s empirical analysis of IT project cost overruns demonstrates why second-order ignorance deserves its own theoretical treatment [Flyvbjerg et al., 2022]. Overruns do not follow a normal distribution. They follow a power law with exponent  $\alpha \approx 1.0$ , which implies that both the variance and the mean are infinite. Standard risk management tools (Monte Carlo simulations, confidence intervals, contingency buffers) assume finite moments. They are mathematically inapplicable to the actual distribution. The overruns that bankrupt projects and occasionally companies are not tail events in a well-behaved distribution. They are the expected behavior of a heavy-tailed process driven by hidden interdependencies that compound nonlinearly once they begin to interact. The Standish Group’s finding that fewer than one in three software projects meet their success criteria is the aggregate consequence of this dynamic [Standish Group, 2015].

The organizational theory literature offers a structural explanation for why second-order ignorance persists. Luhmann’s theory of autopoietic systems establishes that organizations are operationally closed: they reproduce their own elements (decisions, communications, roles) from their own elements [Luhmann, 1995]. This closure creates a boundary that is simultaneously the organization’s identity and its blind spot. The system cannot observe the distinction between itself and its environment from within, because that distinction is the condition of its own observation. An organization contains structural truths about itself—hidden couplings, emergent dependencies, unwritten workflows—that its own reporting mechanisms cannot surface, because those mechanisms are constructed from the same operational elements whose interactions produce the hidden truths.

The pattern is general. The same compression-selection-equilibrium sequence operates wherever intelligent systems coordinate through compressed representations. Individual cognition compresses approximately one billion sensory bits per second into approximately ten bits of conscious throughput [Zheng & Meister, 2025]; beliefs that fit the compressed model feel true regardless of their correspondence with reality. Organizations compress through hierarchy, metrics, and process; signals that fit the reporting frame survive, and the frame confirms itself. AI systems compress training corpora into parameter weights; outputs that score well on preference proxies propagate regardless of alignment with the values those proxies were meant to represent. Academic publishing compresses careers into citation counts and journal prestige; work that fits the paradigm survives, and the paradigm confirms itself [Smaldino & McElreath, 2016]. The substrates share little except compression and selection. The pattern appears across all of them because compression and selection are the cause.

The consequences are quantifiable. Agent-based simulation of tournament-based performance evaluation—the most widely adopted formal evaluation mechanism in technology organizations—demonstrates that forced ranking produces classification errors of 32% under idealized conditions and 53% under realistic team composition variance, with the mechanism systematically punishing the best-managed teams [McEntire, 2025c]. The evaluation hierarchy cannot accurately assess its own members because it evaluates a global population through local frames. An organization whose formal evaluation apparatus produces outcomes statistically indistinguishable from random allocation cannot be expected to perceive hidden structural couplings that its evaluation apparatus was not designed to detect.

Organizations do not passively fail to see. They actively select against the signals that would save them. The mechanism driving this selection is *dysmemic pressure*: the compound force that emerges when strategic communication degradation, adverse selection in idea markets, and transmission bias interact within environments shaped by compressed representations [McEntire, 2025a]. Signals optimized for internal fitness rather than correspondence with external reality—dysmemes—flood the information environment. The resulting equilibrium, where organizational reality decouples from external reality while remaining internally consistent, is stable because no individual can profitably deviate. The Cage, as we term this equilibrium, is the default state of any organization that compresses information to coordinate at scale. It requires no conspiracy, no malice, and no stupidity. It requires only compression and selection operating over time.

## 2 Why Existing Approaches Cannot Address Second-Order Ignorance

Existing tools fail at second-order ignorance, and they fail for a reason that better engineering cannot fix. The dominant paradigm in AI-assisted knowledge work is retrieve-then-reason. A user formulates a query. A retrieval system surfaces relevant documents. A language model synthesizes a response. The architecture assumes the user knows what to look for, can express it as a query, and will recognize the answer when it arrives. Retrieval-augmented generation, enterprise search, and chatbot interfaces all inherit this assumption. They are powerful tools for first-order ignorance, where the gap is known and the question is formulable. They are structurally incapable of addressing second-order ignorance, because no query is ever issued.

Multi-agent AI frameworks represent the current frontier, and they reproduce the problem at a different level. Frameworks like CrewAI, AutoGen, and LangGraph decompose complex tasks into subtasks assigned to specialized agents that communicate through message passing. The architecture mirrors a human organization: a coordinator distributes work, specialists execute, results flow back for integration. The designers are modeling what they know, and what they know is hierarchical task decomposition. The resulting system inherits the information-theoretic limitations of the structure it imitates.

Crawford and Sobel’s 1982 model of strategic information transmission proves that when a sender’s preferences diverge from a receiver’s by bias parameter  $b$ , the maximum number of distinguishable partitions in the communication is bounded by

$$N^* = \left\lceil -\frac{1}{2} + \frac{1}{2}\sqrt{1 + \frac{2}{b}} \right\rceil. \quad (1)$$

At  $b = 0$  (perfect alignment), the sender can communicate with arbitrary precision. At  $b \geq 1/4$ ,  $N^* = 1$ : the communication collapses to a babbling equilibrium where the receiver can extract no information from the sender’s message [Crawford & Sobel, 1982]. The bias need not be large. A project manager whose performance review depends on project velocity has a small but nonzero divergence from the engineering lead whose performance review depends on system

reliability. That divergence, multiplied across every communication relay in a reporting chain, produces measurable information loss.

The empirical magnitude of this loss has been quantified. Liberti and Mian’s study of hierarchical lending organizations found that as loan approval authority moves up the hierarchy, decision sensitivity to soft information (subjective assessments, local knowledge, character judgments) drops precipitously [Liberti & Mian, 2009]. The decay is not gradual. It exhibits a structural break between the second and third hierarchical levels, where reliance on soft information collapses and decisions revert to hard metrics alone. The hierarchy does not slowly degrade information. It kills subjective signal at a specific organizational seam, and everything above that seam operates on the skeleton of quantitative data stripped of the context that gave it meaning. The Data Processing Inequality provides the information-theoretic foundation: for any Markov chain  $X \rightarrow Y \rightarrow Z$ , the mutual information  $I(X; Z) \leq I(X; Y)$  [Cover & Thomas, 2006]. No post-processing of a compressed signal can recover information that was lost at the compression stage. If the raw signal is a pull request and the compressed signal is a status update, the status update has strictly less information. Processing the status update more carefully, with more sophisticated AI, through more elaborate agent architectures, cannot recover what was discarded.

Multi-agent frameworks optimize the processing of already-degraded signals. The coordinator agent reads the project manager’s summary, which has already lost the coupling information visible in the raw commit history. The specialist agent receives a task description abstracted away from the concrete artifact where the problem lives. The architecture is a digital recreation of the organizational reporting chain, complete with its information-theoretic limitations. The loss is structural, not computational. Better models processing worse signals cannot recover what was discarded at the compression stage.

Spence’s signaling theory sharpens the diagnosis [Spence, 1973]. Work artifacts—merged pull requests, passing test suites, deployed services, CI/CD pipeline logs—are costly signals in the economic sense. They require verifiable effort to produce. A team that claims “testing is complete” produces a cheap signal: the claim costs nothing and carries no binding commitment. A CI/CD pipeline log showing 847 passing tests across 12 modules is a costly signal: producing it required the tests to actually pass. The information content of costly signals is strictly higher than cheap signals because the cost of production provides a credibility guarantee that cheap signals lack. An AI system that reads status reports is reading cheap talk. An AI system that reads build logs is reading costly signals. The observation channel determines the information ceiling before any analysis begins.

Elliott’s empirical research on stigmergic collaboration identifies a phase transition in coordination mechanisms at approximately 25 participants [Elliott, 2007]. Below this threshold, direct social negotiation can sustain coherent collaboration. Above it, the combinatorial explosion of pairwise communication channels overwhelms the cognitive capacity of participants. The connection to Crawford–Sobel is formal: as group size  $N$  increases, the accumulated bias across communication relays pushes the effective  $b$  past the  $1/4$  threshold where  $N^* = 1$ . Direct communication collapses into babbling. Stigmergic coordination—indirect coordination through shared environmental traces—is the only mechanism that survives above the Elliott threshold, because it removes the sender-receiver interface where strategic degradation occurs. The participants interact with the shared environment rather than with each other. The environment cannot strategically distort.

### 3 Ambient Structure Discovery: The Paradigm

The alternative to retrieve-then-reason is a system that does not wait for queries because it does not need them. Ambient Structure Discovery operates on three principles that collectively address the structural limitations identified above.

The first principle is continuous ingestion of sematectonic traces. The term is Wilson’s, from the entomological literature on stigmergy [Wilson, 1971]. Sematectonic stimuli are physical modifications to the environment that carry information about the work that produced them, as distinct from marker-based stimuli, which are signals added to the environment specifically for communication purposes. Code commits, pull request diffs, CI/CD pipeline outputs, deployment manifests, and database migration files are sematectonic traces of software engineering work. They are indexical signs in Peirce’s semiotic framework: causally connected to the actions that produced them, carrying structural information about those actions regardless of whether anyone intended to communicate that information. The system reads these traces continuously, without being asked, and without requiring the traces’ producers to do anything differently.

The second principle is self-organizing topological specialization. Rather than assigning signals to pre-defined categories (which would require someone to know what categories exist, reintroducing first-order assumptions), the system routes signals through a competitive mesh of processing nodes. Each node develops affinity for signals it has successfully processed. Over time, nodes differentiate: one becomes expert in authentication-related traces, another in payment processing, a third in infrastructure configuration. This differentiation is not designed. It emerges from the interaction between signal content and routing dynamics, the way a river channel emerges from the interaction between water flow and terrain. The topology of the resulting mesh is itself a structural model of the organization’s actual work patterns, as distinct from its formal organizational chart.

The third principle is geometric pattern detection without semantic pre-specification. The mesh’s topology contains structural information that no individual node possesses. Two nodes that process signals with high spectral similarity (meaning they activate overlapping regions of the mesh) are structurally related regardless of whether their signals share vocabulary. A booking service and a reservation service that both touch the same database views will produce traces that route through overlapping mesh regions, making their coupling visible in the mesh geometry even if the word “booking” never appears in the reservation service’s codebase. The detection happens at the level of topological structure, which is why it can surface relationships that keyword search, semantic search, and even human inspection would miss. The relationship lives in the shape of the mesh that the content produced.

These three principles collectively address second-order ignorance because they remove the requirement for a query at every stage. Ingestion does not require anyone to specify what to ingest. Specialization does not require anyone to define specializations. Detection does not require anyone to hypothesize what patterns might exist. The system functions as a perceptual organ that sees organizational structure the way a retina sees light: continuously, without being asked, producing representations that are available for interpretation but independent of prior hypotheses about what to look for.

The paradigm addresses two distinct failure modes. The first is Discovery: hidden interdependencies that no one knows exist. The second is Normalized Deviance: known risks that everyone discusses and nobody acts on. Both are failures of organizational perception, but their spectral signatures differ (Section 4.5) and the mechanisms that sustain them are distinct (Section 6). A system designed for one must also address the other, because surfaced findings that are not structurally resolved decay from the first failure mode into the second (Section 7).

The theoretical foundations of this paradigm draw from multiple disciplines that have independently formalized aspects of the problem. Grassé’s stigmergy [Grassé, 1959], Grossberg’s Adaptive Resonance Theory [Grossberg, 1976], Beer’s Viable System Model [Beer, 1972], Shannon’s rate-distortion theory [Shannon, 1959], Simon’s bounded rationality [Simon, 1956], and Crawford and Sobel’s strategic communication model [Crawford & Sobel, 1982] each contribute formal results that map directly to architectural mechanisms. The following section describes the architecture and those mappings.

## 4 The Stigmergic Mesh

### 4.1 Signal Flow

A signal enters the system as a structured representation of a sematectonic trace. It carries a source identifier, a timestamp, an action type, and content extracted from the work artifact. The source and action type are open strings, not enumerated types. This is a deliberate design choice grounded in Wittgenstein’s observation that meaning is use [Wittgenstein, 1953]: an action called `deploy_to_staging` means what the organization’s deployment practices make it mean, not what a predefined ontology declares. The open taxonomy allows the system to discover categories that no designer anticipated, which is the precondition for surfacing second-order ignorance.

The signal first passes through a constraint filter that enforces organizational policy mechanically. Signals containing personally identifiable information, credentials, or content matching configurable sensitivity patterns are killed (discarded with audit trail), redacted (sensitive content removed, structure preserved), or passed unchanged. The filter operates as a deterministic gatekeeper that applies pattern-matching rules without discretion, ensuring that policy enforcement does not depend on the judgment of downstream components. In Beer’s Viable System Model, this filter functions as the System 2 coordination mechanism: an anti-oscillatory device that prevents pathological signals from propagating through the mesh [Beer, 1972].

The signal then enters a deduplication layer. Near-identical signals (same source, same action, same content within a configurable similarity window) are collapsed. This is variety attenuation in Beer’s framework: reducing unnecessary variety so that the mesh’s finite processing capacity is allocated to genuinely distinct signals rather than redundant copies of the same event.

### 4.2 Competitive Routing

The deduplicated signal enters the routing layer, which implements a competitive breadth-first search across all active workers. Each worker computes a familiarity score for the incoming signal from a weighted combination of term overlap, source affinity, temporal proximity, and signal credibility—where credibility implements Crawford–Sobel costly signaling directly in routing, weighting code commits above comments and diff magnitude above assertions. The first worker whose score exceeds its adaptive threshold accepts the signal. Routing terminates on acceptance; no exhaustive search over all workers occurs.

This is Simon’s satisficing, implemented as a routing mechanism [Simon, 1956]. The system does not evaluate all workers to find the optimal match. It finds the first match that exceeds a threshold and commits. Simon’s formal result establishes that satisficing outperforms optimization when evaluation is costly, the search space is combinatorial, and the environment is non-stationary. All three conditions hold for organizational signal streams: evaluating every worker against every signal would be quadratically expensive, the space of possible signal-worker assignments is combinatorial, and the organization’s structure changes continuously as teams form, projects launch, and priorities shift.

Gigerenzer’s ecological rationality research provides additional grounding [Gigerenzer, 2008]. The familiarity threshold is a fast-and-frugal heuristic in Gigerenzer’s taxonomy: a one-reason decision rule (“is this signal familiar enough?”) that exploits environmental structure (signals from related activities share vocabulary) to achieve performance that matches or exceeds exhaustive scoring. The bias-variance tradeoff explains why: in high-variance environments, simple rules that accept higher bias produce lower total error because they avoid overfitting to noise.

The threshold is adaptive. It scales with the worker’s fullness, a normalized measure of how much of its processing capacity has been consumed. As a worker accumulates signals and develops a richer term profile, its threshold rises. A worker that has processed hundreds of

authentication-related signals requires a higher familiarity score before accepting a new signal, because its existing context is already dense and the marginal value of an additional similar signal is lower. This adaptive scaling serves the same function as complement coding in Carpenter and Grossberg’s Fuzzy ART [Carpenter et al., 1991]: it prevents established categories from absorbing everything and collapsing into over-general representations. Without it, the first worker to develop any specialization would attract all subsequent signals, monopolizing the mesh and preventing differentiation.

### 4.3 Gap Detection and Worker Spawning

When a signal fails to exceed any worker’s adaptive threshold—when no existing node in the mesh finds it sufficiently familiar—the system creates a new worker and assigns the signal as its founding context. This is the mechanism that allows the mesh to grow in response to genuinely novel organizational activity.

The formal identity with Adaptive Resonance Theory is exact. In ART, an input vector  $I$  is compared against each stored category template  $w_J$  via the vigilance criterion:

$$\frac{|I \cap w_J|}{|I|} \geq \rho, \tag{2}$$

where  $\rho$  is the vigilance parameter [Carpenter & Grossberg, 1987]. If the match exceeds  $\rho$  for some category  $J$ , resonance occurs and the category’s template updates to incorporate the input (match-based learning). If no category satisfies the vigilance criterion, a reset wave propagates and the system recruits an uncommitted node as a new category.

Grossberg proved that match-based learning—updating templates only when resonance occurs—stabilizes in arbitrary non-stationary input streams [Grossberg, 1976]. This is the stability-plasticity guarantee that distinguishes ART from gradient-descent approaches. Backpropagation and its descendants use mismatch-based learning: they update weights precisely when the output differs from the target. This produces catastrophic forgetting when the input distribution shifts. The mesh’s workers update their profiles only when accepting signals (resonance), not when rejecting them. Rejected signals produce a weight-shift, a lossy impression that slightly adjusts the worker’s profile without full learning, but the fundamental update mechanism is match-based. The formal consequence is that the mesh’s category structure converges to a stable partition of the signal space even when the underlying organizational activity is non-stationary, which it always is.

The vigilance parameter  $\rho$  controls the granularity of categories. Low vigilance ( $\rho$  near 0) produces broad, abstract categories: a single worker for all “backend” activity. High vigilance ( $\rho$  near 1) produces tight, specific categories: separate workers for authentication, authorization, session management, and token refresh. In the implementation, each worker’s adaptive threshold serves as its vigilance parameter, rising with fullness to prevent established categories from absorbing everything. A separate, lower floor—the `gap_threshold`—governs the reset mechanism: when no worker’s vigilance criterion is satisfied, the system recruits a new node only if all scores fall below this floor, preventing spurious category creation from signals that are merely unfamiliar to the nearest worker but within the mesh’s existing representational capacity. Tuning the adaptive threshold adjusts the resolution of the organizational map the mesh produces; tuning the gap floor adjusts the mesh’s sensitivity to genuinely novel activity.

### 4.4 Learning Modes and Energy Dynamics

When a worker accepts a signal, it processes the content through one of three learning modes determined by the signal’s information density and the worker’s current state. RAG-indexed learning stores the full content for later retrieval, producing maximum information retention at maximum storage cost. Context-summarized learning compresses the content into the worker’s

existing context representation, preserving semantic structure while reducing storage. Weight-shifted learning adjusts the worker’s term profile without retaining the content itself, functioning as a lossy impression that influences future routing without expanding the worker’s knowledge base.

This tiered learning implements Bateson’s definition of information as “a difference that makes a difference” [Bateson, 1972]. A RAG-indexed signal makes a large difference: it adds retrievable knowledge. A weight-shifted signal makes a small difference: it nudges routing behavior. A rejected signal makes no difference to the rejecting worker’s state, which is precisely Bateson’s criterion for non-information.

Each worker carries an energy level that decays exponentially in the absence of incoming signals. Workers that stop receiving signals gradually lose energy until they fall below an archival threshold and are removed from active routing, their accumulated context preserved but no longer influencing signal flow. This mechanism implements Prigogine’s insight about dissipative structures [Prigogine, 1977]: systems far from thermodynamic equilibrium maintain internal order through continuous energy (information) input. When the input stops, the structure decays toward equilibrium, which in this context means undifferentiation. The signal stream is the energy that sustains the mesh’s organizational topology. If a project goes dormant, its workers archive. If it reactivates, new signals spawn new workers or reactivate archived ones.

We conjecture favorable Lyapunov stability properties for this energy system. Define  $V(x)$  as the total representation error across all active workers (the sum of mismatches between worker profiles and the signals they have processed). The energy decay suggests  $\dot{V}(x) < 0$  in the absence of perturbation: representation error decreases monotonically as stale workers archive and active workers refine. New signals perturb the system, but each perturbation is absorbed through the competitive routing mechanism, which directs the signal to the worker best equipped to incorporate it. Empirically, the system converges to a new equilibrium that incorporates the perturbation rather than oscillating or diverging; a formal proof remains future work.

## 4.5 Spectral Fingerprinting and Structural Detection

The mesh’s topology contains information that transcends any individual worker’s knowledge. Two signals that activate similar patterns of workers across the mesh (high familiarity scores from the same set of workers) are structurally related, regardless of whether they share vocabulary. The vector of familiarity scores that a signal receives from all workers constitutes its spectral fingerprint: its position in the mesh’s knowledge space.

This approach to structural identity draws from the GraphWave framework for graph-based structural role detection [Donnat et al., 2018]. In GraphWave, nodes in a graph are characterized by the heat-kernel signature of their local topology—specifically how a diffusion process initiated at the node propagates through its neighborhood. Nodes with similar diffusion patterns occupy similar structural roles even if they are distant in the graph. The mesh’s spectral fingerprint serves an analogous function: signals with similar activation patterns across workers occupy similar structural positions in the organization, even if they originate from different teams, different repositories, or different communication channels.

Anomalous structural patterns manifest as spectral energy shifts. Under normal operation, connected nodes in the mesh exhibit homophily: signals from related activities route through overlapping worker sets, concentrating spectral energy in low-frequency components (smooth variations across the mesh topology). When structural anomalies emerge—a new coupling between previously independent systems, a knowledge silo forming where information flow has been interrupted, a coordination fracture between teams that should be collaborating—the local homophily breaks. High-frequency components appear in the spectral energy distribution, representing abrupt differences between adjacent regions of the mesh. The Beta Wavelet Graph Neural Network framework formalizes this [Tang et al., 2022]: standard graph convolution networks are low-pass filters that smooth signals and hide anomalies, while band-pass filters focused

on high-frequency bands reveal the structural discordances where anomalies reside.

The spectral energy right-shift—an increase in high-frequency spectral energy—serves as the mesh’s algedonic signal in Beer’s Viable System Model. The algedonic channel is a pain/pleasure signal that bypasses the normal reporting hierarchy to reach System 5 (policy) directly when organizational viability is threatened. The spectral right-shift bypasses intermediate interpretive layers: it is a geometric signal derived from the mesh’s topology, not a semantic judgment about what the anomaly means. The meaning is for System 5—the human recipient—to determine. The mesh’s job is to make the anomaly perceptible.

The converse spectral signature—eigenvalue *collapse*, a left-shift toward low-frequency concentration—carries distinct diagnostic meaning. Where the right-shift signals unexpected heterophily (Discovery: structural couplings that should not exist), the left-shift signals suspicious homophily (Normalized Deviance: independent evaluation modes merging into consensus). When the mesh’s correlation matrix loses rank, the spectral energy concentrates in fewer, lower-frequency modes. Topics that were previously discussed with distinct framing converge toward uniform language. Variance in reports decreases while variance in outcomes does not. Dimensionality drops. The mesh’s topology is compressing, and the compression is a symptom of conformity rather than convergence. Section 6 formalizes this signature and its organizational interpretation.

## 4.6 Insight Agents and Consensus

The spectral detection mechanisms identify where structural anomalies exist. Insight agents interpret what those anomalies might mean. Four agent tiers operate at increasing levels of abstraction: indexers that catalog and cross-reference signals, surfacers that evaluate the implications of indexed patterns, correlators that identify relationships across worker boundaries, and profilers that characterize the organizational entities (people, teams, systems) whose activity produces the observed patterns.

Each agent tier receives pre-structured observations rather than raw data. The observation context includes the relevant workers’ term profiles, the signals that triggered the observation, the spectral fingerprints of the involved regions, and the specific structural question being asked (what does this coupling mean? who depends on this interface? what would break if this node were removed?). The agents produce assessments with explicit confidence levels and reasoning traces.

No single agent’s assessment is surfaced directly to the human operator. Assessments enter a consensus mechanism where multiple agents evaluate the same structural question independently. Consensus is reached when a configurable threshold of agents agree on the assessment’s direction and the aggregate confidence exceeds a minimum. Cross-source diversity in consensus is an emergent property of ART’s specialization dynamics: workers that develop affinity for signals from a particular source (code repositories, issue trackers, communication channels) naturally produce source-diverse quorums when multiple workers evaluate the same finding. The current architecture does not enforce source diversity as a quorum constraint, relying instead on the tendency of competitive routing to produce heterogeneous specialization. This mechanism addresses Crawford and Sobel’s partition resolution: by requiring multiple independent evaluations, the system increases the effective partition count of the communication channel between the mesh and the human operator, producing finer-grained distinctions than any single agent could reliably achieve.

The consensus mechanism also serves as the mesh’s System 3 (control) in the VSM mapping. System 3 optimizes resource allocation across operational units. The consensus layer decides which insights warrant human attention—a scarce resource subject to the attention economics that Ocasio’s Attention-Based View of the Firm identifies as the binding constraint on organizational decision-making [Ocasio, 1997]. Surfacing too many insights produces alert fatigue; surfacing too few misses critical structural shifts. The consensus threshold is the tuning

parameter that navigates this tradeoff.

The mesh’s output is itself a costly signal. The Crawford–Sobel babbling equilibrium applies to the mesh’s output channel as much as to the organization’s internal channels: if the mesh produces cheap talk (narrative reports, ungrounded alerts), the receiver has no reason to update. The mesh escapes this trap through artifact generation. Its output is a deposited artifact in the work environment—a sematectonic trace of the mesh’s own processing, verifiable against the signal stream that produced it. The assessment includes the specific signals, the workers that processed them, the spectral fingerprints that triggered detection, and the consensus scores that passed the threshold. Producing this artifact required the full mesh pipeline. Fabricating it would require fabricating a consistent signal stream, routing history, and spectral analysis. The cost of production provides the credibility guarantee that transforms the mesh’s output from cheap talk into a costly signal that merits attention.

#### 4.7 Architectural Completeness

The architecture maps to Beer’s Viable System Model as follows. Workers processing signals in their specialized domains constitute System 1 (operations). The routing layer’s competitive accept-reject mechanism, together with the constraint filter and deduplication layer, constitutes System 2 (coordination): the anti-oscillatory mechanisms that prevent pathological signal distribution. The consensus mechanism and agent assessment layer constitute System 3 (control): optimization and resource allocation across operational units. The insight agents, spectral analysis, and gap detection mechanisms constitute System 4 (intelligence): the scanning function that monitors the boundary between the organization and its environment, detecting threats and opportunities.

System 5 (policy) is the human operator who receives surfaced insights and decides organizational response.

This distribution is deliberate. Systems 1 through 4 are computational. System 5 is human. The architecture does not attempt to close the loop autonomously, and this is a design virtue rather than a limitation. The Conant-Ashby Good Regulator Theorem states that every good regulator of a system must contain a model of that system [Conant & Ashby, 1970]. The mesh *is* the model. Traditional organizational charts are “bad models” in Conant-Ashby’s sense: they lack requisite variety (they are too simple to capture the organization’s actual complexity) and isomorphism (they do not match the organization’s actual topology). The mesh captures the heavy-tailed, high-entropy, topologically complex reality of how work actually flows through the organization.

A system that included its own System 5 would be an autonomous agent capable of overriding human judgment about organizational policy. The organizational incompleteness argument—that formal systems cannot fully characterize their own properties from within—applies to the mesh itself. A mesh that decided its own policy would be subject to its own blind spots, which it could not observe from within its own operational closure. Delegating System 5 to a human operator preserves the second-order observation relationship that gives the architecture its theoretical power. The mesh is a prosthetic for organizational perception. Luhmann’s framework would call it a “system for observing systems,” and the gap between observation and action is where human intelligence properly resides.

#### 4.8 Self-Observing Coherence

The mesh’s anomaly detection apparatus generalizes to monitoring its own input channel. The signal stream arriving at the mesh is itself an organizational signal environment subject to compression, selection, and decay—a source can go silent because the organization stopped producing traces in that channel, because an integration failed, or because access was revoked.

These are structurally identical to the organizational blind spots the mesh is designed to detect: information that was present and ceased to be, without anyone noticing.

When a source goes silent asymmetrically—code repository activity ceases while issue trackers and communication channels remain active on a workday—the mesh’s own topology deforms. Workers specialized in the silent source lose energy; others do not. The spectral fingerprints shift. Rather than building a separate monitoring system to detect this deformation, the mesh reads its own state: worker energy differentials, fingerprint divergence, and source distribution entropy. The same detection apparatus that identifies organizational blind spots identifies degradation in its own observation channel.

The cadence monitor discretizes observation history into day-of-week bins with hierarchical fallback: day-specific baselines, weekday/weekend baselines, and all-time baselines. Anomaly detection compares current source distribution entropy against the baseline for the corresponding period rather than against a global average. This eliminates false positives from predictable organizational rhythms—weekend quiet, sprint-boundary surges, standup spikes—while preserving sensitivity to genuine asymmetric degradation. Spectral resolution grows with observation depth. A first run uses the all-time baseline. After one week, day-of-week baselines emerge. After a sprint cycle, sprint-phase baselines become available. The system’s self-knowledge improves with runtime without architectural changes, following the same match-based learning dynamics that govern worker specialization.

Detection uses an asymmetry score:

$$A(E, v) = \sqrt{\frac{|E_{\text{baseline}} - E_{\text{current}}|}{E_{\text{baseline}}}} \cdot \sqrt{\frac{v_{\text{current}}}{v_{\text{baseline}}}} \quad (3)$$

where  $E$  is source distribution entropy and  $v$  is signal volume. When entropy drops while volume holds—asymmetric compression, the pathological case where one source dies while others continue— $A$  amplifies. When both entropy and volume drop—symmetric quiet, the benign case of uniform reduced activity—the volume term suppresses. The gradient is continuous; no threshold tuning is required.

## 5 The Scoring Function

The mesh produces structural observations continuously. The question is which observations warrant human attention. A scoring function prioritizes findings by four components, each grounded in a distinct formal tradition. The components are multiplicative, not additive, because each represents a necessary condition: a finding with zero bridge distance is already known, a finding with zero entity confidence cannot be trusted, a finding with zero risk coupling does not matter, and a finding with zero communicability cannot be transmitted. An additive function would allow a high score when any component is zero, which is wrong—a perfectly communicated finding about a resolved dependency is worthless. The multiplicative form implements logical conjunction: every component must be nonzero for the finding to warrant attention.

$$S(f) = d_{\text{bridge}}(f) \cdot c_{\text{entity}}(f) \cdot r_{\text{coupling}}(f) \cdot \gamma(f), \quad (4)$$

where  $d_{\text{bridge}}$  is the topological bridge distance,  $c_{\text{entity}}$  is entity resolution confidence,  $r_{\text{coupling}}$  is risk coupling, and  $\gamma$  is communicability.

*Bridge distance*  $d_{\text{bridge}}(f)$  measures how far outside the organization’s known subspace  $K$  the finding lives. Burt’s structural holes theory provides the conceptual grounding [Burt, 2004]: a finding that spans a gap between disconnected clusters in the communication graph occupies a position of low network constraint and high effective size—precisely where Granovetter’s non-redundant information flows [Granovetter, 1973]. The bridge distance is computed as the

shortest path in the communication graph  $G_C$  between the actors named in the finding, normalized through a sigmoid that maps distance to  $[0, 1]$ . The finding itself constitutes the  $G_D$  edge—the dependency that the finding identifies. High  $G_C$  distance between actors connected by a finding is therefore direct evidence of  $G_D/G_C$  divergence: a structural coupling exists that the communication topology does not span.

Formally, bridge distance is the divergence between two graphs: the dependency graph  $G_D$  (where code references, API calls, data flows, and deployment dependencies actually connect components) and the communication graph  $G_C$  (where discussion, review, and coordination connect the teams responsible for those components). Where  $G_D$  has an edge and  $G_C$  does not, a structural dependency exists that no one is talking about. This is load-bearing silence—not the absence of something to say, but the absence of communication where a dependency demands it. Bridge distance is maximal precisely at these points: components that are structurally coupled in the work but disconnected in the discourse. The mesh detects this divergence because it ingests both graphs simultaneously—sematectonic traces from code produce  $G_D$ , while traces from reviews, issues, and channels produce  $G_C$ . The gap between them is where Discovery findings live.

*Entity confidence*  $c_{\text{entity}}(f)$  measures the system’s confidence that the entities involved (people, teams, services, repositories) are correctly resolved. Entity resolution across heterogeneous trace sources is inherently ambiguous. A finding connecting “the auth service” to “the identity provider” is valuable only if those names map to the correct organizational entities. Confidence decays with ambiguity: multiple candidate resolutions reduce the score proportionally.

*Risk coupling*  $r_{\text{coupling}}(f)$  measures the topological proximity of the finding to a control surface—a deployment boundary, a financial transaction, a safety-critical path. The communicability matrix from spectral graph theory [Estrada & Hatano, 2008] provides the formal mechanism: the  $(i, j)$  entry of  $e^A$ , where  $A$  is the adjacency matrix, counts the weighted sum of all walks between nodes  $i$  and  $j$ , giving a natural measure of how closely a finding in the mesh is coupled to a node representing operational risk.

*Communicability*  $\gamma(f)$  replaces the binary gate implied by Crawford–Sobel with a continuous function:

$$\gamma(f) = \sigma(N^*(b_f) - 1), \quad (5)$$

where  $\sigma$  is a sigmoid and  $N^*(b_f)$  is the Crawford–Sobel partition count for the estimated bias  $b_f$  along the communication channel between the finding’s location and the human recipient. In practice,  $b_f$  is not directly observable; it is estimated from the finding’s specificity (the number of named actors and resolved bridge distances) and confidence, producing a proxy for  $N^*$  that preserves the sigmoid’s qualitative behavior. At  $N^* \geq 3$ , communicability is approximately 1: the finding can be transmitted with useful precision. At  $N^* = 1$  (babbling equilibrium), communicability approaches zero but does not reach it—the gate is continuous, not binary, because even degraded channels carry some information through costly signaling.

These four components define a space in which findings distribute across four quadrants (Table 1).

Discovery findings have high  $d_{\text{bridge}}$  and high  $r_{\text{coupling}}$ : they are structurally remote from the organization’s known knowledge and close to something that matters. Normalized Deviance findings have low  $d_{\text{bridge}}$  and high  $r_{\text{coupling}}$ : they are familiar—the organization already discusses them—but coupled to risk that is not being managed. The scoring function surfaces both, but the detection mechanisms and organizational implications differ.

A structural property of the scoring function deserves emphasis: all four components are topological. Bridge distance is a graph metric. Entity confidence is a resolution property of the organizational ontology. Risk coupling is computed from the communicability matrix of the mesh adjacency. Communicability is derived from the Crawford–Sobel partition count. None of these components measure satisfaction, sentiment, or self-report. They measure the shape of the organization’s work. This makes the scoring function immune to the Prendergast capture

Table 1: The four quadrants of organizational perception.

	<b>High Risk Coupling</b>	<b>Low Risk Coupling</b>
<b>High Bridge Distance</b>	<b>Discovery.</b> Hidden interdependencies near control surfaces. The mesh’s primary target.	<b>Noise.</b> Structural surprises that do not touch operational risk. Correctly filtered out.
<b>Low Bridge Distance</b>	<b>Normalized Deviance.</b> Known risks that everyone discusses and nobody acts on. Section 6.	<b>Ambient.</b> Known conditions with low risk coupling. Correctly ignored.

dynamic [Prendergast, 1993], in which agents distort their reports to match what they believe the principal wants to hear. An engineer cannot reduce the bridge distance of a finding by expressing concern about it. A manager cannot lower the risk coupling by adding the risk to a dashboard. The only way to change the score is to change the underlying topology—to actually resolve the structural coupling, close the communication gap, or reduce the dependency. The scoring function is not gameable by the mechanisms that dysmemic pressure employs, because dysmemic pressure operates on representations and the scoring function operates on structure.

## 6 Normalized Deviance Detection

### 6.1 The Approach to Babbling

Vaughan’s concept of normalization of deviance describes the process by which an organization redefines previously unacceptable risk as acceptable through repeated exposure without consequence [Vaughan, 1996]. Each incident that does not produce catastrophe shifts the baseline. The anomaly becomes expected variation, then routine, then invisible. Vuori and Huy document the complementary dynamic in upward communication: middle managers, afraid of top managers, systematically distort reports until the board receives consistent reassurance while the organization’s competitive position deteriorates beyond recovery [Vuori & Huy, 2016].

In the Crawford–Sobel framework, Normalized Deviance is the transition toward babbling equilibrium. The partition count  $N^*$  is not a static property of the channel; it changes over time as preferences, incentives, and norms evolve. When an organization first encounters a risk, discussion may be precise—specific failure modes, specific probabilities, specific remediation plans. As the risk persists without consequence, the discussion coarsens. “Component X fails under condition Y” becomes “we have a component X issue” becomes “component X is on the risk register.” The partition count is declining. The communication channel is losing resolution. The risk coupling  $r_{\text{coupling}}$  is unchanged or increasing, but the information content of the communication about it is decreasing.

The organizational response becomes the organizational substitute for action. The discussion itself becomes the deliverable. Status meetings review the risk. Dashboards display it. Quarterly reports mention it. The forms of attention continue. The substance drains away. Talk substitutes for action, and the substitution is invisible because the organization’s measurement systems track the talk.

## 6.2 Eigenvalue Collapse

The spectral signature of Normalized Deviance is the opposite of Discovery’s right-shift. Where Discovery manifests as unexpected high-frequency energy (heterophily between mesh regions that should be smooth), Normalized Deviance manifests as eigenvalue *collapse*—a left-shift toward low-frequency concentration that signals the loss of independent evaluation modes.

Consider the mesh’s correlation matrix  $C$  across workers processing signals about a particular risk domain. Under healthy conditions, different workers maintain distinct perspectives—different teams discuss the risk differently, emphasize different aspects, propose different mitigations. The matrix  $C$  has full rank. Independent modes of evaluation exist. As Normalized Deviance progresses, framing converges. Teams that previously discussed the risk with distinct vocabulary begin using the same phrases. Reports that previously offered different assessments begin reaching the same conclusions. The correlation matrix loses rank. Eigenvalues that were previously distinct merge. The effective dimensionality of the organization’s perception of the risk decreases even as the risk itself continues.

The observable signatures are specific and measurable: communication diversity declining across teams discussing the same topic, topic framing converging toward shared vocabulary, variance in assessments decreasing while variance in outcomes remains constant or increases. The mesh can track these signatures because it ingests the raw traces from multiple teams simultaneously. The convergence is visible in the mesh’s topology as regions that were previously distinct begin to overlap.

A related mechanism that accelerates eigenvalue collapse is *taxonomic competence*: the development of fluent shared vocabularies for recurring failures. When an organization develops standardized category names, established triage flows, and practiced diagnostic routines for a class of failures, the competence at categorizing the failure prevents the failure from registering as a crisis. The shared failure vocabulary is simultaneously the linguistic compression and the framing convergence—the taxonomy *is* the collapsed eigenspace. The organization becomes expert at describing the problem, and the expertise substitutes for solving it. Naming replaces resolving. The mechanism is self-reinforcing: each use of the standardized vocabulary confirms the category’s legitimacy as a managed condition rather than an unresolved crisis, further reducing the pressure gradient that would otherwise force resolution.

## 6.3 Phase Locking and Rogue Waves

The catastrophe that emerges from Normalized Deviance is rarely a single large risk materializing. It is individually modest risks synchronizing. Conformity bias is the synchronizing force.

The dynamic follows coupled oscillator mechanics. Each normalized risk is a frequency in the organization’s risk spectrum. The degraded component, the environmental sensitivity, the schedule pressure, the delivery culture—each one a quiet oscillation. The power spectrum shows the individual risks. It does not show the correlations between them. The correlations are where the rogue wave forms.

The coupling is reflexive. The narrative that a risk is acceptable changes behavior: inspections become less thorough, mitigations are deferred, workarounds become standard procedure. Changed behavior reinforces the narrative: if the workaround holds, the risk was indeed acceptable. The system locks into a positive feedback loop that appears stable because the variance in reports is decreasing—everyone agrees the risk is managed—while the variance in outcomes is increasing as independent safety margins erode in parallel. The NK landscape model formalizes why this is catastrophic [Kauffman, 1993]: when the coupling between components  $K$  increases, the fitness landscape becomes increasingly rugged until a *complexity catastrophe* makes the system globally impossible to optimize even though it appears locally optimizable from any single vantage point.

A phase-lock detector monitors whether independent risks are synchronizing. The required measurements are: topic-variance over time (is the framing of a known risk converging across teams?), action-correlation (are discussions producing state changes—commits, ticket closures, deployments—or has mention frequency decoupled from action frequency?), and channel migration (are discussions moving from public to private channels while the underlying work activity continues unchanged?). The divergence between the communication graph and the dependency graph is the signal: when the first is losing edges while the second is not, the organization is losing visibility into its own risk landscape.

Two distinct mechanisms produce organizational blind spots, and the distinction matters for detection. The first is passive: blind spots settle where executive attention does not reach, the way debris accumulates on a highway shoulder. The road alligator sits in the complement of the attention surface—nobody steered it there; it arrived because nobody steered it away. These passive blind spots are detectable through load-bearing silence (Section 5): structural dependencies that exist in the work but generate no corresponding communication.

The second mechanism is active. Dymemic pressure does not merely fail to illuminate uncomfortable topics; it generates selection pressure that pushes evaluation vectors *away* from them. The more dangerous the truth, the stronger the repulsive force. This explains channel migration: when a risk becomes politically charged, discussion does not simply fade—it relocates from public channels to direct messages, from written records to verbal hallway conversations, from explicit risk registers to implicit shared understanding among insiders. The observable signature is specific: topics whose public-channel frequency declines while the underlying work activity (commits, deployments, dependency changes) continues unchanged. The communication graph is losing edges where the dependency graph is not. Passive settling produces uniform silence. Active repulsion produces *shaped* silence—gaps with a characteristic pattern that correlates with political sensitivity, and the mesh can detect the shape.

## 6.4 Linguistic Compression as Leading Indicator

The transition from active risk management to Normalized Deviance has a measurable linguistic signature. The shift from direct language to hedged, nominalized, passive prose precedes operational failure—it is a leading indicator, not a lagging one.

Hedging density increases: “will fail” becomes “may experience degradation.” Passive voice ratio rises: “the team decided to defer” becomes “the decision was made to defer,” removing the agent from the action. Nominalization frequency grows: “the system malfunctions” becomes “malfunction events have been observed,” converting active processes into static nouns. Specificity of commitments decreases: “we will replace the component by March” becomes “remediation efforts are ongoing.” The shift from “malfunction” to “anomaly” to “observation” is Normalized Deviance made lexical.

These features are measurable in the text that the mesh ingests. The measurement architecture follows the same principle as the mesh itself: the LLM agents that already process every signal through reflection and evaluation possess the contextual understanding that rigid pattern-matching cannot. An agent reading “the deployment was successful” knows this is not backward-looking language; an agent reading “anomaly” in an engineering context knows whether it is precise terminology or bureaucratic substitution. The assessment is folded into the existing agent calls—the correlator already reasons about batches of signals, and assessing compression costs near-zero marginal tokens while gaining the contextual discrimination that no word list can match. The agents assess eight properties across two tiers.

The first tier captures rhetorical compression: hedging density (the degree to which commitments are qualified, conditions multiplied, and accountability diffused), passive voice saturation (the extent to which agents are removed from actions, converting “the team decided” into “the decision was made”), nominalization (the conversion of active processes into static nouns, making ongoing failures sound like managed conditions), and specificity (the presence or absence

of concrete referents—dates, names, numbers, direct verbs—that would make a commitment falsifiable). These four compose a compression index on  $[0, 1]$ .

The second tier captures information-theoretic properties validated against empirical corpora: lexical diversity (unique tokens divided by total tokens—pre-constraint organizational filings show LD in the 0.16–0.19 range; post-constraint filings compress to 0.13–0.15), Shannon entropy (relative change over time is the signal, since absolute values scale with document length), temporal orientation (the balance between forward-looking and backward-looking language on  $[-1, +1]$ , where sustained negative values indicate an organization that has shifted from describing its future to cataloging its past), and bureaucratic density (the prevalence of euphemistic substitution—language whose function is to redescribe a problem in terms that make it sound managed rather than unresolved).

An exponential moving average tracker monitors compression trends per topic and per channel across sessions, flagging channels whose compression index is increasing as active normalization candidates. A “Born Caged” detector records the initial compression index for each channel when first observed; channels that begin above a threshold entered the mesh’s observation window already compressed—analogue to the modern regulatory filing that arrives pre-loaded with defensive boilerplate, limiting the temporal baseline available for trend detection.

Changes in these metrics across the signal stream for a particular risk domain over time constitute a leading indicator for eigenvalue collapse. The linguistic compression precedes the spectral compression, because the language changes before the topology does. The words shift first; the structure follows. This lead-lag relationship has empirical precedent in aerospace safety: the Rogers Commission [Rogers Commission, 1986] and the Columbia Accident Investigation Board [CAIB, 2003] documented that the normalization of anomaly language—O-ring erosion redescribed as “acceptable risk,” foam strikes redescribed as “in-family”—preceded the operational failures by years, and that seventeen years of reform between disasters failed to reverse the linguistic pattern because the reforms targeted procedures rather than the communication culture that the procedures were expressed in.

## 7 The Decay Path

### 7.1 Discovery to Normalized Deviance Migration

The mesh surfaces a finding. The organization responds: “interesting.” A ticket is created. Nobody picks it up. It appears in the next retrospective. It is discussed. It is acknowledged. Six months later it is a known risk that everyone recognizes and nobody owns. The conversion from second-order ignorance (nobody knew) to first-order ignorance (everybody knows) succeeded. The conversion from first-order ignorance to action did not.

This migration is the default trajectory. Discovery findings decay into Normalized Deviance through a predictable sequence: surfacing, acknowledgment, deferral, normalization. The scoring function (Equation 4) captures the transition:  $d_{\text{bridge}}$  decreases as the finding becomes known (it is no longer structurally remote from the organization’s awareness), while  $r_{\text{coupling}}$  remains unchanged (the risk has not been mitigated). The finding migrates from the Discovery quadrant to the Normalized Deviance quadrant in Table 1.

The relocation thesis applies: surfacing a finding without structurally resolving it relocates the second-order ignorance from “nobody knows” to “everybody knows and nobody acts.” The organization now believes the problem has been addressed because it has been surfaced, discussed, and documented. The documentation becomes the organizational response. The problem persists, but the organization’s model of itself has updated to include the problem as a managed risk, which reduces the perceived urgency of actual mitigation.

## 7.2 Dysmemic Pressure on Countermeasures

The finding that the mesh surfaces is itself a countermeasure to organizational blindness. Countermeasures face dysmemic pressure. “Any structure that counterweights dysmemic pressure faces continuous pressure toward absorption back into the dysmemic equilibrium” [McEntire, 2025a].

Proposition 4 from the companion work on dysmemic pressure is directly applicable: interventions that change expressed norms without changing payoff structures exhibit initial improvement followed by regression to the pre-intervention equilibrium. A finding that changes awareness without changing incentives will decay. The organization’s acknowledgment of the finding changes the expressed norm (“we take this seriously”). It does not change the payoff structure (nobody is rewarded for fixing it, nobody is penalized for deferring it). The prediction is regression: the finding will lose salience until the next crisis forces rediscovery.

Organizational change initiatives fail at rates between 60 and 80 percent [Beer & Nohria, 2000], and the rate is predictable: the initiatives target behavior without targeting the selection environment that produces the behavior. Reform fails because reform proposals are themselves signals subject to the selection environment they aim to change. The finding is a signal. It enters the organizational information environment and competes with other signals for attention, resources, and action. Signals fitting the current frame gain resources. Signals challenging the current frame face friction. The finding, which by definition challenges the frame (it surfaced something the frame missed), faces maximum resistance from the selection environment whose limitations it exposes.

## 7.3 Stigmergic Deposition as Anti-Decay

The mesh’s delivery mechanism is designed to resist the decay path. Three delivery modes exist, and the distinction between them is sharper than it first appears.

*Alert-based delivery* faces attention economics. Ocasio’s framework establishes that situated attention—what people attend to in specific contexts—is the binding constraint on organizational decision-making [Ocasio, 1997]. Alerts interrupt. They compete with the current task for cognitive resources. They are easy to dismiss, because dismissal removes the interruption. The selection pressure on alerts is fierce: only alerts that are immediately actionable survive; everything else is triaged into a backlog that functions as a graveyard.

*Annotation-based delivery*—depositing findings as comments on pull requests, issue trackers, or other existing work artifacts—appears categorically different but collapses into alert-based delivery in digital environments. The reason is geometric. Physical stigmergy works because the environment has spatial structure: the ant encounters the pheromone because walking past it is part of doing the work. The encounter is a side effect of the activity, not an interruption of it. Digital environments lack this geometry. A pull request comment, a ticket annotation, and a log entry are all rows in databases. The “encounter” is determined by which interface a person opens, not by physical proximity to the work. A bot comment on a pull request is structurally an alert that arrives through a different channel. Engineers develop antibodies to automated annotations faster than to any other form of organizational communication. After the third bot comment that is not immediately relevant, every subsequent one is resolved without reading. The channel saturates and the finding decays through exactly the attention economics the delivery mechanism was designed to circumvent. The converse is also observable: organizations stigmergically deposit their own artifacts that *normalize* deviance rather than combat it. A dedicated communication channel for a class of failures that should be rare enough not to need one is a sematectonic trace that shapes future behavior—toward managing the failure efficiently rather than eliminating it. The channel makes the failure cheaper to process, reducing the pressure gradient that would otherwise force resolution. Stigmergic deposition is a mechanism, not a virtue. It can entrench dysfunction as readily as it can surface it.

*Artifact generation* is the mode the theory demands. The finding is not deposited as a comment about a gap. The finding *is* the organizational knowledge whose absence constituted the gap. A parent ticket linking eight calendar synchronization failures across five vendors, with specific unanswered questions and a recommended triage owner, is not a notification—it is a coordination artifact that did not previously exist. An interface contract documenting an undocumented dependency is not commentary—it is the organizational model that was missing. A dependency diagram making a hidden structural relationship explicit is not a report—it is the regulatory model that the Conant-Ashby theorem requires [Conant & Ashby, 1970]. The organization becomes a better regulator of the system because the artifact now exists as part of the organization’s regulatory apparatus. The distinction between a comment (“this coupling exists”) and a model artifact (the interface contract specifying the coupling) is the distinction between describing a gap and closing it.

The correct intermediate architecture, while artifact generation capability matures, is a mirror surface: an independent observation point that leadership consults deliberately, in a context where attention for organizational health is already allocated. The finding registry, with its decay tracking and actionable structure, is encountered during the morning review or the weekly sync—contexts where situated attention is oriented toward organizational health rather than task execution. This satisfies the Mirror properties of Section 8.1 more cleanly than annotation-based delivery: the observation surface is structurally decoupled from the interrupt-driven flow of execution.

A computational system that generates organizational artifacts occupies a unique position in the Mirror architecture. It can ask the questions that no human asks—not because the questions are difficult, but because asking them carries political cost. “This policy contradiction has already cost money and nobody owns it” is a statement that a human team member pays for politically. A non-person pays nothing. Retaliation against a process is incoherent; you can only turn it off, which is a visible organizational decision rather than a quiet political one. The mesh’s non-personhood provides structural insulation that no human mirror can achieve, making it the natural vehicle for the hard questions that dysmemic pressure selects against.

## 8 The Mirror Architecture

### 8.1 Three Properties of a Mirror

The organizational theory literature has long identified the need for external observation without providing a scalable mechanism for it. Consultants observe intermittently, at high cost, with limited access to operational data, and subject to their own incentive structures. Auditors observe periodically, against checklists that the organization helped design, through channels the organization controls. Board members observe through management presentations that have already passed through the compression and selection pipeline that creates the blind spots.

Companion work on rigorous inquiry identifies three properties that any structure must possess to serve as an effective mirror for organizational self-observation [McEntire, 2026]. *Insulation*: the mirror’s survival cannot depend on the approval of those it observes. Its budget, tenure, and authority must be protected from retaliation. *Access*: the mirror must see information that primary frames discard—the raw data before compression, the complaints that never reach the dashboard, the variance that metrics smooth away. *Authority*: the mirror must be able to surface findings without passing through the filters it is examining. If the report goes to the manager who caused the problem, the mirror is broken.

Most corporate oversight fails because it lacks one of the three properties. The internal audit that reports to the CFO lacks insulation. The board that sees management presentations lacks access. The whistleblower hotline that routes to HR lacks authority. Even dedicated contrarian analysis units—chartered specifically to challenge prevailing assessments—fail when

their findings must pass through the same evaluation hierarchy they are designed to question. Insulation and access without authority that bypasses the existing fitness landscape reproduces the selection pressure the unit was designed to circumvent.

The mesh satisfies all three. It does not participate in the communication hierarchy (insulation—its operation does not depend on organizational approval of its findings). It reads sematectonic traces rather than reports (access—it observes the work product itself, which is the information that the reporting chain discards). And artifact generation bypasses the reporting chain (authority—the mesh produces organizational knowledge artifacts directly rather than transmitting findings through the channels whose limitations the findings expose). Its non-personhood provides a form of insulation unavailable to human mirrors: it can surface politically uncomfortable findings without facing the retaliation that dysmemic pressure directs at human messengers.

## 8.2 Orthogonal Evaluation

The mesh provides evaluation in dimensions orthogonal to the organization’s own evaluation criteria. The endogenous evaluation criteria—the metrics, dashboards, status reports, and review processes that the organization uses to assess itself—span a subspace  $K$  of the full evaluation space. This subspace is determined by the organization’s history, priorities, and frame. The mesh generates evaluation vectors outside  $K$  through topological bridging: connecting structural regions that the organization’s own evaluation criteria do not connect.

Lawvere’s fixed-point theorem provides the formal constraint [Lawvere, 1969]: in a sufficiently expressive system, complete self-evaluation is impossible. Every evaluation requires a distinction, and the distinction being used to evaluate cannot itself be evaluated at that moment. The organization’s blind spots are not random—they are structurally determined by the evaluation criteria the organization uses. The null space of the organization’s evaluation matrix is where second-order ignorance lives.

The mesh’s topological bridge distance (Equation 4) is a measure of how far outside  $K$  a finding lives. High bridge distance means the finding connects structural regions that the organization’s own instruments cannot connect, because the instruments are made of the same material as the knowledge they are designed to assess.

A limitation follows directly from the formal constraint. The relocation thesis, supported by evidence from immunological self-evaluation and organizational reform [McEntire, 2025a], establishes that each proposed fix moves the blind spot rather than eliminating it. Spencer-Brown’s re-entering form produces temporal oscillation rather than static contradiction [Spencer-Brown, 1969]: the blind spot is not fixed in place. It rotates. The mesh cannot find the blind spot once and declare victory. It must track the blind spot’s movement across successive measurement periods, detecting when a previously surfaced finding has been absorbed into the organizational frame without being resolved.

## 8.3 The Discipline of Rigorous Inquiry, Computationally

Rigorous inquiry is defined by three properties: independence (the check originates outside the frame being tested), specificity (the check generates predictions specific enough to be wrong), and closure (the result propagates back to modify the frame) [McEntire, 2026]. The mesh implements all three computationally.

*Independence:* the mesh reads sematectonic traces, not reports. The traces are produced by the work itself, not by the reporting processes that compress and select the work’s representation. The observation channel is structurally independent of the communication channel where dysmemic pressure operates. A manager who does not report a problem creates a gap in the organization’s self-model. The manager’s code commits, deployment patterns, and build failures still produce traces that the mesh can observe.

*Specificity*: the scoring function (Equation 4) produces quantified, falsifiable claims. A finding with  $d_{\text{bridge}} = 4.2$ ,  $c_{\text{entity}} = 0.91$ ,  $r_{\text{coupling}} = 0.78$ , and  $\gamma = 0.85$  is a specific prediction that can be evaluated against organizational reality. The bridge distance is either accurate (the coupling exists) or it is not. The risk coupling is either real (the finding touches a control surface) or it is not. The scoring function creates surface area for reality to push back.

*Closure*: artifact generation forces findings into organizational reality. The finding does not remain in a report that can be filed and forgotten. The mesh produces the organizational knowledge whose absence constituted the gap—parent tickets, interface contracts, dependency diagrams—artifacts that practitioners encounter in the course of their work because they *are* part of the work. A finding registry tracks decay: findings that are not acted on within configurable windows are flagged as decaying, and findings that persist across multiple measurement periods without state change are classified as candidates for Normalized Deviance (Section 6). Closure is the hardest property to maintain, because it requires that results actually change something. The decay tracking mechanism is the architectural response to the observation that “the Cage is patient—it waits for the rigor to tire” [McEntire, 2026].

## 9 Implementation

The reference implementation consists of a Python 3.12+ codebase with 1,024 tests covering mesh operations, routing, familiarity scoring, constraint enforcement, energy dynamics, lifecycle management, consensus, insight generation, spectral analysis, policy enforcement, identity resolution, linguistic compression metrics, and attention modeling. Data validation uses Pydantic; numerical operations use NumPy. All I/O is asynchronous. The design philosophy privileges determinism where possible and emergence where necessary: routing is deterministic given the current mesh state, while the resulting topology emerges from the interaction between signal content and routing dynamics.

The four-agent tiering (indexer, surfer, correlator, profiler) uses LLM-backed reasoning with mechanical fallbacks. If an LLM call fails, times out, or produces output that fails structural validation, the system falls back to deterministic heuristics that produce lower-quality but structurally valid results. The insight agents never block the routing pipeline. Signals continue flowing and workers continue differentiating regardless of whether higher-order analysis succeeds. This separation ensures that the mesh’s perceptual function—its ability to see organizational structure—does not depend on the availability or reliability of any single AI model.

The communication graph—the entity-relationship structure extracted from signal metadata—provides the substrate for bridge distance computation, channel migration detection, and phase-lock analysis. Entity resolution maps heterogeneous identifiers in sematectonic traces (commit author names, PR reviewer handles, deployment service names, CI/CD pipeline identifiers) to a unified organizational ontology through a multi-provider identity resolution system. Static authoritative sources (team rosters, email aliases) establish canonical identities; live providers (GitHub organization membership, Linear workspace users, Slack workspace users) enrich profiles with platform-specific handles and UUIDs; and a runtime alias store persists identities discovered during signal processing across sessions. Resolution proceeds through progressive normalization—exact match, handle-prefix stripping, CamelCase decomposition, username pattern matching, and fuzzy matching—with explicit system-identifier detection to exclude bots and automation from human identity graphs. A PII sensitivity layer classifies fields by exposure level (public, internal, confidential) and redacts above the configured threshold at output boundaries, ensuring that the mesh’s internal identity model never leaks into logs or external integrations. The architecture achieved 95% resolution rate (70 canonical identities, 4 unresolved—all automation accounts) in continuous deployment, progressing from 54% with code-repository-only metadata to 83% with communication channel data to 95% with multi-provider enrichment. The communication graph and the mesh topology provide complementary views: the mesh

topology reflects structural similarity of signal content, while the communication graph reflects the organizational relationships between the entities producing the signals. Divergence between the two—when structurally similar signals come from organizationally distant teams, or when organizationally proximate teams produce structurally dissimilar signals—is itself diagnostic.

The Normalized Deviance detection framework (Section 6) is operationalized through five automated detectors. An *action-correlation* detector classifies signals as discussion (messages, comments) or action (commits, deployments, ticket closures, state transitions) and computes the ratio per topic cluster; ratios below 0.1 indicate decoupling—the approach to babbling where the communication channel retains resolution for describing problems but not for resolving them. The detector additionally tracks March’s 1991 exploration-exploitation balance: repeated actions on the same systems constitute exploitation, while actions on novel systems constitute exploration, and high action-correlation composed entirely of exploitation signals the vulnerability that March predicts. A *channel-migration* detector measures the ratio of private to public channel discussion per topic, weighted by concurrent work activity; high private ratio with continued work activity indicates knowledge occlusion. A *phase-lock* detector clusters findings by Jaccard similarity of term sets and flags clusters where average similarity exceeds 0.6 with action-correlation below 0.2—converging risk framing without corresponding action, the synchronized oscillator dynamic that produces rogue waves. A *solo-ownership* detector identifies topics spanning multiple systems with exactly one actor producing action signals; multi-system failures assigned to individual engineers without escalation paths represent structural isolation at capacity ceilings. A *linguistic-compression* detector implements the leading indicator described in Section 6.4.

The open taxonomy design uses string identifiers for signal sources, action types, worker competencies, and insight categories. No enumerated types constrain what the system can discover. An action type like `hotfix_deployed_without_review` can emerge from the traces themselves without anyone having predicted that category would be needed. Alexander’s pattern language framework provides the theoretical justification [Alexander, 1977]: design quality emerges from the interaction of semi-independent patterns (competency-to-primitive mappings), not from top-down specification of all possible categories.

Cost structure was a first-order design constraint. The mesh’s self-organizing topology operates on term-frequency comparisons and threshold arithmetic, not on LLM inference. Only the insight agents require LLM calls, and those calls operate on pre-structured observation contexts rather than raw signal streams. The system can process thousands of signals per hour at commodity compute costs, with LLM inference costs scaling with the frequency of structural assessments rather than with signal volume.

## 10 Evaluation

The architecture is deployed with organizational consent at a vacation rental marketplace and property management platform with 29–31 engineers across multiple timezones. No individual-identifying information is reported; all findings are described at the team and system level. The deployment ingests sematectonic traces from GitHub (commits, pull requests, reviews, CI/CD pipeline outputs) and Linear (issues, comments, status changes) across 91 repositories.

In the full 30-day deployment, the mesh ingested 3,520 signals (709 from GitHub, 2,811 from Linear) and processed the complete signal stream. The mesh self-organized from 3 initial workers to 47, with emergent specializations—PMS calendar synchronization, booking and reservation management, unit and listing data operations, payment processing, bug fixes and testing—corresponding to the organization’s functional domains without being given that structure. Acceptance rate was 59% (2,081 accepted, 1,439 duplicates). Batch deduplication in the correlator eliminated 89% of redundant LLM calls (8,840 calls avoided). Total cost was \$42.49; processing cost was \$0.012 per signal. Three policy interventions enforced budget constraints

on repositories exceeding differential line thresholds.

Twenty-three findings were surfaced across five categories (RISK, CONTRADICTION, TREND, OVERLAP, DEPENDENCY) with confidence scores ranging continuously from 40% to 95%. The raw correlator produced 3,146 candidate observations from overlapping analysis windows; the surfer collapsed these to 23 deduplicated key findings—an 11% survival rate from candidate to promoted finding, evidence of genuine selectivity rather than undifferentiated surfacing. The gradation discriminates meaningfully: 40–65% for structural overlap within a single ticket scope, 70–80% for cross-ticket correlations with partial evidence, 85% for multi-source convergence with direct textual evidence of conflict, and 90–95% for patterns with quantitative evidence, named organizational impact, and multiple corroborating signal sources.

Four findings illustrate distinct theoretical claims. A systemic calendar synchronization failure (85% TREND) connected six tickets across different properties and three external system vendors—all involving calendar state divergence between the internal management system and external platforms, worked independently with no visible coordination. The mesh detected the systemic pattern from sematectonic traces alone; the structural conclusion—that one synchronization protocol is architecturally insufficient for a class of properties—was not visible in any single ticket. A feature-level decision vacuum (90% RISK, 85% CONTRADICTION) identified a multi-party fragmented negotiation around a financially impactful feature where a committed code solution (a per-unit configuration toggle resolving all competing positions) existed in the dependency graph  $G_D$  while the organizational impasse persisted in the communication graph  $G_C$ —the gap between what the organization had built and what it knew it had built. This is the  $G_D/G_C$  divergence formalized in Section 5 operating at the feature-decision level. A policy communication failure (85% CONTRADICTION) surfaced a case where one team was unaware of a policy change that other teams had been executing for months, with financial transactions already processed under the assumed new policy before the contradiction was detected. A P1 ticket unassigned for 26 days was being actively cited in conversations to explain related failures—quantified at over \$1,000 per incident with multiple known examples—but never discussed as a work item. The ticket’s existence substituted for resolution: action decoupling in which the organization’s tracking system itself became a normalization mechanism.

One result is notable for what it did *not* produce. The coupling detector, which in the initial run generated 100 false positives at activation 1.00, produced zero false couplings after centered cosine correlation with variance guards—the correct result when no hidden couplings exist in the processed sample.

The mesh’s findings were subjected to independent cross-source validation. An AI agent with access to the organization’s internal communication records—a data source the mesh does not ingest—was given the mesh’s specific findings and asked to determine whether communication records would confirm, complicate, or refute them. All five investigated findings were confirmed and three were extended. The calendar synchronization finding was confirmed with a critical temporal comparison: the engineer who independently connected the same tickets did so months after the mesh had correlated them from sematectonic traces, and the engineer’s architectural conclusion remained in a conversation thread rather than propagating to the issue tracker as a decision. The decision vacuum was worse than detected: four contradictory positions rather than three, spread across two communication channels with no cross-reference, and the existing code solution undiscovered by any of the four parties. The policy failure was confirmed with a complete timeline: the policy decision was made in September, a participant explicitly warned 77 days later that the executing team had not been informed, and the predicted failure materialized 89 days after the warning—the executing team confirmed they were unaware of the change. Even the participant who knew about the change had imprecise recall of when it occurred.

A consistent structural pattern emerged across the validation: 5 of 6 tickets investigated had richer resolution context in communication records than in the issue tracker, and 0 of 6 had that context reflected back. The mesh, reading costly signals (code commits, ticket state changes),

detected patterns that existed in the organization’s cheap-signal channels (conversations) but could not propagate through them. The information existed in the organization. It was not legible to the organization’s formal tracking systems. The mesh made it legible from a structurally independent observation channel.

The cross-source validation provides preliminary detection latency evidence: in the calendar synchronization case, the mesh detected the systemic pattern from a 30-day signal window in a single automated pass, while the human engineer who independently detected the same pattern did so after months of accumulated experience with individual incidents. The mesh’s detection was both earlier and persistent—a structured, scored artifact with signal provenance rather than a conversation that the organization could not act on. Action-correlation—whether surfaced findings produce organizational state changes—requires longitudinal follow-up observation beyond the initial deployment window.

A subsequent 7-day run over a subset of the same organizational period tested cross-run stability. The mesh self-organized from 3 initial workers to 7—fewer than the 30-day run’s 47, but at 49% average fullness versus 22%, reflecting tighter specialization under lower signal volume. The organizational communication graph was monotonically accumulative: the 7-day run’s graph was a strict superset of the 30-day run’s (34 nodes, 58 edges versus 31 nodes, 53 edges), confirming that the mesh does not forget structure across runs. The mesh’s vocabulary doubled from 2,935 terms to 6,126 despite processing only 39% as many signals—empirical evidence that ART’s match-based learning produces non-linear representational improvement when prior categories exist to anchor new signals. The insight funnel for the 7-day run showed 125 raw candidate observations collapsing to 95 after agent-level deduplication and 14 promoted findings—the same 11% survival rate as the full run, suggesting the selectivity threshold is a stable property of the architecture rather than an artifact of signal volume.

A three-source deployment added organizational communication channels alongside code repositories and the issue tracker. The cadence monitor (Section 4.8) learned distinct weekday and weekend entropy profiles within the first observation cycle, correctly distinguishing between symmetric quiet—reduced activity across all sources on weekends—and asymmetric quiet—a single engineer’s weekend code repository activity producing a characteristically different but organizationally normal entropy signature on Saturdays. No configuration of expected source distributions or activity schedules was required; the mesh learned the organization’s operational rhythms from observation alone. The addition of a third signal source also demonstrated the identity resolution improvements that enable cross-source corroboration: entity resolution improved from 54% to 83% with communication channel data providing the display-name-to-handle mappings that repository and tracker metadata lack, and subsequently to 95% (70 canonical identities resolved, 4 unresolved—all automation accounts) with multi-provider enrichment from GitHub organization membership, Linear workspace users, and Slack workspace users operating alongside static authoritative sources. The system detected 48 cross-source corroborations linking issue tracker activity with communication channel discussions via resolved actor identities.

Continuous operation over a 15-day live window processed 3,138 signals from 64 GitHub repositories and 7 Linear teams. The mesh self-organized from 3 initial workers to 50 and pruned 6 idle workers to a stable topology of 44 at 22% average fullness—the same equilibrium density as the initial 30-day run, suggesting that fullness stabilizes as a property of the organization rather than of signal volume. Batch deduplication in the correlator eliminated 87% of redundant LLM calls (9,007 calls avoided); total cost was \$51.58 (\$0.016 per signal). The policy engine recorded 2,968 traces and enforced 4 budget interventions on repositories exceeding differential line thresholds. The insight funnel produced 339 raw candidate observations, collapsing to 185 after agent-level deduplication and 36 promoted findings—an 11% survival rate, replicating the selectivity observed in both prior runs and confirming it as a stable architectural property rather than an artifact of deployment duration. Critically, bridge distance scores in this deployment ranged from 0.37 to 0.98, resolving the initial deployment’s null result: the

three-source identity resolution produced sufficient graph density for the  $d_{\text{bridge}}$  component to discriminate meaningfully. The highest-scoring finding ( $S(f) = 0.498$ ,  $d_{\text{bridge}} = 0.98$ ) identified three engineers independently investigating the same authentication failure from different hypotheses—structurally disconnected actors converging on a single root cause without visible coordination, with one engineer’s fix potentially rendering the others’ investigation redundant.

Supplementary analysis of the organization’s internal communication records tested the three Normalized Deviance signatures (Section 6) against five recurring problem areas. The analysis was conducted independently: the communication records (several months of conversational history across twelve channels) constitute a data source the mesh did not ingest, and the analysis was not informed of the mesh’s specific findings. It was asked to search for ND signatures in the communication data and converged on the same organizational topics the mesh had independently identified from engineering traces. This cross-source convergence—two instruments observing different data, arriving at the same structural patterns—is the cross-source consistency checking described in Section 11.2 operating in practice.

*Action decoupling* was the dominant signature, present at high confidence across all five topics examined. In every case, discussion volume was high—threads of 50 to 166 replies on individual incidents—while resolution evidence was absent or slow. The organization is fluent at describing its problems. The description has become the deliverable. One topic showed a dedicated communication channel for a class of failures that should be rare enough not to need one; the channel’s existence is an organizational artifact that encodes the assumption the failures will continue, making management of ongoing failure more efficient and thereby reducing the pressure to eliminate the failure itself. Across all five topics, the ratio of discussion messages to action messages (where “action” denotes a reference to a concrete state change: a commit, a deployment, a ticket closure, a decision with an owner and a deadline) was consistent with the pattern formalized in Section 6.1: the communication channel’s partition count  $N^*$  remains high for describing the problem but approaches babbling for resolving it.

*Linguistic compression* was present at moderate confidence in four of five topics, limited by temporal depth in the conversational record. The most striking example involved a data type confusion producing six-figure financial errors, described as “should be an easy fix.” The containment framing—locating the problem in a known subsystem, minimizing remediation effort, hedging commitment—is precisely the linguistic signature Section 6.4 predicts: specificity of commitments decreasing, nominalization converting active processes into static nouns, hedging density increasing. The analysis noted that the thinness of the conversational record for an error of that magnitude may itself be the most concerning signal: problems large enough to warrant alarm that generate little traceable alarm have typically already been deeply normalized.

*Framing convergence* was the weakest signature, present at low-to-moderate confidence. Most topics were dominated by a single voice rather than multiple independent perspectives, limiting the measurement. Where multiple speakers were present, the convergence was visible—two engineers describing a systemic availability problem in identical terms, acknowledging scope (“this is happening for almost all”) while not escalating. In a non-normalized culture, that sentence initiates an incident declaration. Here it reads as status.

Three cross-cutting findings extend the theoretical framework. First, *taxonomic competence functions as a normalization mechanism*. Teams have developed fluent vocabularies for recurring failures—standardized category names, established triage flows, practiced diagnostic routines. The competence at categorizing failures prevents those failures from being treated as the crises they represent. Naming the problem has become a substitute for solving it. This mechanism bridges linguistic compression and framing convergence: the shared vocabulary *is* the convergence. Second, *solo ownership of systemic problems* appeared in two of five topics—multi-system failures assigned to individual engineers without reinforcement, where confusion or capacity limits produce stalls with no escalation path. Third, the “*easy fix*” frame appeared across topics: describing a months-old, high-impact failure as simple suppresses urgency by implying the prob-

lem persists by choice rather than by difficulty, which is precisely the reframing that Vaughan’s normalization of deviance predicts [Vaughan, 1996].

The evaluation framework spans three quadrant-specific measurement domains.

*Discovery detection.* The primary measure is instances of second-order ignorance surfaced: structural couplings, knowledge silos, and coordination fractures that were not previously visible to any team member. Secondary measures include the ratio of actionable insights to total insights surfaced (precision), the time between a structural anomaly emerging in the trace data and the mesh surfacing a perceptible signal (detection latency), and the correspondence between the mesh’s emergent topology and the organization’s actual work structure, validated by team leads who can confirm or deny the structural relationships the mesh discovers. A null result—zero discovery-classified findings in a deployment window—does not by itself indicate a system limitation; it may indicate that the organization did not produce discoverable second-order ignorance in that window, particularly in small, communicative teams where knowledge boundaries overlap significantly. The distinction between “the system cannot discover” and “there was nothing to discover here” is empirically testable through cross-organizational deployment.

*Normalized Deviance detection.* Five automated detectors (Section 9) operationalize the three longitudinal signatures. Eigenvalue trend tracks the effective dimensionality of the mesh’s correlation matrix for specific risk domains. Linguistic compression trend monitors eight agent-assessed metrics (Section 6.4) with exponential moving average smoothing and Born Caged baseline detection, with display thresholds calibrated to empirical baselines from SEC filing analysis (LD compression onset at 0.16, severe compression below 0.13). Action-correlation computes the ratio of discussion to state change per topic cluster, with the exploration-exploitation balance tracked separately to identify March’s 1991 vulnerability: topics showing high action-correlation composed entirely of exploitation (repeated actions on the same systems, no novel system engagement). Channel migration and solo ownership provide supplementary structural signatures. A finding that scores high on  $r_{\text{coupling}}$  and low on  $d_{\text{bridge}}$  while exhibiting decreasing eigenvalue trend and increasing linguistic compression is a candidate for Normalized Deviance.

*Decay tracking.* The finding registry records finding-to-acknowledgment latency (how long after surfacing does the organization respond?), acknowledgment-to-action latency (how long after response does a state change occur?), and decay events (findings that were acknowledged but produced no state change within a configurable window). The migration rate from Discovery to Normalized Deviance—findings that were initially high- $d_{\text{bridge}}$  and transitioned to low- $d_{\text{bridge}}$  without a corresponding decrease in  $r_{\text{coupling}}$ —provides a quantitative measure of the decay path.

The value of individual findings is evaluated as Bayesian surprise: the Kullback–Leibler divergence  $D_{\text{KL}}(P_{\text{post}}||P_{\text{prior}})$  between the organization’s dependency model before and after incorporating the finding. A finding that changes nothing in the model—either because it was already known or because it is irrelevant—has zero KL divergence. A finding that forces a substantial revision of the dependency model has high divergence. This provides a theoretically grounded measure of epistemic value that is independent of subjective novelty or satisfaction ratings: it measures how much the organization’s model of itself had to change when the finding was incorporated. Counterfactual evaluation via Inverse Propensity Scoring will estimate the aggregate value of surfaced insights against a baseline of no-mesh operation.

The mesh’s findings self-classify into four tiers of epistemic value. *Retrieval* re-describes knowledge already present in the artifact stream—restating the content of existing tickets or pull requests. *Correlation* connects signals across source or actor boundaries that had not been previously linked—identifying that two engineers are independently working on related subsystems, or that an issue tracker pattern corresponds to a communication channel discussion. *Amplification* quantifies a known problem with new evidence—attaching magnitude, frequency, or structural context to an issue the organization was aware of but had not measured. *Discovery* surfaces second-order ignorance: structural patterns that no individual actor could have identified because the constituent signals resided in separate knowledge boundaries. In the deployment

reported here, the observed distribution was 57% retrieval, 38% correlation, 5% amplification, and 0% discovery. The mesh’s primary empirical contribution was coordination failure detection and normalized deviance identification. Discovery-classified findings were not observed, consistent with a small, communicative organization operating below Elliott’s coordination threshold where knowledge boundaries overlap significantly and most structural relationships are already visible to at least one team member. The system correctly self-classified retrieval findings as such—an honest result that distinguishes between the mesh functioning as a search engine (retrieval) and functioning as a structural detector (correlation and above).

## 11 Discussion

### 11.1 Computational Formalization of External Observation

The architecture’s theoretical contribution is the computational formalization of external observation for organizations. Luhmann’s autopoiesis theory establishes that organizations, as operationally closed systems, cannot observe their own blind spots from within [Luhmann, 1995]. Lawvere’s fixed-point theorem formalizes the constraint: no system admits a complete self-evaluation function [Lawvere, 1969]. The mesh provides what neither the organization’s internal processes nor its internal members can provide: a structural model of the organization that is not subject to the organization’s own operational closure.

This claim requires precision. The mesh is not omniscient. It has its own blind spots: signals it cannot ingest, patterns below its spectral resolution, structural relationships that do not manifest in sematectonic traces. The claim is narrower and more defensible: the mesh’s blind spots are structurally independent of the organization’s blind spots, because the mesh observes traces rather than participating in the communication processes that create organizational blindness. A manager who does not report a problem creates a gap in the organization’s self-model. The manager’s code commits, deployment patterns, and build failures still produce sematectonic traces that the mesh can observe. The information suppressed in the communication channel persists in the work channel, and the mesh reads the work channel.

### 11.2 Immunity to Dymemic Pressure

The mesh is immune to dymemic pressure because it does not participate in the communication processes that create it. It reads sematectonic traces, which are costly signals that cannot be degraded by the sender’s incentive structure without also degrading the work product itself. An engineer who inflates a status report faces no cost. An engineer who inflates a test suite—adding hundreds of passing tests that test nothing meaningful—faces the cost of writing and maintaining those tests. The differential cost of manipulation between cheap signals (reports) and costly signals (artifacts) is the mesh’s primary defense against strategic distortion.

This defense is not absolute. The verification games literature identifies scenarios where even costly signals can be strategically shaped: commits restructured to hide coupling, activity inflated to dominate a worker’s term space, logging selectively configured to create false spectral signatures. The mesh’s resilience to these attacks depends on cross-source consistency checking—the principle that manipulating one trace source while leaving correlated sources unmanipulated creates spectral discordances that the mesh can detect. An engineer who restructures commits to hide a dependency cannot simultaneously restructure the deployment manifest, the CI/CD pipeline configuration, and the database migration files without the effort exceeding the benefit. The attack surface is proportional to the number of independent trace sources the mesh ingests, which is why source diversity is an architectural priority.

### 11.3 Dimensional Independence Rather Than Exogeneity

The mesh does not claim to be an external observer in any absolute sense. No such observer exists for an embedded system. Lawvere’s fixed-point theorem prohibits complete self-evaluation, but the prohibition targets universality, not partial evaluation. The companion work on orthogonal evaluation establishes that exogeneity is not binary—it exists on a spectrum defined by the degree of informational, structural, and logical independence between the evaluation criterion and the process being evaluated [McEntire, 2026b]. Mathematical orthogonality (error-correcting codes operating in the null space of the generator matrix) is nearly as reliable as true exogeneity. Socially constructed orthogonality (financial auditing) predictably collapses. The mesh operates in the middle of this spectrum: its evaluation criteria are structurally independent of the organization’s own criteria (it reads traces, not reports; it measures topology, not sentiment) but not logically independent (it is configured and deployed by humans within the organizational context).

The claim is dimensional, not positional. The organization’s endogenous evaluation criteria span a subspace  $K$ . The mesh generates evaluation vectors outside  $K$ —not because it occupies a magical external vantage point, but because topological bridging, spectral analysis, and graph divergence produce measurements that the organization’s own instruments cannot produce, since those instruments are constructed from the same knowledge that defines  $K$ . Each new vector dimension provides partial escape from the self-referential trap. The escape is not complete. It does not need to be. Surfacing one structural coupling that no team knew existed, detecting one instance of normalized deviance where discussion substituted for action, is more than any existing tool achieves. The contribution is not solving all second-order ignorance. It is solving some, in a domain where the prior state of the art solves none.

The Conant-Ashby Good Regulator Theorem requires that the regulator’s internal model be isomorphic to the system being regulated [Conant & Ashby, 1970]. The mesh’s emergent topology satisfies this requirement in a way that no designed model can: it grows and changes in response to the same signal stream that reflects the organization’s actual work, producing a representation whose structure is determined by the organization’s structure rather than by a designer’s assumptions about what that structure might be. Ashby’s Law of Requisite Variety—“only variety can destroy variety”—is satisfied by the mesh’s open taxonomy and dynamic topology, which can represent arbitrary complexity rather than being constrained to a fixed schema [Ashby, 1956].

### 11.4 Bounded Rationality as Architectural Principle

The mesh’s competitive routing implements bounded rationality as a design choice rather than a limitation. Simon demonstrated that satisficing outperforms optimization under the conditions that characterize real organizational environments: high evaluation costs, time pressure, and non-stationary distributions [Simon, 1956]. The mesh’s first-to-accept routing is optimal in Gigerenzer’s sense of ecological rationality: a simple heuristic that exploits the structure of its environment (related signals share vocabulary, so familiar workers are likely to be appropriate workers) to achieve performance that matches or exceeds exhaustive assignment.

Ocasio’s Attention-Based View provides the organizational counterpart [Ocasio, 1997]. Attention is the binding constraint on organizational decision-making. The mesh’s spectral detection and consensus mechanisms function as an attentional prosthetic: they do not replace human attention but redirect it. By surfacing structural anomalies that would otherwise remain below the threshold of organizational perception, the mesh alters what Ocasio calls “situated attention”—the specific issues that decision-makers attend to in specific contexts. The mesh makes visible what was previously invisible, and visibility is the precondition for action.

## 11.5 Probabilistic Emission

The mesh’s value is exclusively inferential. Every signal it ingests was produced by someone who already knows it. Every ticket, pull request, and message was authored by a person with direct access to that fact. If the mesh restates any first-order fact, it is functioning as a search engine. The only output worth surfacing is the implication that no individual actor could derive alone—the structural pattern that emerges from combining facts distributed across separate knowledge boundaries. This means the emission decision must distinguish between what an actor has *access to* and what an actor *actually knows*. Access is binary: a person was tagged on a ticket or was not. Attention is a probability distribution shaped by audience size, signal volume, temporal decay, contextual load, and role relevance. The same fact presented in a 6-person design review and a 40-person all-hands has different probabilities of registering with any given attendee. An actor tagged on every ticket in a project effectively receives no signal from any individual tag—the information-theoretic capacity of the attention channel is finite, and per-signal absorption decreases as volume increases.

The surfacing decision is an expected-loss calculation. Let  $P_k(a, f)$  denote the probability that actor  $a$  already knows finding  $f$ , estimated from the attention model, and let  $I(f) = S(f)$  denote the finding’s importance from Equation 4. Let  $C_r$  denote the cost of redundancy—annoyance, attention fatigue, erosion of trust in the system. The emission criterion surfaces finding  $f$  to actor  $a$  when the expected cost of not surfacing exceeds the expected cost of surfacing:

$$I(f) \cdot (1 - P_k(a, f)) > C_r \cdot P_k(a, f) \quad (6)$$

Rearranging:  $I(f)/C_r > P_k(a, f)/(1 - P_k(a, f))$ . The left side is the importance-to-annoyance ratio; the right side is the odds that the actor already knows. High importance with low  $P_k$  always surfaces—the cost of the actor not knowing dominates. Low importance with high  $P_k$  never surfaces—the cost of redundancy dominates. Critically, high importance with *moderate*  $P_k$  still surfaces: a finding discussed three weeks ago in a large meeting, which may or may not have registered, warrants a redundant reminder when the cost of being wrong about  $P_k$  exceeds the cost of annoyance. Findings are ranked by the margin  $I(f)(1 - P_k) - C_r \cdot P_k$ , so actors encounter the highest expected-value items first. The probability  $P_k$  is estimated via noisy-OR over exposure events—each prior encounter (authoring, reviewing, being tagged, being present in a channel) independently contributes a probability of absorption, with temporal decay, channel-size discount, and volume discount as modifiers. The model calibrates through feedback: “I already knew that” revises the actor’s attention estimate upward; “I had no idea” revises it downward—with asymmetric learning rates, because the cost of under-informing is invisible while the cost of over-informing is self-correcting. Over time, the mesh learns each actor’s actual attention footprint rather than their nominal access profile.

The emission model addresses a self-referential threat. A mesh that surfaces too many redundant findings trains the organization to normalize ignoring it. The mesh’s communication channel degrades through the same mechanism the mesh was designed to detect in the organization: dysmemic pressure selects against attention to the signal because the signal-to-noise ratio has fallen below the threshold that justifies the cognitive cost of engagement. The emission criterion is the mesh’s defense against becoming its own Normalized Deviance—the system everyone has learned to disregard. The formal connection to Crawford–Sobel is direct: even a sender with zero incentive divergence produces zero information transfer if the receiver’s prior already contains the message. The attention model estimates the receiver’s prior, ensuring that the mesh’s limited bandwidth is allocated to findings with the highest expected information gain. A structural consequence: when a high-importance finding has high  $P_k$  for the direct team but low  $P_k$  for leadership, the emission criterion routes the finding upward without an escalation process—the mathematics dissolve the political cost of escalation by surfacing findings to the actors with the lowest  $P_k$  weighted by authority to act, routing around the hierarchical com-

pression that Liberti and Mian show collapses sensitivity at Level 3. The reflexivity extends further: the mesh must apply the same credibility calculus to its own users that it applies to organizational signals, weighting observed behavior (decisions that contradict a finding) above stated feedback (“I already knew that”), because the user’s self-report is itself a cheap signal subject to the bounded rationality the emission model exists to address.

Organizational dysfunction manifests as clusters of individually sub-threshold findings rather than single dramatic failures. The emission model addresses this through compound risk scoring: findings that share actors, terms, or channels are grouped by single-linkage clustering, and the cluster score  $S_c = \max(S_i) + \delta \sum_{j \neq i^*} S_j$  (where  $\delta = 0.3$ ) can trigger escalation when no individual finding exceeds the surfacing threshold. Four findings each scoring 0.3 produce  $S_c = 0.3 + 0.9 \times 0.3 = 0.57$ —surfaced as a portfolio when the individual findings would be suppressed. This formalizes Flyvbjerg’s observation that complex-system failures accumulate through normalized deviation at the individual level while the compound risk grows unmonitored.

The emission model also addresses temporal dynamics through re-emission—event-driven re-triggering of findings that have decayed below the surfacing threshold. Three independent triggers operate: *new evidence* re-surfaces a deferred finding when fresh signals overlap its term set, preventing premature closure; *compound risk* re-surfaces when a pair of individually sub-threshold findings (sharing actors or terms) jointly cross the threshold, connecting the portfolio mechanism to the temporal lifecycle; and *decay escalation* provides a last-chance alarm when a finding enters the normalized stage without having produced an organizational state change. The decay escalation trigger operationalizes Vaughan’s normalization of deviance at the system level: a finding that has been surfaced, acknowledged, and not acted upon is migrating from Discovery into the organization’s ambient background, and the re-emission mechanism forces it back into visibility before the transition completes.

The redundancy cost  $C_r$  is not static. An adaptive mechanism scales  $C_r = C_{\text{base}} \times (0.5 + r_{\text{fp}})$  where  $r_{\text{fp}}$  is the ratio of “already knew” feedback to total feedback for that actor. An actor who consistently reports knowing surfaced findings sees  $C_r$  increase (up to  $1.5 \times C_{\text{base}}$ ), suppressing marginal findings. An actor who consistently reports surprise sees  $C_r$  decrease (down to  $0.5 \times C_{\text{base}}$ ), lowering the surfacing threshold. The adaptation prevents the mesh from converging on a single surfacing policy across actors with heterogeneous attention footprints. Complementing the adaptive threshold, the system discloses its filtering decisions: each surfacing pass reports how many findings were suppressed and why—“ $N$  suppressed as likely known,  $M$  below significance threshold”—with the four scoring components ( $d_{\text{bridge}}$ ,  $c_{\text{entity}}$ ,  $r_{\text{coupling}}$ ,  $\gamma$ ) displayed alongside each finding’s composite score. The disclosure prevents the mesh from becoming a black box whose suppressions are themselves a form of organizational blindness: if the system filters silently, it can induce second-order ignorance about the user’s own information diet.

## 11.6 Substrate Independence as Contribution

If the compression-selection-equilibrium mechanism is substrate-independent, the mesh design may transfer to domains beyond organizational knowledge management. AI alignment faces the same structural obstacle: language models compress training corpora into parameter weights, preference signals select on the compressed representation, and the resulting outputs drift toward proxy fitness rather than genuine alignment. The mesh architecture—observing costly traces rather than cheap reports, using competitive routing rather than centralized coordination, detecting spectral anomalies rather than semantic categories—addresses the structural mechanism rather than the substrate-specific symptoms.

Academic peer review exhibits the same dynamics. Publication compresses research into standardized formats; citation counts and journal prestige select on the compressed representation; work that fits the paradigm survives while paradigm-challenging work faces friction [Smaldino & McElreath, 2016]. A mesh that ingested the sematectonic traces of research

activity—data collection patterns, replication attempts, methodological choices visible in code repositories—rather than the published representations of that activity might detect the same structural anomalies in academic knowledge production that the organizational mesh detects in software engineering.

The substrate independence claim is testable. If correct, interventions proven effective in one domain should transfer when structurally analogous versions are applied in others. If organizational Mirrors improve accuracy, analogous structures in AI systems (evaluation outside the training loop, separated from deployment optimization) should improve alignment. The tests are specified. The predictions are falsifiable.

The self-observing coherence mechanism (Section 4.8) provides concrete evidence for the substrate independence claim. The mesh’s anomaly detection generalizes to self-monitoring because the mesh’s input channel is structurally identical to the organizational communication channels it observes—both are signal environments subject to compression, selection, and decay. The detection mechanism is indifferent to whether the signal source is messages between engineers or events arriving at the mesh’s own ingestion layer. The architecture does not distinguish between “organizational signal that has gone silent” and “mesh input source that has gone silent” because the distinction is substrate-specific rather than structural. A reflexive note is warranted: the self-observation capability was identified not by the system itself but by an external observer who recognized that input coherence checking was reimplementing detection logic the architecture already possessed. This is itself an instance of second-order ignorance resolved by external observation—the system did not know what it did not know about its own capabilities. The mesh needed a Mirror.

## 11.7 Oscillating Blind Spots

Spencer-Brown’s re-entering form establishes that self-reference produces temporal oscillation rather than static contradiction [Spencer-Brown, 1969]. An organization that discovers a blind spot has altered its evaluation surface. The blind spot does not disappear; it moves. The evaluation criteria that surfaced the original blind spot now constitute part of the organization’s frame, and the new frame has its own null space.

The implication for the mesh is that no single finding is permanent. A structural coupling surfaced and resolved changes the organizational topology. The mesh’s next measurement period confronts a different organization, whose blind spots are determined by the new frame that includes the resolution of the previous finding. The mesh must track movement, not position. The finding registry (Section 7) provides the longitudinal data: is the same structural gap appearing in different locations over successive measurement periods? If so, the organization is oscillating rather than converging, and the oscillation pattern is itself a diagnostic.

## 11.8 Limitations

The mesh cannot close the loop autonomously. It surfaces findings and deposits them in the work environment. It cannot ensure that anyone acts on them. The decay path from Discovery to Normalized Deviance (Section 7) is the system’s primary failure mode, and it is structural: the mesh can make findings visible, but it cannot alter the payoff structure that determines whether visibility leads to action.

The mesh does not address incentive structures directly. If the organization’s reward system punishes the engineer who acts on a surfaced finding (because the fix disrupts a deadline, because the finding implicates a political ally, because the remediation is expensive), the finding will decay regardless of its environmental presence. Stigmergic deposition raises the cost of ignoring the finding but does not eliminate it.

The mesh has its own blind spots. It can observe only what produces sematectonic traces. Work that leaves no trace—verbal agreements, unrecorded decisions, knowledge held in the

minds of individual practitioners—is invisible. The mesh’s blind spots are structurally independent of the organization’s blind spots (they are determined by the trace sources available rather than by the organization’s frame), but they exist. The claim is partial escape from self-referential evaluation failure, not complete escape. Lawvere’s theorem is not negotiable.

Bridge distance computation required sufficient graph density before producing discriminative values. In the initial deployment window, the communication graph was too sparse; subsequent three-source deployment with identity resolution populating cross-source edges produced bridge distances ranging from 0.37 (actors in overlapping team clusters) to 0.98 (actors in structurally disconnected departments investigating the same failure independently). The metric discriminates meaningfully: a 95% correlation finding with  $d_{\text{bridge}} = 0.95$  correctly identifies two managers who independently created duplicate incident tickets within one minute—structurally disconnected actors performing identical triage without coordination—while a 30% retrieval finding with  $d_{\text{bridge}} = 0.50$  correctly identifies planned cross-repository work by a single developer. The consensus architecture is implemented but requires topological diversity that a single small organization may not produce; validation of quorum dynamics requires deployment across multiple organizations or a substantially larger engineering team.

The mesh’s persistence architecture separates semantic state (vocabulary, affinities, communication graph, attention model, finding registry) from topological state (worker count, specializations, selectivity thresholds, neighbor relationships). Semantic state is persisted across sessions; topological state is not. The mesh restarts from seed workers and must re-earn its differentiated topology through signal flow. This is a deliberate tradeoff: fresh workers at low vigilance thresholds accept signals readily and avoid the fragmented mesh that would result from restoring high-threshold workers into a shifted signal distribution. But it means that a mesh that self-organized to 44 specialized workers over 15 days of continuous operation restarts as 3 generic workers. ART guarantees that stable categories will re-form under continued input, but category formation is path-dependent—the presentation order of signals influences which categories emerge. Topology serialization with graduated re-acceptance (restoring workers at a fraction of their prior selectivity) would preserve the emergent structure while retaining the stability benefits of fresh-start acceptance.

The evaluation to date reflects deployment at a single organization. The architecture’s formal properties—ART stability, spectral detection, competitive routing—are organization-independent, but empirical validation of detection rates, false positive characteristics, and decay dynamics across diverse organizational structures remains future work.

## 11.9 Active Perturbation as Future Work

The current architecture is passive: it observes the traces that work produces and detects structural anomalies in the resulting topology. A natural extension is *active perturbation*—controlled informational interventions designed to discriminate load-bearing silence from ambient silence. The distinction matters: not every gap between the dependency graph and the communication graph represents a blind spot. Some dependencies are well-understood and simply do not generate discussion because they work. Active perturbation would introduce a small, controlled informational signal near a suspected blind spot and measure the organizational response. If the signal propagates normally—discussed, assessed, resolved or dismissed on its merits—the silence was ambient. If the signal is absorbed without trace, redirected, or generates a disproportionate defensive response, the silence was load-bearing, and dysmemic pressure is actively maintaining it. This connects the variance compression insight from the companion work on dysmemic pressure [McEntire, 2025a]: under constraint, the system cannot maintain all fictions simultaneously, and where variance persists or spikes under perturbation reveals real structure. The design of such perturbations—their magnitude, targeting, ethical constraints, and measurement protocols—is beyond the scope of this paper but represents a natural next step from passive observation to active diagnostic.

## 12 Related Work

The architecture synthesizes formal results from multiple disciplinary traditions, each of which addresses a component of the problem.

Stigmergy research, from Grassé’s original observations of termite construction [Grassé, 1959] through Heylighen’s theory of cognitive stigmergy [Heylighen, 2016], establishes that indirect coordination through environmental traces scales beyond the cognitive limits that constrain direct communication. Elliott’s empirical finding that stigmergic collaboration is the only viable coordination mechanism above approximately 25 participants provides the scaling argument [Elliott, 2007].

Self-organizing map research, from Kohonen’s original SOM [Kohonen, 1982] through Fritzke’s Growing Neural Gas [Fritzke, 1995] and Alahakoon’s Growing SOM, provides convergence proofs for competitive learning. Cottrell and Fort’s 1987 proof that one-dimensional competitive learning converges as a Markov chain extends to the mesh’s routing dynamics [Cottrell & Fort, 1987]. UbiSOM’s utility-based dead unit recycling maps to the mesh’s energy decay and archival mechanism.

Adaptive Resonance Theory, from Carpenter and Grossberg’s ART-1 [Carpenter & Grossberg, 1987] through Fuzzy ART [Carpenter et al., 1991] and ARTMAP, provides stability-plasticity guarantees that no other self-organizing architecture offers. The formal proof that match-based learning stabilizes in arbitrary non-stationary streams is the theoretical foundation for the mesh’s gap detection mechanism [Grossberg, 1976].

Signaling theory and strategic communication, from Spence’s costly signals [Spence, 1973] through Crawford and Sobel’s partition model [Crawford & Sobel, 1982] to Kamenica and Gentzkow’s Bayesian persuasion [Kamenica & Gentzkow, 2011], provides the economic formalization of why organizational communication degrades and why direct artifact observation is information-theoretically superior.

The Viable System Model, from Beer’s original cybernetic framework [Beer, 1972] through subsequent developments in variety engineering and organizational diagnosis, provides the completeness criterion for regulatory architectures. The five-system structure establishes what components are required for viability and how they must interact.

Rate-distortion theory and the Data Processing Inequality [Cover & Thomas, 2006] provide information-theoretic bounds on what any hierarchy-mediated observation system can achieve. Liberti and Mian’s empirical demonstration that hierarchies produce structural breaks in soft-information sensitivity transforms the theoretical argument into a quantitative one [Liberti & Mian, 2009].

Spectral graph theory, from the graph Laplacian and its eigenvalue properties through Graph-Wave [Donnat et al., 2018] and BWGNN [Tang et al., 2022], provides the mathematical framework for structural role detection and anomaly identification in the mesh topology.

The companion paper on dysmemic pressure [McEntire, 2025a] provides the formal adversary model: the compound selection force that the mesh is designed to circumvent. The three component dynamics (strategic communication degradation, adverse selection in idea markets, transmission bias) specify the mechanisms by which organizations produce self-deception. The mesh’s architectural choices—reading traces instead of reports, using competitive routing instead of hierarchical delegation, depositing findings stigmergically instead of alerting through channels—are point-for-point responses to those mechanisms.

Compressed sensing [Candès & Tao, 2006] establishes that if the signal of interest is sparse in some basis, far fewer measurements than classical theory demands are sufficient for recovery, provided the measurements are incoherent with that basis. If consequential blind spots are sparse at any given time—a reasonable assumption for most organizations—the mesh’s cross-source correlation functions as an approximation of incoherent measurement: each source provides a perspective orthogonal to any single team’s view, and the combination recovers structure that

no individual source could reveal.

NK landscape theory [Kauffman, 1993] formalizes the relationship between coupling and optimization difficulty. As the coupling parameter  $K$  increases, the fitness landscape becomes rugged and the number of local optima grows exponentially. The organizational parallel is direct: as hidden interdependencies accumulate, the organization enters a complexity catastrophe where local optimization (each team optimizing independently) produces global dysfunction. The mesh’s spectral detection of hidden couplings is an attempt to reduce the effective  $K$  by making the coupling visible before the catastrophe materializes.

Vaughan’s formalization of normalization of deviance [Vaughan, 1996] provides the canonical description of how organizations redefine unacceptable risk as routine. The mesh’s Normalized Deviance detection (Section 6) extends Vaughan’s qualitative account into a quantitative framework based on eigenvalue collapse, linguistic compression, and action-correlation metrics.

Network sociology, from Granovetter’s strength of weak ties [Granovetter, 1973] through Burt’s structural holes [Burt, 2004], provides the formal foundation for the mesh’s bridge distance computation. The mesh’s topological bridging operationalizes Burt’s insight that non-redundant information flows through positions of low network constraint.

Socio-technical congruence (STC) research directly formalizes the  $G_D/G_C$  comparison that the mesh’s bridge distance operationalizes. Cataldo, Herbsleb, and Carley demonstrated that misalignment between task dependency structures and coordination activities predicts integration failures and quantified the productivity impact [Cataldo et al., 2008]. Amrit and van Hillegersberg’s TESNA tool mines software repositories to compare architectural dependencies with developer social networks, identifying structure clashes automatically [Amrit & van Hillegersberg, 2010]. Tamburri, Palomba, and Kazman detect “community smells”—organizational anti-patterns such as lone wolves, radio silence, and organizational silos—from the same repository and communication data the mesh ingests [Tamburri et al., 2021]. The mesh differs from this tradition in three respects. First, STC tools operate on predefined pattern templates (specific smell definitions, specific congruence metrics); the mesh discovers emergent patterns not pre-specified by the analyst. Second, STC analysis is periodic—a snapshot of alignment at a point in time; the mesh operates continuously on streaming signals with findings that accumulate, decay, and migrate between states. Third, STC measures congruence as an aggregate property of the organization; the mesh surfaces specific findings with structural scores that discriminate significance, enabling prioritization of individual coordination gaps rather than aggregate assessment.

Prediction markets and Tetlock’s forecasting research [Tetlock, 2005] represent alternative approaches to surfacing second-order ignorance through aggregation of diverse perspectives. The mesh differs in that it does not require participants to formulate predictions; it observes traces and surfaces structure without requiring anyone to hypothesize what that structure might be.

What distinguishes this work from each of these traditions individually is the integration: a single architecture in which stigmergic data acquisition feeds self-organizing topology that is regulated by ART-style vigilance, monitored by spectral analysis, interpreted through bounded rationality, delivered through a cybernetically complete (to System 4) regulatory structure, and designed to resist the specific adversary model that dysmemic pressure formalizes. The contribution is the demonstration that these mechanisms compose into a coherent system that addresses a problem—second-order organizational ignorance—that none of them addresses alone.

## 13 Conclusion

Organizations are operationally closed systems. They reproduce their own elements from their own elements, and the conditions of their self-observation are constructed from the same processes they attempt to observe. This structural condition produces second-order ignorance: the hidden couplings, the emergent dependencies, the coordination fractures that no one knows to

look for because the organization’s own perceptual apparatus was not designed to see them. Every tool that requires a query inherits this limitation. The mesh does not require a query. It reads the traces that work leaves behind, organizes them into a topology that reflects the organization’s actual structure, and surfaces the geometric patterns that indicate where something unexpected is happening.

The architecture is not an assembly of convenient parts. It is a set of mechanisms each of which exists because a specific formal constraint demands it, and removing any one produces a specific, identifiable failure.

The mesh must ingest work artifacts rather than reports, because the Data Processing Inequality guarantees that reports carry strictly less information than the artifacts they summarize, and Liberti-Mian demonstrates that the loss is not gradual but catastrophic at the third hierarchical level. It must use competitive routing rather than centralized assignment, because Crawford-Sobel proves that communication channels with incentive divergence  $b \geq 1/4$  collapse to babbling, and Elliott’s threshold establishes that above 25 participants the accumulated bias across relays pushes past that bound. The routing must be match-based rather than gradient-based, because Grossberg’s ART stability proof is the only convergence guarantee that holds in non-stationary input streams—gradient descent produces catastrophic forgetting when the organization’s activity distribution shifts, which it does continuously. The mesh must use spectral detection rather than semantic classification, because second-order ignorance means the semantic categories that would identify the finding do not yet exist in the organization’s vocabulary; spectral signatures (right-shift for Discovery, left-shift for Normalized Deviance) are geometric properties of the topology that require no semantic pre-specification. The scoring function must be topological rather than sentiment-based, because Prendergast’s analysis of yes-men dynamics proves that any scoring mechanism based on self-report is gameable by the same dysmemic forces the system is designed to detect. The delivery must generate artifacts rather than annotations, because digital environments lack the geometry that makes physical stigmergy work—a bot comment on a pull request is structurally an alert arriving through a different channel, and Ocasio’s attention economics applies to it identically. The generated artifact must be a model artifact rather than a comment, because the Conant-Ashby Good Regulator Theorem requires that the organization’s regulatory model be isomorphic to the system being regulated—a comment describes the gap, but a parent ticket, an interface contract, or a dependency diagram *closes* it.

Strip any one of these mechanisms and the architecture fails at a specific point. Without costly signal ingestion, the mesh reads the same degraded information the organization already has. Without ART stability, the mesh’s categories collapse under distributional shift. Without spectral detection, the mesh can only find what it already has words for. Without topological scoring, the mesh is gameable. Without artifact generation, the mesh’s findings are annotations that decay through the same attention economics the architecture was designed to circumvent—comments about gaps rather than the organizational knowledge whose absence constituted them. The composition is not additive. It is structural. Each mechanism addresses a formal limitation that the others cannot address, and together they constitute a system that does something none of them can do alone: surface structural truths about an organization that the organization’s own instruments are constitutionally unable to produce.

The paper’s most novel contribution is the formalization of Normalized Deviance as a distinct detection target. Discovery—finding what nobody knows—is the mesh’s advertised purpose. Detecting what everybody knows and nobody acts on is harder, because the information is not hidden. It is visible, discussed, documented, and ignored. The eigenvalue collapse signature, the linguistic compression leading indicator, and the phase-lock detection framework provide the mesh with a formal apparatus for detecting this second failure mode, and the decay path analysis (Section 7) connects the two: surfaced findings that are not structurally resolved migrate from the first failure mode into the second.

The substrate independence of the underlying mechanism—compression, selection, equilibrium—implies that the architecture may transfer to domains beyond organizational knowledge management. If the mechanism is general, the structural response should be too.

Deployment validates the architecture’s core claims with increasing empirical strength. The mesh’s emergent topology recovers organizational structure without being given it—47 workers self-organizing from 3 into specializations that mirror the organization’s functional domains. Cross-source geometric matching surfaces systemic patterns invisible to any individual team—six tickets, three vendors, one structural failure that no single vantage point could see—and independent cross-source validation confirmed that the mesh detected the same patterns a human engineer independently identified, but earlier, from different evidence, and in a form that persists as a structured artifact rather than a conversation the organization cannot act on. The  $G_D/G_C$  divergence detects load-bearing silence: a code solution to a multi-party organizational impasse existed in the dependency graph while the impasse persisted in the communication graph, undiscovered by any party. An 89-day timestamped trail from policy decision through explicit prediction of communication failure to realized failure provides a measured instance of Crawford-Sobel degradation across an organizational boundary. Confidence calibration discriminates meaningfully across the 40%–95% range, the coupling detector produces honest empty sets rather than false positives, and bridge distance—initially null under sparse graph conditions—now discriminates across the full [0.37, 0.98] range with three-source identity resolution, correctly separating planned single-developer cross-repository work from structurally disconnected engineers converging on the same failure without coordination. Supplementary analysis of internal communication records validates the Normalized Deviance detection framework: action decoupling—discussion substituting for resolution—was present at high confidence across all five recurring problem areas tested, with taxonomic competence functioning as a previously unidentified normalization mechanism: the organization’s fluency at categorizing failures prevents those failures from being treated as the crises they represent. What remains is the longitudinal demonstration: action-correlation across both quadrants, and the decay tracking that will determine whether surfaced findings produce organizational state changes or migrate from Discovery into Normalized Deviance. Preliminary detection latency evidence is encouraging—the mesh correlated a systemic pattern from 30 days of sematectonic traces that a human engineer independently detected only after months of accumulated experience—but the distribution across findings requires additional measurement periods. The theory makes specific, falsifiable predictions. The code is built. The traces are flowing. The predictions are being tested.

## References

- Alexander, C. (1977). *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press.
- Amrit, C. & van Hillegersberg, J. (2010). Exploring the impact of socio-technical core-periphery structures in open source software development. *Journal of Information Technology*, 25(2):216–229.
- Armour, P. G. (2000). The five orders of ignorance. *Communications of the ACM*, 43(10):17–20.
- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Bateson, G. (1972). *Steps to an Ecology of Mind*. University of Chicago Press.
- Beer, M. & Nohria, N. (2000). Cracking the code of change. *Harvard Business Review*, 78(3):133–141.
- Beer, S. (1972). *Brain of the Firm*. Allen Lane.

- Burt, R. S. (2004). Structural holes and good ideas. *American Journal of Sociology*, 110(2):349–399.
- Candès, E. J. & Tao, T. (2006). Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Transactions on Information Theory*, 52(12):5406–5425.
- Cataldo, M., Herbsleb, J. D., & Carley, K. M. (2008). Socio-technical congruence: A framework for assessing the impact of technical and work dependencies on software development productivity. In *Proceedings of the Second ACM-IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*, pp. 2–11.
- Carpenter, G. A. & Grossberg, S. (1987). A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing*, 37(1):54–115.
- Carpenter, G. A., Grossberg, S., & Rosen, D. B. (1991). Fuzzy ART: Fast stable learning and categorization of analog patterns by an adaptive resonance system. *Neural Networks*, 4(6):759–771.
- Conant, R. C. & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2):89–97.
- Cottrell, M. & Fort, J.-C. (1987). Étude d’un processus d’auto-organisation. *Annales de l’Institut Henri Poincaré*, 23(1):1–20.
- Cover, T. M. & Thomas, J. A. (2006). *Elements of Information Theory*. Wiley, 2nd edition.
- Crawford, V. P. & Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.
- Estrada, E. & Hatano, N. (2008). Communicability in complex networks. *Physical Review E*, 77(3):036111.
- Donnat, C., Zitnik, M., Hallac, D., & Leskovec, J. (2018). Learning structural node embeddings via diffusion wavelets. In *Proceedings of the 24th ACM SIGKDD*, pp. 1320–1329.
- Elliott, M. (2007). *Stigmergic Collaboration: A Theoretical Framework for Mass Collaboration*. PhD thesis, University of Melbourne.
- Flyvbjerg, B., Budzier, A., Lee, J. S., Keil, M., Lunn, D., & Bester, D. W. (2022). The empirical reality of IT project cost overruns: Discovering a power-law distribution. *Journal of Management Information Systems*, 39(3):607–639.
- Granovetter, M. S. (1973). The strength of weak ties. *American Journal of Sociology*, 78(6):1360–1380.
- Fritzke, B. (1995). A growing neural gas network learns topologies. In *Advances in Neural Information Processing Systems 7*, pp. 625–632.
- Gigerenzer, G. (2008). *Rationality for Mortals: How People Cope with Uncertainty*. Oxford University Press.
- Grassé, P.-P. (1959). La reconstruction du nid et les coordinations interindividuelles chez *Belligositermes natalensis* et *Cubitermes sp.* *Insectes Sociaux*, 6(1):41–80.
- Grossberg, S. (1976). Adaptive pattern classification and universal recoding: I. Parallel development and coding of neural feature detectors. *Biological Cybernetics*, 23(3):121–134.

- Heylighen, F. (2016). Stigmergy as a universal coordination mechanism: I. Definition and components. *Cognitive Systems Research*, 38:4–13.
- Kauffman, S. A. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press.
- Kamenica, E. & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.
- Lawvere, F. W. (1969). Diagonal arguments and cartesian closed categories. In *Category Theory, Homology Theory and their Applications II*, Lecture Notes in Mathematics 92, pp. 134–145. Springer.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43(1):59–69.
- Liberti, J. M. & Mian, A. R. (2009). Estimating the effect of hierarchies on information use. *The Review of Financial Studies*, 22(10):4057–4090.
- Luhmann, N. (1995). *Social Systems*. Stanford University Press.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1):71–87.
- McEntire, J. (2025). Dysmemic pressure: Selection dynamics in organizational information environments. *arXiv preprint arXiv:2512.14716*. Also available at SSRN: <https://ssrn.com/abstract=6015814>.
- McEntire, J. (2026). Compression, selection, and organizational self-deception: A unified theory. Available at SSRN: <https://ssrn.com/abstract=6015774>.
- McEntire, J. (2025). Tournament-based performance evaluation and systematic misallocation: Why forced ranking systems produce random outcomes. *arXiv preprint arXiv:2512.06583*.
- McEntire, J. (2026). The discipline of rigorous inquiry. Preprint. Available at <https://cageandmirror.com/research/rigorous-inquiry>.
- Ocasio, W. (1997). Towards an attention-based view of the firm. *Strategic Management Journal*, 18(S1):187–206.
- Prigogine, I. (1977). *Self-Organization in Nonequilibrium Systems*. Wiley.
- Prendergast, C. (1993). A theory of “yes men.” *American Economic Review*, 83(4):757–770.
- Shannon, C. E. (1959). Coding theorems for a discrete source with a fidelity criterion. In *IRE National Convention Record*, Part 4, pp. 142–163.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2):129–138.
- Smaldino, P. E. & McElreath, R. (2016). The natural selection of bad science. *Royal Society Open Science*, 3(9):160384.
- Spencer-Brown, G. (1969). *Laws of Form*. Allen & Unwin.
- Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374.
- Standish Group (2015). *CHAOS Report 2015*. The Standish Group International.

- Tang, J., Li, J., Gao, Z., & Li, J. (2022). Rethinking graph neural networks for anomaly detection. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, pp. 21076–21089.
- Tamburri, D. A., Palomba, F., & Kazman, R. (2021). Exploring community smells in open-source: An automated approach. *IEEE Transactions on Software Engineering*, 47(3):630–652.
- Tetlock, P. E. (2005). *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton University Press.
- Presidential Commission on the Space Shuttle *Challenger* Accident (1986). *Report of the Presidential Commission on the Space Shuttle Challenger Accident*. Washington, DC.
- Columbia Accident Investigation Board (2003). *Report of the Columbia Accident Investigation Board*, Volume 1. NASA.
- Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. University of Chicago Press.
- Vuori, T. O. & Huy, Q. N. (2016). Distributed attention and shared emotions in the innovation process: How Nokia lost the smartphone battle. *Administrative Science Quarterly*, 61(1):9–51.
- Wilson, E. O. (1971). *The Insect Societies*. Harvard University Press.
- Wittgenstein, L. (1953). *Philosophical Investigations*. Blackwell.
- Zheng, J. & Meister, M. (2025). The unbearable slowness of being: Why do we live at 10 bits/s? *Neuron*, 113(2):192–204.