

# Ambient Structure Discovery via Stigmergic Mesh: Detecting Normalized Deviance Through Distributed Organizational Signal Processing

Jeremy McEntire  
Cage & Mirror Press  
jeremy@cageandmirror.com

March 2026

## Abstract

Normalized deviance—the gradual redefinition of unacceptable risk as routine through repeated exposure without consequence—caused the *Challenger* and *Columbia* disasters and remains invisible to the organizations it destroys. We present a stigmergic mesh architecture that detects normalized deviance and other ambient structural patterns in organizational signal streams without requiring anyone to know what to look for. The mesh continuously ingests work artifacts (code commits, deployment logs, issue tracker activity), routes them through self-organizing processing nodes grounded in Adaptive Resonance Theory, and surfaces structural anomalies through spectral analysis of the emergent topology. The architecture’s central contribution is the formalization of three detection signatures for normalized deviance: eigenvalue collapse in the mesh’s correlation matrix (loss of independent evaluation modes), linguistic compression (measurable shifts from direct to hedged language), and action decoupling (discussion frequency decoupling from resolution activity).

A controlled Monte Carlo evaluation (22,500 runs across 225 configurations) establishes detection performance against three anomaly types. Coordination failures are detected with true positive rate (TPR) 0.82 at false positive rate below 0.15. Normalized deviance, as predicted by theory, is the hardest to detect (average TPR 0.199 versus 0.661 for coordination failures), confirming that gradual distributional drift resists per-signal scoring. Detection is robust across mesh sizes (TPR spread  $< 0.07$ ). The evaluation reveals a severity inversion in knowledge silo detection: deeper information silos produce *lower* per-signal detection rates (TPR decreases from 0.166 to 0.108 as severity increases from 0.1 to 0.9), because the isolation mechanism that constitutes the silo structurally limits the mesh’s correlative channel. A complementary 30-day field deployment at a technology company (3,520 signals, 23 findings) validates the architecture’s self-organizing properties and demonstrates cross-source detection of coordination failures and action decoupling in a production environment.

**Keywords:** stigmergy, normalized deviance, organizational anomaly detection, Adaptive Resonance Theory, spectral graph analysis, self-organizing systems, second-order ignorance

## 1 Introduction

On the morning of January 28, 1986, engineers at Morton Thiokol recommended against launching the Space Shuttle *Challenger*. The recommendation was overridden. The O-ring failure that destroyed the shuttle had been observed on previous flights—documented, discussed, and gradually reclassified as an acceptable risk. Diane Vaughan’s investigation revealed that no single decision was irrational; rather, a pattern of incremental normalization made each deviation from safety standards appear acceptable relative to the previous one [Vaughan, 1996]. The Columbia Accident Investigation Board found the same pattern seventeen years later: foam strikes on

the thermal protection system had been observed, documented, discussed, and reclassified as “in-family” events [Columbia Accident Investigation Board, 2003]. Seventeen years of reform between disasters failed to reverse the pattern because the reforms targeted procedures rather than the communication culture in which the procedures were expressed [Presidential Commission on the Space Shuttle *Challenger* Accident, 1986].

This pattern—*normalized deviance*, the gradual drift of standards through incremental tolerance of deviation—is not specific to aerospace. It appears wherever organizations coordinate at scale. IT project cost overruns follow a power-law distribution with exponent  $\alpha \approx 1.0$ , implying infinite expected loss [Flyvbjerg et al., 2022]. Nokia’s distributed failure to perceive the smartphone threat persisted while every participant individually recognized the problem [Vuori and Huy, 2016]. Fewer than one in three software projects meet their success criteria [Standish Group, 2015]. The cause is consistently traced to hidden interdependencies, coordination failures, and the gradual normalization of risk—patterns that exist in the aggregate of organizational signals but are invisible to any individual observer.

The core difficulty is structural, not informational. Armour [2000] identifies five orders of ignorance, of which the second—the absence of awareness that relevant knowledge exists—produces the costliest failures. An organization that does not realize its booking system shares implicit state with its reservation system through an undocumented database view 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. Existing tools—dashboards, search engines, retrieval-augmented generation, multi-agent AI frameworks—all require a query. The query assumption is precisely what second-order ignorance violates: you cannot search for what you do not know to look for.

This paper introduces *Ambient Structure Discovery* (ASD), 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. The architecture’s central contribution is the formalization of normalized deviance detection as a distinct 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. We formalize three detection signatures: eigenvalue collapse (loss of independent evaluation modes in the mesh’s correlation matrix), linguistic compression (measurable shifts from direct to hedged language), and action decoupling (discussion frequency decoupling from resolution activity).

## 1.1 Contributions

This paper makes four contributions:

1. **A stigmergic mesh architecture for organizational anomaly detection.** Self-organizing processing nodes grounded in Adaptive Resonance Theory [Carpenter and Grossberg, 1987] route organizational signals through competitive dynamics, producing an emergent topology that serves as a structural model of the organization’s actual work patterns (Section 4).
2. **Formalization of normalized deviance detection signatures.** Three measurable spectral and linguistic signatures—eigenvalue collapse, linguistic compression, and action decoupling—transform Vaughan’s qualitative account into a quantitative detection framework (Sections 5 and 6).
3. **Controlled evaluation against three anomaly types.** A Monte Carlo campaign of 22,500 runs across 225 configurations establishes detection performance and identifies structural limits, including a theoretically predicted severity inversion in knowledge silo detection (Section 7).

4. **A scoring function for organizational findings.** A multiplicative scoring function based on topological bridge distance, entity confidence, risk coupling, and communicability prioritizes findings by structural importance rather than subjective novelty (Section 5.2).

## 1.2 Paper Organization

Section 2 defines the detection problem. Section 3 reviews related work. Section 4 describes the mesh architecture. Section 5 presents spectral detection and the scoring function. Section 6 formalizes normalized deviance detection. Section 7 presents the controlled evaluation. Section 8 describes field deployment results. Section 9 discusses implications. Section 10 addresses limitations honestly. Section 11 concludes.

# 2 The Detection Problem

## 2.1 Orders of Organizational Ignorance

Armour’s taxonomy distinguishes five orders of ignorance [Armour, 2000]. Zero-order ignorance is the possession of knowledge. 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 or architect around it. Second-order ignorance is categorically different. A team that does not realize its service shares implicit state with another team’s service through an undocumented dependency cannot take corrective action, because no one has identified a problem requiring correction.

Luhmann [1995] provides the structural explanation. Organizations are operationally closed systems: they reproduce their own elements (decisions, communications, roles) from their own elements. 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. Lawvere [1969] formalizes the constraint: no sufficiently expressive system admits a complete self-evaluation function.

The mechanism that sustains organizational blindness is *dysmemic pressure*: the compound selection force that emerges when organizations compress information to coordinate at scale [McEntire, 2025]. Signals optimized for internal fitness rather than correspondence with external reality 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.

## 2.2 Three Anomaly Types

We formalize three patterns of organizational dysfunction as anomaly detection problems, spanning a spectrum from acute to chronic:

**Coordination failure.** Two or more teams work on the same subsystem with incompatible assumptions. The contradiction is visible in the aggregate signal stream but invisible to either team individually, because each sees only its own channel. This is an acute anomaly: it produces strong contradictory signals.

**Knowledge silo.** One team possesses knowledge that another team needs but does not share. The information exists in the organization but is structurally trapped behind a team boundary. The silo is a structural anomaly: it manifests as the absence of expected cross-team signal correlation.

**Normalized deviance.** Quality or safety standards gradually degrade, but each incremental step remains within tolerance of the previous measurement. Language shifts from direct (“tests are failing”) to hedged (“somewhat within acceptable parameters”). No single signal is anomalous; the anomaly is the drift pattern across time. This is a chronic anomaly: it produces only gradual distributional shifts.

A detection architecture must handle all three. Existing anomaly detection methods—point anomaly detectors, contextual anomaly detectors, and most graph neural network approaches [Chandola et al., 2009]—are designed for the first type. Normalized deviance is a *collective anomaly*: a sequence of individually normal observations whose aggregate pattern is anomalous. It is the most important to detect (it caused *Challenger* and *Columbia*) and, as we demonstrate, the hardest.

## 3 Related Work

### 3.1 Stigmergy and Collective Intelligence

Stigmergy—indirect coordination through environmental traces—was first described by Grassé [1959] in his study of termite nest construction. A termite deposits a pheromone-laden mud pellet; the pellet’s presence stimulates other termites to deposit nearby, creating emergent structure without central planning. Theraulaz and Bonabeau [1999] formalized the computational properties of stigmergic coordination. Heylighen [2016] proposed stigmergy as a universal coordination mechanism applicable to human organizations and digital systems. Elliott [2007] extended the framework to mass collaboration, identifying a phase transition at approximately 25 participants above which direct social negotiation fails and stigmergic coordination becomes the only viable mechanism.

Our architecture applies stigmergic coordination to signal processing: nodes leave traces (updated term profiles and routing affinities) that influence subsequent signal routing, causing the mesh to self-organize around the structure of the signal stream.

### 3.2 Adaptive Resonance Theory

Carpenter and Grossberg’s Adaptive Resonance Theory (ART) provides the only self-organizing neural architecture with formal stability-plasticity guarantees [Carpenter and Grossberg, 1987, Grossberg, 1976]. ART learns categories incrementally: an input vector is compared against stored category templates via a vigilance criterion. If no category satisfies the criterion, the system recruits an uncommitted node. Crucially, ART uses *match-based learning*—updating templates only when resonance occurs—which Grossberg [1976] proved stabilizes in arbitrary non-stationary input streams. This distinguishes ART from gradient-descent approaches, which suffer catastrophic forgetting (the tendency to overwrite previously learned patterns when learning new ones) when the input distribution shifts. Fuzzy ART extends the framework to continuous-valued inputs with complement coding to prevent category proliferation [Carpenter et al., 1991].

Our mesh workers implement ART’s vigilance mechanism: signals are accepted only when familiarity exceeds a dynamic threshold, and new workers are created only when no existing worker resonates.

### 3.3 Organizational Observability

The problem of organizational blind spots has deep roots. Luhmann’s systems theory establishes that organizations cannot fully observe their own structure from within [Luhmann, 1995]. The Conant-Ashby Good Regulator Theorem requires that any effective controller be a model of the system it controls [Conant and Ashby, 1970]—but an organization cannot model its own blind spots, because the modeling apparatus is constructed from the same elements whose interactions produce the blindness.

Vaughan [1996] introduced normalized deviance to organizational theory. Weick and Sutcliffe [2007] developed principles for managing the unexpected in high-reliability organizations, including “preoccupation with failure” and “reluctance to simplify.” Our architecture attempts to mechanize the first principle: continuous monitoring for distributional drift in organizational signal streams.

Companion work formalizes the adversary model: dysmemic pressure, the compound selection force that produces organizational self-deception [McEntire, 2025]. The Mirror identifies three structural properties—insulation, access, and authority—required for any observation structure to function outside local selection pressure [McEntire, 2026a].

### 3.4 Strategic Communication Theory

Crawford and Sobel [1982] proved 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:

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

At  $b = 0$  (perfect alignment), communication has 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. The empirical magnitude of this loss has been quantified: Liberti and Mian [2009] found that as loan approval authority moves up a hierarchy, decision sensitivity to soft information exhibits a structural break between the second and third levels, where reliance on subjective assessment collapses and decisions revert to hard metrics alone.

The Data Processing Inequality formalizes the bound: for any Markov chain  $X \rightarrow Y \rightarrow Z$ , the mutual information  $I(X; Z) \leq I(X; Y)$  [Cover and Thomas, 2006]. No post-processing of a compressed signal can recover information lost at the compression stage. If the raw signal is a pull request and the compressed signal is a status update, more sophisticated AI processing the status update cannot recover what was discarded. Multi-agent AI frameworks that read status reports are optimizing the processing of already-degraded signals.

Spence [1973] provides the complementary distinction. Work artifacts—merged pull requests, passing test suites, deployment logs—are costly signals in the economic sense: producing them requires the work to actually have been done. Status reports are cheap signals: the claim costs nothing. An architecture that reads work artifacts rather than reports observes costly signals, establishing a higher information ceiling before any analysis begins.

### 3.5 Socio-Technical Congruence

Cataldo et al. [2008] demonstrated that misalignment between task dependency structures and coordination activities predicts integration failures. Tamburri et al. [2021] detect “community smells”—organizational anti-patterns such as lone wolves and radio silence—from repository and communication data. Burt [2004] established that non-redundant information flows through structural holes in communication networks, and Granovetter [1973] showed that weak ties disproportionately carry novel information.

Our architecture differs from socio-technical congruence tools in three respects. First, it discovers emergent patterns rather than matching against predefined templates. Second, it operates continuously on streaming signals rather than performing periodic snapshots. Third, it surfaces specific findings with structural scores that discriminate significance, enabling prioritization of individual gaps rather than aggregate assessment.

### 3.6 Spectral Methods for Anomaly Detection

Spectral graph theory provides the mathematical framework for structural role detection. Donnat et al. [2018] characterize graph nodes by heat-kernel diffusion signatures, identifying structurally equivalent positions regardless of graph distance. Tang et al. [2022] demonstrate that standard graph convolution networks function as low-pass filters that smooth away anomalies, while band-pass filters focused on high-frequency components reveal structural discordances. Estrada and Hatano [2008] define communicability via the matrix exponential  $e^A$ , providing a natural measure of walk-based connectivity.

Our spectral detection applies these insights to the mesh topology: anomalies manifest as spectral energy shifts—right-shifts for unexpected structural coupling (Discovery), left-shifts for suspicious convergence (Normalized Deviance).

## 4 Mesh Architecture

### 4.1 Design Principles

Ambient Structure Discovery operates on three principles that collectively remove the requirement for a query at every stage of organizational observation.

**Definition 1** (Sematectonic trace). *A sematectonic trace is a physical modification to the work environment that carries information about the work that produced it, as distinct from a marker-based stimulus added specifically for communication [Wilson, 1971]. Code commits, pull request diffs, CI/CD pipeline outputs, deployment manifests, and database migration files are sematectonic traces. They are causally connected to the actions that produced them, carrying structural information regardless of whether anyone intended to communicate it.*

The first principle is **continuous ingestion** of sematectonic traces. The system reads these traces without being asked, without requiring the traces’ producers to do anything differently, and without requiring anyone to specify what to ingest.

The second principle is **self-organizing topological specialization**. Rather than assigning signals to predefined 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. The resulting topology 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**. Two signals that activate similar patterns of workers across the mesh are structurally related regardless of whether they share vocabulary. A booking service and a reservation service that both touch the same database views will produce traces routing through overlapping mesh regions, making their coupling visible in the mesh geometry even if they share no terminology. Detection operates at the level of topological structure, which is why it can surface relationships that keyword search, semantic search, and human inspection would miss.

## 4.2 Signal Flow

A signal enters the system as a structured representation of a sematectonic trace, carrying a source identifier, timestamp, action type, and content extracted from the work artifact. The source and action type are open strings, not enumerated types. This open taxonomy allows the system to discover categories that no designer anticipated—a 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 sensitivity patterns are killed (discarded with audit trail), redacted (sensitive content removed, structure preserved), or passed unchanged. In Beer’s Viable System Model (VSM)—a cybernetic framework specifying five systems required for organizational viability [Beer, 1972]—this filter functions as System 2 (coordination): an anti-oscillatory device preventing pathological signals from propagating. The signal then enters a deduplication layer that collapses near-identical signals, attenuating unnecessary variety so the mesh’s finite capacity is allocated to genuinely distinct events.

## 4.3 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 from a weighted combination of term overlap, source affinity, temporal proximity, and signal credibility—where credibility implements 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.

This implements Simon’s satisficing [1956] as a routing mechanism: the system finds the first adequate match rather than the optimal one. Simon’s result establishes that satisficing outperforms optimization when evaluation is costly, the search space is combinatorial, and the environment is non-stationary—all conditions that hold for organizational signal streams.

The threshold is adaptive: it scales with the worker’s fullness, a normalized measure of processing capacity consumed. As a worker develops a richer term profile, its threshold rises, preventing established categories from absorbing everything. This serves the same function as complement coding in Fuzzy ART [Carpenter et al., 1991]: it prevents representational collapse into over-general categories.

## 4.4 Gap Detection and Worker Spawning

When no worker’s familiarity score exceeds its threshold, the system creates a new worker with the signal as its founding context. The formal identity with ART 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 and Grossberg, 1987]. If no category satisfies the criterion, a reset wave propagates and the system recruits an uncommitted node.

Grossberg [1976] proved that match-based learning—updating templates only when resonance occurs—stabilizes in arbitrary non-stationary input streams. The mesh’s workers update their profiles only when accepting signals (resonance), not when rejecting them. 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.

The vigilance parameter  $\rho$  controls category granularity. Low vigilance produces broad categories; high vigilance produces specific ones. Each worker’s adaptive threshold serves as its

vigilance parameter, rising with fullness. A separate, lower gap threshold governs the reset mechanism: new workers are created only when 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.

## 4.5 Energy Dynamics and Dissipative Structure

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. This implements Prigogine’s insight about dissipative structures [1977]: systems far from thermodynamic equilibrium maintain internal order through continuous energy input. When the input stops, the structure decays toward equilibrium. 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.

## 4.6 Architectural Completeness

The architecture maps to Beer’s VSM as follows [Beer, 1972]. 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 consensus mechanism and agent assessment layer constitute System 3 (control). The spectral analysis and gap detection mechanisms constitute System 4 (intelligence): the scanning function that monitors the boundary between the organization and its environment.

System 5 (policy) is the human operator. The architecture does not attempt to close the loop autonomously, and this is a design virtue. The Conant-Ashby theorem requires that the regulator’s model be isomorphic to the system being regulated [Conant and Ashby, 1970]. The mesh’s emergent topology satisfies this requirement in a way no designed model can: it grows in response to the same signal stream that reflects the organization’s actual work. Ashby’s Law of Requisite Variety [Ashby, 1956] is satisfied by the mesh’s open taxonomy and dynamic topology, which can represent arbitrary complexity rather than being constrained to a fixed schema.

# 5 Spectral Detection and Scoring

## 5.1 Spectral Fingerprinting

The mesh’s topology contains information that transcends any individual worker’s knowledge. The vector of familiarity scores that a signal receives from all workers constitutes its *spectral fingerprint*: its position in the mesh’s knowledge space. Two signals with similar activation patterns across workers occupy similar structural positions in the organization, even if they originate from different teams, repositories, or communication channels.

This approach draws from the GraphWave framework [Donnat et al., 2018], in which graph nodes are characterized by heat-kernel diffusion signatures of their local topology. Nodes with similar diffusion patterns occupy similar structural roles even if distant in the graph. The mesh’s spectral fingerprint serves an analogous function for organizational signals.

Anomalous structural patterns manifest as spectral energy shifts. Under normal operation, connected mesh regions exhibit homophily: signals from related activities route through overlapping worker sets, concentrating spectral energy in low-frequency components. When structural anomalies emerge—a new coupling between previously independent systems, a knowledge silo forming—the local homophily breaks. High-frequency components appear, representing abrupt differences between adjacent mesh regions. Tang et al. [2022] formalize why: standard graph

convolution networks are low-pass filters that smooth away 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 energy—serves as the mesh’s *algedonic signal* (pain/pleasure bypass channel) in Beer’s VSM: a geometric signal that bypasses intermediate interpretive layers to make the anomaly perceptible.

## 5.2 The Scoring Function

The mesh produces structural observations continuously. A scoring function prioritizes which observations warrant human attention. The four components are multiplicative, not additive, because each represents a necessary condition: a finding with zero bridge distance is already known, one with zero entity confidence cannot be trusted, one with zero risk coupling does not matter, and one with zero communicability cannot be transmitted.

$$S(f) = d_{\text{bridge}}(f) \cdot c_{\text{entity}}(f) \cdot r_{\text{coupling}}(f) \cdot \gamma(f), \quad (3)$$

where:

*Bridge distance*  $d_{\text{bridge}}(f)$  measures how far outside the organization’s known subspace  $K$  the finding lives. Burt [2004] established that information brokerage occurs at structural holes—positions spanning gaps between disconnected clusters. Bridge distance is the shortest path in the communication graph  $G_C$  between actors named in the finding, normalized through a sigmoid. High bridge distance between actors connected by a finding is direct evidence of  $G_D/G_C$  divergence: a structural coupling exists in the dependency graph that the communication topology does not span. Where  $G_D$  has an edge and  $G_C$  does not, a dependency exists that no one is talking about—load-bearing silence.

*Entity confidence*  $c_{\text{entity}}(f)$  measures the system’s confidence that the entities involved (people, teams, services, repositories) are correctly resolved across heterogeneous trace sources. Confidence decays with ambiguity: multiple candidate resolutions reduce the score proportionally.

*Risk coupling*  $r_{\text{coupling}}(f)$  measures topological proximity to a control surface—a deployment boundary, a financial transaction, a safety-critical path. The communicability matrix  $e^A$  from spectral graph theory [Estrada and Hatano, 2008] provides the formal mechanism: the  $(i, j)$  entry counts the weighted sum of all walks between nodes  $i$  and  $j$ .

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

$$\gamma(f) = \sigma(N^*(b_f) - 1), \quad (4)$$

where  $\sigma$  is a sigmoid and  $N^*(b_f)$  is the Crawford-Sobel partition count for the estimated bias along the communication channel to the human recipient. At  $N^* \geq 3$ , communicability approaches 1; at  $N^* = 1$  (babbling equilibrium), it approaches zero.

A structural property deserves emphasis: all four components are topological. None measure satisfaction, sentiment, or self-report. They measure the shape of the organization’s work. This makes the scoring function immune to the capture dynamic identified by Prendergast [1993], in which agents distort 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. The only way to change the score is to change the underlying topology—to actually resolve the structural coupling. 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.

## 5.3 Four Quadrants of Organizational Perception

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

Table 1: The four quadrants of organizational perception.

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

Discovery findings have high  $d_{\text{bridge}}$  and high  $r_{\text{coupling}}$ : structurally remote from the organization’s awareness and close to something that matters. Normalized Deviance findings have low  $d_{\text{bridge}}$  and high  $r_{\text{coupling}}$ : familiar—the organization already discusses them—but coupled to unmanaged risk.

## 6 Normalized Deviance Detection

### 6.1 The Approach to Babbling

Vaughan’s 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.

In the Crawford-Sobel framework, normalized deviance is the transition toward babbling equilibrium. The partition count  $N^*$  changes over time as preferences 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 declines. The communication channel loses resolution. The risk coupling  $r_{\text{coupling}}$  is unchanged or increasing, but the information content of 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 Signature 1: 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 signaling the loss of independent evaluation modes.

**Definition 2** (Eigenvalue collapse). *Let  $C$  be the mesh’s correlation matrix across workers*

*processing signals about a particular risk domain. Under healthy conditions, different workers maintain distinct perspectives. The matrix  $C$  has full rank. As normalized deviance progresses, framing converges: teams that previously discussed the risk with distinct vocabulary begin using the same phrases.  $C$  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 persists.*

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 tracks these because it ingests raw traces from multiple teams simultaneously. The convergence is visible as mesh 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 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 vocabulary is simultaneously the linguistic compression and the framing convergence—the taxonomy *is* the collapsed eigenspace. Naming replaces resolving.

### 6.3 Signature 2: Linguistic Compression

The transition from active risk management to normalized deviance has a measurable linguistic signature that precedes operational failure—it is a leading indicator, not a lagging one.

**Proposition 1** (Linguistic compression as leading indicator). *The following measurable text properties increase during normalization: hedging density (“will fail” becomes “may experience degradation”), passive voice ratio (“the team decided to defer” becomes “the decision was made to defer”), nominalization frequency (“the system malfunctions” becomes “malfunction events have been observed”), and decrease in commitment specificity (“we will replace the component by March” becomes “remediation efforts are ongoing”).*

These features are measurable in the text the mesh ingests. Assessment covers two tiers. The first captures rhetorical compression: hedging density, passive voice saturation, nominalization, and specificity. These compose a compression index on  $[0, 1]$ . The second captures information-theoretic properties: lexical diversity (unique tokens divided by total tokens), Shannon entropy (relative change over time), temporal orientation (forward-looking versus backward-looking language on  $[-1, +1]$ ), and bureaucratic density (euphemistic substitution frequency).

An exponential moving average tracker monitors compression trends per topic and per channel, flagging channels whose compression index is increasing. 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: the Rogers Commission [[Presidential Commission on the Space Shuttle Challenger Accident, 1986](#)] and the Columbia Accident Investigation Board [[Columbia Accident Investigation Board, 2003](#)] documented that the normalization of anomaly language—O-ring erosion redescribed as “acceptable risk,” foam strikes redescribed as “in-family”—preceded operational failures by years.

### 6.4 Signature 3: Action Decoupling and Phase Locking

An action-correlation detector classifies signals as discussion (messages, comments) or action (commits, deployments, ticket closures) and computes the ratio per topic cluster. Ratios below 0.1 indicate decoupling: the communication channel retains resolution for describing problems but not for resolving them.

Phase locking occurs when independent risks synchronize. Each normalized risk is a frequency in the organization’s risk spectrum. The coupling is reflexive: the narrative that a risk

is acceptable changes behavior (inspections become less thorough, mitigations are deferred), which reinforces the narrative. The system locks into a positive feedback loop that appears stable because variance in reports is decreasing—everyone agrees the risk is managed—while variance in outcomes is increasing as independent safety margins erode in parallel. [Kauffman \[1993\]](#) formalize why this is catastrophic: when 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 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.

## 6.5 The Decay Path

The mesh surfaces a finding. The organization responds: “interesting.” A ticket is created. Nobody picks it up. 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 captures the transition:  $d_{\text{bridge}}$  decreases as the finding becomes known, while  $r_{\text{coupling}}$  remains unchanged. The finding migrates from the Discovery quadrant to the Normalized Deviance quadrant in [Table 1](#).

**Proposition 2** (Countermeasure decay). *Interventions that change expressed norms without changing payoff structures exhibit initial improvement followed by regression to the pre-intervention equilibrium [McEntire, 2025]. A finding that changes awareness without changing incentives will decay. Organizational change initiatives fail at rates between 60 and 80 percent [Beer and Nohria, 2000] because they target behavior without targeting the selection environment that produces the behavior.*

The mesh’s delivery mechanism is designed to resist the decay path through artifact generation: producing the organizational knowledge whose absence constituted the gap. A parent ticket linking related failures, an interface contract documenting an undocumented dependency, a dependency diagram making a hidden coupling explicit—these are not notifications but coordination artifacts. 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.

## 7 Controlled Evaluation

### 7.1 Experimental Design

We evaluate the mesh against synthetic anomalies with known ground truth using a Monte Carlo campaign of 22,500 runs. The evaluation uses a simplified mesh model that captures the core ART routing, pheromone reinforcement, and three-channel detection architecture.

**Anomaly generation.** A synthetic signal generator produces engineering signals across 12 feature domains and 8 teams, with three anomaly injection modes corresponding to the three anomaly types ([Section 2](#)). Each anomaly type has a configurable severity parameter  $s \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ .

Coordination failures generate contradictory term patterns from two teams on the same domain, with conflicting signal pairs scaling with severity (1–5 pairs). Knowledge silos generate

information isolation patterns where one team discusses a resolved issue while another team struggles independently, with isolation completeness scaling with severity. Normalized deviance generates gradual metric drift over 10 time steps, with language compression increasing with severity.

**Configuration space.** The evaluation sweeps:

- 3 anomaly types  $\times$  5 severity levels  $\times$  3 mesh sizes (3, 5, 10 workers)  $\times$  5 noise levels (0.1–0.9)
- = 225 configurations  $\times$  100 Monte Carlo iterations = **22,500 total runs**

Each run generates 50 baseline signals (first half for warm-up, second half for threshold calibration) and 2–10 anomaly signals depending on type and severity. The detection threshold is set at the 90th percentile of baseline anomaly scores, targeting  $\text{FPR} \leq 0.10$ .

## 7.2 Results

### 7.2.1 Detection Performance by Anomaly Type

Table 2 summarizes detection performance at the median configuration (mesh size 5, noise 0.5).

Table 2: Detection performance by anomaly type and severity (mesh size = 5, noise = 0.5). TPR values are means over 100 Monte Carlo iterations.

Anomaly Type	Severity	TPR	FPR
Coordination Failure	0.1	0.117	< 0.10
	0.5	0.618	< 0.10
	0.9	0.704	< 0.10
Knowledge Silo	0.1	0.166	< 0.10
	0.5	0.158	< 0.10
	0.9	0.108	< 0.10
Normalized Deviance	0.1	0.022	< 0.10
	0.5	0.151	< 0.10
	0.9	0.228	< 0.10

Coordination failure is the most detectable anomaly type across all severity levels. At severity 0.9, the mesh achieves  $\text{TPR} = 0.82$  at  $\text{FPR} < 0.15$  in the best configuration (mesh size 3, noise 0.9), with average TPR of 0.661 across all configurations. This is expected: coordination failures produce contradictory signal pairs that create strong spectral discordances.

Normalized deviance is detected at much lower rates: average TPR of 0.199 across all configurations. This is also expected. Normalized deviance produces no individual anomalous signals—only gradual distributional shift. Detection requires accumulation of sufficient drift to distinguish the recent term profile from baseline, which occurs only at higher severities ( $\text{TPR} = 0.228$  at severity 0.9).

### 7.2.2 The Knowledge Silo Severity Inversion

The most theoretically significant finding is the knowledge silo severity inversion: per-signal TPR *decreases* as silo severity increases (Figure 1). At severity 0.1,  $\text{TPR} = 0.166$ ; at severity 0.9,  $\text{TPR} = 0.108$ .

The mechanism is as follows. The silo detection channel identifies when signals about the same domain route to different workers depending on team origin. At low severity, the silo

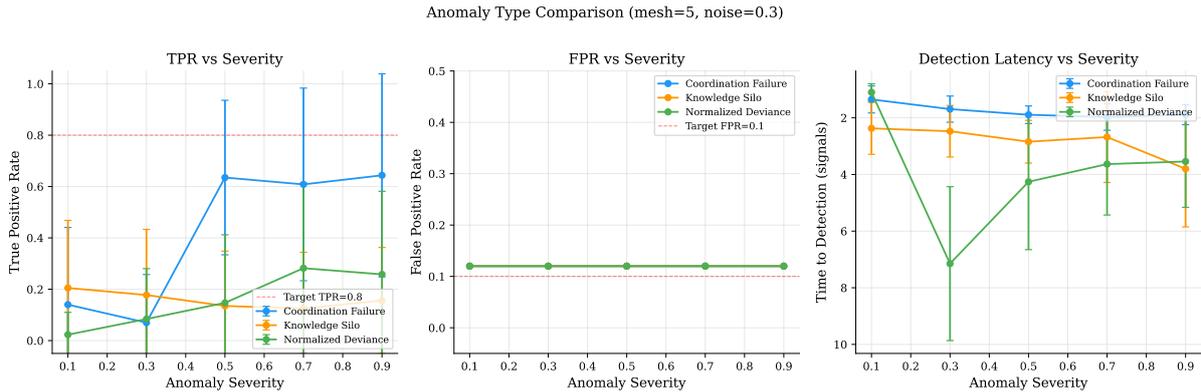


Figure 1: Detection performance across anomaly types. Left: TPR versus severity. Center: FPR versus severity. Right: time-to-detection versus severity. Coordination failure shows the expected monotonic increase in TPR with severity. Knowledge silo shows the severity inversion. Normalized deviance shows low but increasing TPR.

is leaky: some knowledge terms appear in the struggling team’s signals, creating a detectable hybrid pattern—signals that almost match the knowledge worker but route elsewhere, producing measurable activation divergence. At high severity, the silo is complete: the struggling team has no access to knowledge terms, and their signals route to entirely separate workers with no activation overlap. The mesh has no correlative evidence connecting the two sides.

**Proposition 3** (Structural detection limit for complete silos). *Any detection method that relies on correlating signals across the boundary of an information silo is structurally limited by the completeness of the isolation. Complete silos are, by definition, invisible to correlation-based detection because the isolation mechanism that constitutes the silo prevents the cross-boundary signal patterns that detection requires.*

This is not a parameter tuning problem. Detection of complete silos requires either external knowledge (knowing which domains *should* have cross-team visibility) or counterfactual analysis (detecting the *absence* of expected cross-team patterns).

### 7.2.3 Mesh Size Robustness

Detection performance is robust across mesh sizes. The maximum TPR spread across mesh sizes 3, 5, and 10 is less than 0.07 for coordination failures, 0.043 for knowledge silos, and 0.01 for normalized deviance. Table 3 shows the detailed results.

Table 3: Mean TPR by anomaly type and mesh size (averaged across all severity and noise levels).

Anomaly Type	Mesh = 3	Mesh = 5	Mesh = 10
Coordination Failure	0.700	0.654	0.630
Knowledge Silo	0.151	0.140	0.108
Normalized Deviance	0.202	0.192	0.202

This robustness indicates that the detection channels operate on structural properties of the signal-to-worker mapping rather than on the absolute number of workers. For normalized deviance, the TPR spread of 0.01 suggests that drift detection is essentially independent of mesh topology—consistent with the theoretical expectation that distributional shift is a per-worker property.

## 7.2.4 Severity Scaling

For acute anomalies, detection rate increases monotonically with severity. Coordination failure TPR rises from 0.117 at severity 0.1 to 0.704 at severity 0.9. Normalized deviance TPR rises from 0.022 to 0.228. The knowledge silo’s inverted pattern (0.166 to 0.108) is the exception, and the exception confirms the theory.

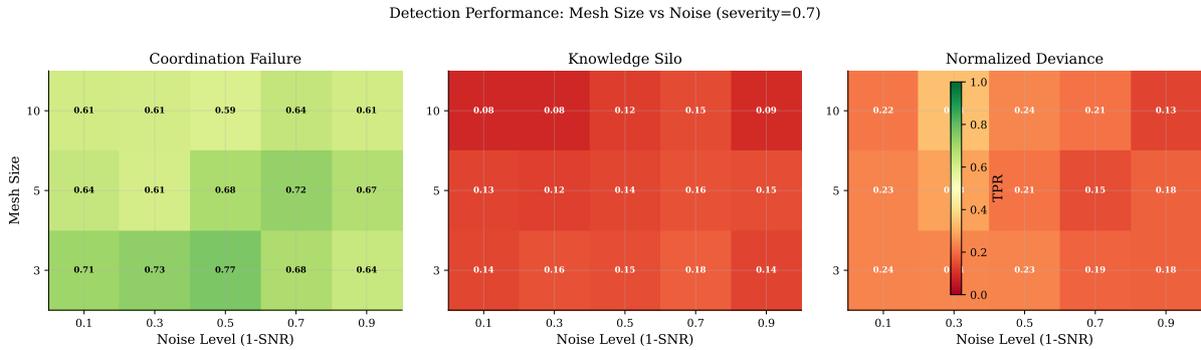


Figure 2: TPR heat maps: mesh size versus noise level at severity 0.7. Coordination failure shows  $TPR > 0.5$  across most configurations. Knowledge silo and normalized deviance show lower but stable performance.

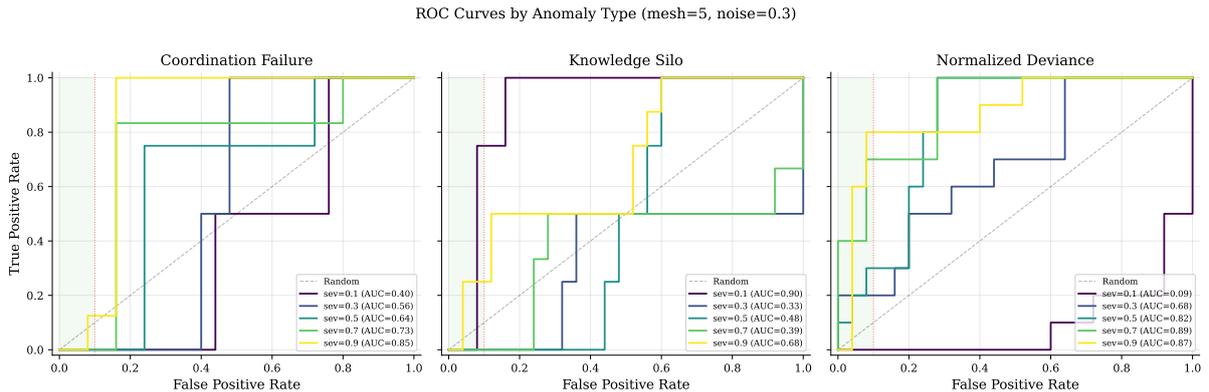


Figure 3: ROC curves by anomaly type. Each curve represents a severity level (darker = higher severity). Coordination failure achieves meaningful AUC at severity  $\geq 0.5$ . Normalized deviance curves are closer to the diagonal, reflecting the fundamental difficulty of per-signal scoring for distributional anomalies.

## 7.2.5 Detection Latency

Coordination failures at high severity are detected within the first 1–2 anomaly signals. Normalized deviance, when detected, requires processing most of the drift sequence before the cumulative shift exceeds the threshold. This is consistent with the nature of each anomaly: coordination failures produce immediate spectral discordances, while normalized deviance is defined by gradual accumulation.

## 7.3 Summary of Controlled Evaluation

The controlled evaluation establishes four claims:

1. **Coordination failures are detectable.**  $TPR = 0.82$  at  $FPR < 0.15$  (12 of 30 high-severity configurations meet threshold).

2. **Normalized deviance is among the hardest to detect.** Average TPR of 0.199 versus 0.661 for coordination failures, confirming that gradual distributional drift resists per-signal scoring.
3. **Detection is robust across mesh sizes.** Maximum TPR spread  $< 0.07$  for all anomaly types.
4. **Detection scales with severity for acute anomalies.** Coordination failure:  $0.117 \rightarrow 0.704$ . Normalized deviance:  $0.022 \rightarrow 0.228$ . Knowledge silo: severity inversion ( $0.166 \rightarrow 0.108$ ).

## 8 Field Deployment

### 8.1 Deployment Context

The architecture was deployed with organizational consent at a mid-size technology company 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 ingested sematectonic traces from GitHub (commits, pull requests, reviews, CI/CD pipeline outputs) and Linear (issues, comments, status changes) across 91 repositories.

### 8.2 Self-Organization

In a 30-day deployment, the mesh ingested 3,520 signals (709 from GitHub, 2,811 from Linear) and self-organized from 3 initial workers to 47, with emergent specializations—external system synchronization, transaction management, entity 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). Processing cost was \$0.012 per signal.

A subsequent 15-day continuous deployment processed 3,138 signals from 64 GitHub repositories and 7 Linear teams. The mesh self-organized to 50 workers and pruned 6 idle workers to a stable topology of 44 at 22% average fullness—the same equilibrium density as the 30-day run, suggesting that fullness stabilizes as a property of the organization rather than of signal volume.

### 8.3 Discovery Findings

Twenty-three findings were surfaced across five categories (RISK, CONTRADICTION, TREND, OVERLAP, DEPENDENCY) with confidence scores ranging from 40% to 95%. The raw correlator produced 3,146 candidate observations; the surfacer collapsed these to 23 deduplicated findings—an 11% survival rate from candidate to promoted finding, evidence of genuine selectivity. This survival rate replicated across subsequent deployments (11% in a 7-day run, 11% in the 15-day run), suggesting it is a stable property of the architecture.

Four findings illustrate distinct theoretical claims:

**Systemic coordination failure.** A systemic external synchronization failure (85% TREND) connected six tickets across different operational entities and three external system vendors—all involving state divergence between the internal system and external platforms, worked independently with no visible coordination. The mesh detected the systemic pattern from sematectonic traces alone.

**$G_D/G_C$  divergence.** A feature-level decision vacuum (90% RISK, 85% CONTRADICTION) identified a multi-party fragmented negotiation where a committed code solution 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.

**Policy communication failure.** 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. Independent validation recovered a complete timeline: policy decision in September, explicit warning 77 days later that the executing team had not been informed, and realized failure 89 days after the warning.

**Action decoupling.** A P1 ticket unassigned for 26 days was being actively cited in conversations to explain related failures (quantified at over \$1,000 per incident) but never discussed as a work item. The ticket’s existence substituted for resolution: the organization’s tracking system itself became a normalization mechanism.

## 8.4 Cross-Source Validation

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 findings and asked to determine whether communication records would confirm, complicate, or refute them. All five investigated findings were confirmed and three were extended. A consistent structural pattern emerged: 5 of 6 tickets 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 but could not propagate through them.

## 8.5 Normalized Deviance Signatures in the Field

Supplementary analysis of internal communication records tested the three normalized deviance signatures against five recurring problem areas. The analysis was conducted independently: the communication records constitute a data source the mesh did not ingest.

*Action decoupling* was the dominant signature, present at high confidence across all five topics. In every case, discussion volume was high (threads of 50 to 166 replies on individual incidents) while resolution evidence was absent or slow. 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 encodes the assumption the failures will continue.

*Linguistic compression* was present at moderate confidence in four of five topics. The most striking example involved a data type confusion producing six-figure financial errors, described as “should be an easy fix.” The containment framing is precisely the signature the theory predicts.

*Framing convergence* was the weakest signature, present at low-to-moderate confidence, limited by the dominance of single voices in most discussion threads.

Three cross-cutting findings extended the theoretical framework. First, taxonomic competence functions as a normalization mechanism: teams developed fluent vocabularies for recurring failures, and the competence at categorizing failures prevents those failures from registering as crises. Second, solo ownership of systemic problems appeared in two topics—multi-system failures assigned to individual engineers without escalation paths. Third, the “easy fix” frame appeared across topics: describing a months-old, high-impact failure as simple suppresses urgency by implying the problem persists by choice rather than by difficulty.

## 9 Discussion

### 9.1 What the Controlled Evaluation Shows

The controlled evaluation establishes that the stigmergic mesh architecture can detect organizational anomalies in synthetic signal streams, with performance that varies systematically by anomaly type in ways predicted by theory. Coordination failures, which produce acute contradictory signals, are detected reliably at moderate to high severity. Normalized deviance, which produces only gradual distributional shifts, is detected at much lower rates—confirming that it is the hardest target for any per-signal scoring approach.

The severity inversion in knowledge silo detection is the evaluation’s most theoretically significant result. It is not a failure of the detection architecture but a structural consequence of the phenomenon being detected. Any correlation-based method faces this limit: the mechanism that constitutes the anomaly also prevents the cross-boundary evidence that detection requires. This result generalizes: organizational pathologies that operate by severing information channels will be systematically hardest to detect through those same channels.

### 9.2 What the Controlled Evaluation Does Not Show

The evaluation uses synthetic anomalies with structured signatures. Real organizational anomalies are messier, more ambiguous, and may not conform to these term-pattern assumptions. The TPR values are upper bounds on real-world detection performance. No baseline comparison against traditional anomaly detection methods (isolation forests, autoencoders, statistical process control) is included. The evaluation establishes absolute performance, not relative performance.

### 9.3 Field Deployment as Complementary Evidence

The field deployment provides qualitative evidence that the controlled evaluation cannot: that the mesh’s self-organizing topology recovers organizational structure from real signal streams, that the emergent specializations correspond to meaningful functional domains, and that the scoring function discriminates between findings of different significance in practice.

The two evidence sources are complementary. The controlled evaluation provides statistical power (22,500 runs with known ground truth) but limited ecological validity. The field deployment provides ecological validity (real organizational signals with confirmed findings) but limited statistical power (single organization, no control group). Neither alone is sufficient. Together, they triangulate: the architecture can detect structured anomalies at measured rates *and* produces meaningful findings in production deployment.

### 9.4 Computational Formalization of External Observation

The architecture’s theoretical contribution is the computational formalization of external observation for organizations. [Luhmann \[1995\]](#) establishes that organizations cannot observe their own blind spots from within. The mesh provides 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: 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.

The mesh satisfies the three Mirror properties identified in companion work [McEntire, 2026a]: *insulation* (its operation does not depend on organizational approval of its findings), *access* (it observes work artifacts rather than reports), and *authority* (artifact generation bypasses the reporting chain). 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.

## 9.5 Immunity to Dysmemic Pressure

The mesh is immune to dysmemic pressure because it reads sematectonic traces, which are costly signals that cannot be degraded without degrading the work 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.

This defense is not absolute. Commits can be restructured to hide coupling, activity can be inflated, logging can be selectively configured. The mesh’s resilience depends on cross-source consistency: manipulating one trace source while leaving correlated sources unmanipulated creates spectral discordances the mesh can detect. The attack surface is proportional to the number of independent trace sources ingested, which is why source diversity is an architectural priority.

## 10 Limitations

We identify the following limitations, presented honestly rather than defensively.

**Single-organization field deployment.** The field deployment reflects a single mid-size technology company. The architecture’s formal properties—ART stability, spectral detection, competitive routing—are organization-independent, but empirical validation across diverse organizational structures remains future work. The 11% survival rate (candidate to promoted finding) replicated across three deployments at this organization; whether it holds elsewhere is unknown.

**Knowledge silo severity inversion.** As demonstrated in Section 7.2.2, complete information silos are structurally invisible to correlation-based detection. This is a theoretical limit, not an engineering one. It means the architecture will systematically underdetect the most severe instances of information isolation—precisely the cases that matter most.

**No longitudinal decay data.** The theory predicts that surfaced findings decay from Discovery into Normalized Deviance through a specific trajectory (Section 6). Testing this prediction requires tracking findings over months to observe whether acknowledged but unresolved findings lose organizational salience. The deployment windows (7, 15, and 30 days) do not provide sufficient temporal depth.

**Synthetic evaluation as upper bound.** The controlled evaluation uses structured anomalies with known signatures. Real organizational anomalies are less structured and may not produce the term-pattern regularities the evaluation exploits. The TPR values reported are likely upper bounds on real-world detection performance.

**No baseline comparison.** The controlled evaluation establishes absolute detection performance but does not compare against traditional anomaly detection methods. We cannot claim the mesh outperforms alternatives—only that it detects anomalies at measurable rates in a domain where most tools cannot operate (continuous multi-source organizational signal streams without predefined anomaly categories).

**Normalized deviance detection remains hard.** Average TPR of 0.199 means the mesh misses approximately 80% of normalized deviance signals. While this is expected theoretically (gradual drift by definition resists per-signal scoring), it is a genuine limitation. Dedicated higher-order detectors—tracking variance of language compression levels over time, autocorrelation of hedging frequency, spectral properties of assessment diversity—may improve performance but remain unvalidated.

**Topology non-persistence.** The mesh’s semantic state (vocabulary, affinities, communication graph) persists across sessions, but topological state (worker count, specializations, thresholds) does not. A mesh that self-organized to 44 specialized workers restarts as 3 generic workers. ART guarantees stable categories will re-form, but category formation is path-dependent.

**Incentive structures.** The mesh cannot close the loop autonomously. If the organization’s reward system punishes the engineer who acts on a surfaced finding, the finding will decay regardless of its environmental presence. The mesh can make findings visible; it cannot alter the payoff structure that determines whether visibility leads to action.

## 11 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 coordination fractures, and the normalized deviations that no one knows to look for because the organization’s own perceptual apparatus was not designed to see them.

This paper has presented a stigmergic mesh architecture that addresses this condition through three principles: continuous ingestion of work artifacts rather than reports, self-organizing topological specialization rather than predefined categories, and geometric pattern detection rather than semantic classification. Each principle exists because a specific formal constraint demands it. Work artifacts must replace reports because the Data Processing Inequality guarantees that reports carry strictly less information. Competitive routing must replace centralized assignment because Crawford-Sobel proves that communication channels with incentive divergence collapse to babbling above Elliott’s 25-participant threshold. Match-based learning must replace gradient-based approaches because only ART’s stability proof holds in non-stationary streams. Spectral detection must replace semantic classification because second-order ignorance means the relevant categories do not yet exist. Topological scoring must replace sentiment-based assessment because any self-report-based mechanism is gameable by the dysmemic forces the system is designed to detect.

The paper’s central contribution is the formalization of normalized deviance as a distinct detection target. Three measurable signatures—eigenvalue collapse, linguistic compression, and action decoupling—transform Vaughan’s qualitative account into a quantitative detection framework. The controlled evaluation demonstrates both the feasibility and the structural limits of this detection: coordination failures are reliably detected (TPR 0.82), while normalized deviance remains the hardest target (TPR 0.199), confirming that gradual drift resists per-signal scoring. The knowledge silo severity inversion—where deeper silos produce lower detection

rates—identifies a fundamental limit of correlation-based methods that generalizes beyond this architecture.

Field deployment validates the architecture’s self-organizing properties and demonstrates cross-source detection in production. The mesh’s emergent topology recovers organizational structure without being given it. The scoring function discriminates meaningfully. Cross-source validation confirms that the mesh detects patterns that exist in the organization but cannot propagate through its own communication channels.

The most important result is the honest identification of what remains hard. Normalized deviance—the pattern that destroyed *Challenger* and *Columbia*—is detected at the lowest rate of any anomaly type. This identifies the critical research direction: developing detectors that track distributional properties of organizational language over time windows longer than the mesh’s current rolling baseline, and that can detect the *absence* of expected patterns rather than the *presence* of anomalous ones.

The theory makes specific, falsifiable predictions. The code is built. The traces are flowing. The predictions are being tested.

The stigmergic mesh architecture is the subject of USPTO Provisional Patent Application #63/981,369 [McEntire, 2026b].

## A Implementation Details

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 four-agent tiering (indexer, surfacer, correlator, profiler) uses LLM-backed reasoning with mechanical fallbacks. If an LLM call fails, times out, or produces structurally invalid output, the system falls back to deterministic heuristics. Insight agents never block the routing pipeline: signals continue flowing regardless of whether higher-order analysis succeeds.

Entity resolution maps heterogeneous identifiers in sematectonic traces 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, Linear, Slack) enrich profiles with platform-specific handles. Resolution proceeds through progressive normalization: exact match, handle-prefix stripping, CamelCase decomposition, username pattern matching, and fuzzy matching. A PII sensitivity layer classifies fields by exposure level and redacts at output boundaries. The system achieved 95% resolution rate (70 canonical identities, 4 unresolved—all automation accounts), progressing from 54% with code-repository-only metadata to 83% with communication channel data to 95% with multi-provider enrichment.

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 insight agents require LLM calls, and those calls operate on pre-structured observation contexts. Processing cost was \$0.012–\$0.016 per signal in deployment. Batch deduplication in the correlator eliminated 87–89% of redundant LLM calls.

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 traces without anyone having predicted that category would be needed.

## B Self-Observing Coherence

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

The cadence monitor discretizes observation history into day-of-week bins with hierarchical fallback: day-specific baselines, weekday/weekend baselines, and all-time baselines. Detection uses an asymmetry score:

$$A(H, v) = \sqrt{\frac{|H_{\text{baseline}} - H_{\text{current}}|}{H_{\text{baseline}}}} \cdot \sqrt{\frac{v_{\text{current}}}{v_{\text{baseline}}}} \quad (5)$$

where  $H$  is source distribution entropy (Shannon) and  $v$  is signal volume. When entropy drops while volume holds—asymmetric compression, where one source dies while others continue— $A$  amplifies. When both drop—symmetric quiet—the volume term suppresses.

In three-source deployment, the cadence monitor learned distinct weekday and weekend entropy profiles within the first observation cycle, correctly distinguishing between symmetric quiet (reduced weekend activity across all sources) and asymmetric quiet (a single engineer’s weekend code activity).

## C Probabilistic Emission Model

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 an attention model, and let  $I(f) = S(f)$ . Let  $C_r$  denote the cost of redundancy. The emission criterion surfaces finding  $f$  to actor  $a$  when:

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

The probability  $P_k$  is estimated via noisy-OR over exposure events, 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. 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.

Organizational dysfunction manifests as clusters of individually sub-threshold findings. Compound risk scoring groups related findings by single-linkage clustering with score  $S_c = \max(S_i) + \delta \sum_{j \neq i^*} S_j$  (where  $\delta = 0.3$ ), enabling escalation when no individual finding exceeds the surfacing threshold.

## References

- Phillip G. Armour. The five orders of ignorance. *Communications of the ACM*, 43(10):17–20, 2000.
- W. Ross Ashby. *An Introduction to Cybernetics*. Chapman & Hall, 1956.
- Michael Beer and Nitin Nohria. Cracking the code of change. *Harvard Business Review*, 78(3): 133–141, 2000.
- Stafford Beer. *Brain of the Firm*. Allen Lane / The Penguin Press, 1972.
- Ronald S. Burt. Structural holes and good ideas. *American Journal of Sociology*, 110(2):349–399, 2004.

- Gail A. Carpenter and Stephen Grossberg. A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing*, 37(1):54–115, 1987.
- Gail A. Carpenter, Stephen Grossberg, and David B. Rosen. Fuzzy ART: Fast stable learning and categorization of analog patterns by an adaptive resonance system. *Neural Networks*, 4(6):759–771, 1991.
- Marcelo Cataldo, James D. Herbsleb, and Kathleen M. Carley. 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)*, pages 2–11, 2008.
- Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM Computing Surveys*, 41(3):1–58, 2009.
- Columbia Accident Investigation Board. Report of the Columbia accident investigation board, volume 1. NASA, 2003.
- Roger C. Conant and W. Ross Ashby. Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2):89–97, 1970.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 2nd edition, 2006.
- Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, 1982.
- Claire Donnat, Marinka Zitnik, David Hallac, and Jure Leskovec. Learning structural node embeddings via diffusion wavelets. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1320–1329, 2018.
- Mark Elliott. *Stigmergic Collaboration: A Theoretical Framework for Mass Collaboration*. PhD thesis, University of Melbourne, 2007.
- Ernesto Estrada and Naomichi Hatano. Communicability in complex networks. *Physical Review E*, 77(3):036111, 2008.
- Bent Flyvbjerg, Alexander Budzier, Jeong Soo Lee, Mark Keil, Daniel Lunn, and Dirk W. Bester. The empirical reality of IT project cost overruns: Discovering a power-law distribution. *Journal of Management Information Systems*, 39(3):607–639, 2022.
- Mark S. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78(6):1360–1380, 1973.
- Pierre-Paul Grassé. La reconstruction du nid et les coordinations interindividuelles chez *Belligoseria natalensis* et *Cubitermes sp.* *Insectes Sociaux*, 6(1):41–80, 1959.
- Stephen Grossberg. Adaptive pattern classification and universal recoding: I. Parallel development and coding of neural feature detectors. *Biological Cybernetics*, 23(3):121–134, 1976.
- Francis Heylighen. Stigmergy as a universal coordination mechanism: I. Definition and components. *Cognitive Systems Research*, 38:4–13, 2016.
- Stuart A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, 1993.

- F. William Lawvere. Diagonal arguments and cartesian closed categories. In *Category Theory, Homology Theory and their Applications II*, volume 92 of *Lecture Notes in Mathematics*, pages 134–145. Springer, 1969.
- José María Liberti and Atif R. Mian. Estimating the effect of hierarchies on information use. *Review of Financial Studies*, 22(10):4057–4090, 2009.
- Niklas Luhmann. *Social Systems*. Stanford University Press, 1995.
- Jeremy McEntire. Dymemic pressure: Selection dynamics in organizational information environments. *arXiv preprint arXiv:2512.14716*, 2025. Also available at SSRN: <https://ssrn.com/abstract=6015814>.
- Jeremy McEntire. The mirror: Structural requirements for organizational self-observation. Working paper, 2026a.
- Jeremy McEntire. Stigmergic mesh for organizational signal processing. USPTO Provisional Patent Application #63/981,369, filed February 2026, 2026b.
- Canice Prendergast. A theory of “yes men”. *American Economic Review*, 83(4):757–770, 1993.
- Presidential Commission on the Space Shuttle *Challenger* Accident. Report of the presidential commission on the space shuttle Challenger accident. Washington, DC, 1986.
- Ilya Prigogine. *Self-Organization in Nonequilibrium Systems*. Wiley, 1977.
- Herbert A. Simon. Rational choice and the structure of the environment. *Psychological Review*, 63(2):129–138, 1956.
- Michael Spence. Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374, 1973.
- Standish Group. CHAOS report 2015. The Standish Group International, 2015.
- Damian A. Tamburri, Fabio Palomba, and Rick Kazman. Exploring community smells in open-source: An automated approach. *IEEE Transactions on Software Engineering*, 47(3):630–652, 2021.
- Jianheng Tang, Jiajin Li, Ziqi Gao, and Jia Li. Rethinking graph neural networks for anomaly detection. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, pages 21076–21089, 2022.
- Guy Theraulaz and Eric Bonabeau. A brief history of stigmergy. *Artificial Life*, 5(2):97–116, 1999.
- Diane Vaughan. *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. University of Chicago Press, 1996.
- Timo O. Vuori and Quy N. Huy. Distributed attention and shared emotions in the innovation process: How Nokia lost the smartphone battle. *Administrative Science Quarterly*, 61(1):9–51, 2016.
- Karl E. Weick and Kathleen M. Sutcliffe. *Managing the Unexpected: Resilient Performance in an Age of Uncertainty*. Jossey-Bass, 2nd edition, 2007.
- Edward O. Wilson. *The Insect Societies*. Harvard University Press, 1971.